

Article

Leveraging Possibilistic Beliefs in Unrestricted Combinatorial Auctions

Jing Chen ^{1,*} and Silvio Micali ²

¹ Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, USA

² Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA 02139, USA; silvio@csail.mit.edu

* Correspondence: jingchen@cs.stonybrook.edu; Tel.: +1-631-632-1827

Academic Editor: Paul Weirich

Received: 9 August 2016; Accepted: 17 October 2016; Published: 26 October 2016

Abstract: In unrestricted combinatorial auctions, we put forward a mechanism that guarantees a meaningful revenue benchmark based on the possibilistic beliefs that the players have about each other's valuations. In essence, the mechanism guarantees, within a factor of two, the maximum revenue that the "best informed player" would be sure to obtain if he/she were to sell the goods to his/her opponents via take-it-or-leave-it offers. Our mechanism is probabilistic and of an extensive form. It relies on a new solution concept, for analyzing extensive-form games of incomplete information, which assumes only mutual belief of rationality. Moreover, our mechanism enjoys several novel properties with respect to privacy, computation and collusion.

Keywords: possibilistic beliefs; unrestricted combinatorial auctions; mutual belief of rationality; incomplete information; extensive-form games; distinguishable dominance

1. Introduction

In this paper, we study the problem of generating revenue in unrestricted combinatorial auctions, solely relying on the players' possibilistic beliefs about each others' valuations. Let us explain.

In a combinatorial auction, there are multiple indivisible goods for sale and multiple players who are interested in buying. A valuation of a player is a function specifying a non-negative value for each subset of the goods. Many constraints on the players' valuations have been considered in the literature for combinatorial auctions¹. We instead focus on combinatorial auctions that are unrestricted. That is, in our auctions, a player's value for one subset of the goods may be totally unrelated to his/her value for another subset and to another player's value for any subset. This is the most general class of auctions. It is well known that, for such auctions, the famous Vickrey-Clarke-Groves (VCG) mechanism [4–6] maximizes social welfare in dominant strategies, but offers no guarantee about the amount of revenue it generates. In fact, for unrestricted combinatorial auctions, no known mechanism guarantees any significant revenue benchmark in settings of incomplete information².

In our setting, the seller has no information about the players' valuations, and each player knows his/her own valuation, but not necessarily the valuations of his/her opponents. Our players, however, have beliefs about the valuations of their opponents. Typically, beliefs are modeled as probability distributions: for instance, it is often assumed that the valuation profile, θ , is drawn from a common prior. Our setting is instead non-Bayesian: the players' beliefs are possibilistic and can

¹ Such as monotonicity, single-mindedness and additivity [1–3].

² In complete information settings (where the players have common knowledge about their valuations), assuming common knowledge of rationality, [7–9] have designed mechanisms that guarantee revenue arbitrarily close to the maximum social welfare.

be arbitrary. That is, a player i 's belief consists of a set of valuation profiles, \mathcal{B}_i , to which he/she believes θ belongs. We impose no restriction on \mathcal{B}_i except that, since i knows his/her own valuation, for every profile $v \in \mathcal{B}_i$, we have $v_i = \theta_i$. In a sense, therefore, such possibilistic beliefs are not assumed to exist, but always exist. For instance, if a player i has no information about his/her opponents, then \mathcal{B}_i consists of the set of all valuation profiles v , such that $v_i = \theta_i$; if i has complete information about his/her opponents, then $\mathcal{B}_i = \{\theta\}$; and if θ is indeed drawn from a common prior D , then \mathcal{B}_i consists of the support of D conditioned on θ_i .

Possibilistic beliefs are much less structured than Bayesian ones. Therefore, it should be harder for an auction mechanism to generate revenue solely based on the players' possibilistic beliefs. Yet, in single-good auctions, the authors of [10] have constructed a mechanism that guarantees revenue at least as high as the second-highest valuation and, sometimes, much higher. In this paper, for unrestricted combinatorial auctions, we construct a mechanism that guarantees, within a factor of two, another interesting revenue benchmark, BB , solely based on the players' possibilistic beliefs.

The benchmark BB is formally defined in Section 3, following the framework put forward by Harsanyi [11] and Aumann [12]. However, it can be intuitively described as follows. Let BB_i (for "best belief") be the maximum social welfare player i can guarantee, based on his/her beliefs, by assigning the goods to his/her opponents. Then, $BB = \max_i BB_i$, and the revenue guaranteed by our main mechanism is virtually $BB/2$. Notice that each BB_i does not depend on θ_i at all, a property that, as we shall see, gives our mechanism some advantage in protecting the players' privacy.

To ease the discussion of our main mechanism, in Section 4, we construct a first mechanism, of normal form, that guarantees revenue virtually equal to $BB/2$ under two-step elimination of weakly-dominated strategies. The analysis of our first mechanism is very intuitive. However, elimination of weakly-dominated strategies is order-dependent and does not yet have a well-understood epistemic characterization. Moreover, our first mechanism suffers from two problems shared by most normal-form mechanisms. Namely, (1) it reveals all players' true valuations, and (2) it requires an amount of communication that is exponential in the number of goods. Both problems may not be an issue from a pure game-theoretic point of view³, but are quite serious in several realistic applications, where privacy and communication are, together with collusion and computational complexity, legitimate concerns [14].

Our main mechanism, the best-belief mechanism, significantly decreases the magnitude of the above problems. This second mechanism is designed and analyzed in Section 6 and is of extensive form. In order to analyze it in settings where the players have possibilistic beliefs, we propose a new and compelling solution concept that only assumes mutual belief of rationality, where the notion of rationality is the one considered by Aumann [15].

The Resiliency of the Best-Belief Mechanism.

Besides guaranteeing revenue virtually equal to $BB/2$ under a strong solution concept, the best-belief mechanism enjoys several novel properties with respect to privacy, computation, communication and collusion.

1. Privacy: People value privacy. Thus, "by definition", a privacy-valuing player i de facto receives some "negative utility" if, in an auction, he/she reveals his/her true valuation θ_i in its entirety, but does not win any goods. Even if he/she wins some goods, his/her traditional utility (namely, his/her value for the goods he/she receives minus the price he/she pays) should be discounted by the loss he/she suffers from having revealed θ_i . One advantage of our best-belief mechanism is that it elicits little information from a player, which presumably diminishes the likelihood that privacy may substantially distort a player's incentives.

³ Indeed, the revelation principle [13] explicitly asks the players to directly reveal all of their private information.

2. **Computation:** Typically, unrestricted combinatorial auctions require the evaluation of complex functions, such as the maximum social welfare. In principle, one may resort to approximating such functions, but approximation may distort incentives⁴. By contrast, our mechanism delegates all difficult computation to the players and ensures that the use of approximation is properly aligned with their incentives.
3. **Communication:** By eliciting little information from the players, our mechanism also has low communication complexity; quadratic in the number of players and the number of goods.
4. **Collusion:** Collusion can totally disrupt many mechanisms. In particular, the efficiency of the VCG mechanism can be destroyed by just two collusive players [14]. By contrast, collusion can somewhat degrade the performance of our mechanism, but not totally disrupt it, unless all players are collusive. As long as collusive players are also rational, at least in a very mild sense, the revenue guaranteed by our mechanism is at least half of that obtainable by “the best informed independent player”.

For a detailed discussion about these properties of our mechanism, see Section 6.2.

2. Related Work

Generating revenue is one of the most important objectives in auction design; see [16,17] for thorough introductions about this area. Following the seminal result of [13], there has been a huge literature on Bayesian auctions [18]. Since we do not assume the existence of a common prior and we focus on the players’ possibilistic rather than probabilistic beliefs, our study is different from Bayesian auctions. Spectrum auctions have been widely studied both in theory and in practice, and several interesting auction forms have been proposed recently; see, e.g., [19–23]. Most existing works consider auctions of restricted forms, such as auctions with multiple identical goods and single-demand valuations [24], valuations with free disposal [25], auctions with additive valuations [26], auctions with unlimited supply [27], etc. Revenue in unrestricted combinatorial auctions has been considered by [28], which generalizes the second-price revenue benchmark to such auctions and provides a mechanism guaranteeing a logarithmic fraction of their benchmark in dominant strategies.

The solution concept developed in this paper refines the notion of implementation in undominated strategies [29] and considers a two-round elimination of dominated strategies. In particular, we extend the notion of distinguishably-dominated strategies [30] from extensive-form games of complete information to extensive-form games of incomplete information and possibilistic beliefs. As shown in [30], iterated elimination of distinguishably-dominated strategies is order independent with respect to histories and characterizes extensive-form rationalizability [31,32]. In [10,33], elimination of strictly-dominated strategies has been extended to deal with possibilistic beliefs, but only for normal-form games. Moreover, [34] leverages the players’ beliefs for increasing the sum of social welfare and revenue in unrestricted combinatorial auctions.

Preserving the privacy of the players’ valuations, or types in general, in the execution of a mechanism has been studied by [35]. The authors present a general method using some elementary physical equipment (i.e., envelopes and an envelope randomizer) so as to execute any given normal-form mechanism, without trusting any one and without revealing any information about the players’ true types, beyond what is unavoidably revealed in the final outcome. An alternative way to protect the privacy of the players that has often been considered for auctions is to use encryption and zero-knowledge proofs. In particular, the authors of [36] make efficient use of cryptography to implement single-good auctions so that, after learning all bids, an untrusted auctioneer can prove who won the good and what price he/she should pay, without having any player learn any information

⁴ For instance, the outcome function of the VCG mechanism is NP-hard to compute even when each player only values a single subset of the goods for \$1 and all other subsets for \$0. Moreover, if one replaces this outcome function with an approximation, then VCG would no longer be dominant-strategy truthful.

about the bid of another player. Moreover, in differential privacy [37], the mechanisms or databases inject noise to the final outcome to preserve the participants' privacy. By contrast, our main mechanism does not rely on envelopes or any other form of physical equipment, nor on cryptography or noise. It preserves the players' privacy because, despite the fact that all actions are public, a player is asked to reveal little information about herself/himself.

Strong notions of collusion-resilient implementation have been studied in the literature, such as coalition incentive compatibility [38] and bribe proofness [39]. However, the authors prove that many social choice functions cannot be implemented under these solution concepts. The collusive dominant-strategy truthful implementation is defined in [40], together with a mechanism maximizing social welfare in multi-unit auctions under this notion. Other forms of collusion resiliency have also been investigated, in particular by [41–46]. Their mechanisms, however, are not applicable to unrestricted combinatorial auctions in non-Bayesian settings. Moreover, the collusion models there assume various restrictions (e.g., collusive players cannot make side-payments to one another or enter binding agreements, there is a single coalition, no coalition can have more than a given number of players, etc.). By contrast, in unrestricted combinatorial auctions, our main mechanism does not assume any such restrictions. The resiliency of our mechanism is similar to that of [28], where the guaranteed revenue benchmark is defined only on independent players' valuations when collusion exists.

3. Preliminaries and the Best-Belief Revenue Benchmark

A combinatorial auction context is specified by a triple (n, m, θ) : the set of players is $\{1, \dots, n\}$; the set of goods is $\{1, \dots, m\}$; and the true valuation profile is θ . Adopting a discrete perspective, we assume that a player's value for a set of goods is always an integer. Thus, each θ_i , the true valuation of i , is a function from the powerset $2^{\{1, \dots, m\}}$ to the set of non-negative integers \mathbb{Z}^+ , with $\theta_i(\emptyset) = 0$. The set of possible valuations of i , Θ_i , consists of all such functions, and $\Theta = \Theta_1 \times \dots \times \Theta_n$. After constructing and analyzing our mechanisms, we will discuss the scenarios where values are real numbers.

An outcome of a combinatorial auction is a pair of profiles (A, P) . Here, A is the allocation, with $A_i \subseteq \{1, \dots, m\}$ being the set of goods each player i gets, and $A_i \cap A_j = \emptyset$ for each player $j \neq i$; and P is the price profile, with $P_i \in \mathbb{R}$ denoting how much each player i pays; if $P_i < 0$, then i receives $-P_i$ from the seller. The set of all possible outcomes is denoted by Ω .

The utility function of i , u_i , maps each valuation $t_i \in \Theta_i$ and each outcome $\omega = (A, P)$ to a real: $u_i(t_i, \omega) = t_i(A_i) - P_i$. The social welfare of ω is $SW(\omega) \triangleq \sum_i \theta_i(A_i)$, and the revenue of ω is $REV(\omega) \triangleq \sum_i P_i$. If ω is a probability distribution over outcomes, then $u_i(t_i, \omega)$, $SW(\omega)$ and $REV(\omega)$ denote the corresponding expectations.

Definition 1. An augmented combinatorial auction context is a four-tuple $(n, m, \theta, \mathcal{B})$, where (n, m, θ) is a combinatorial auction context and \mathcal{B} is the belief profile: for each player i , \mathcal{B}_i , the belief of i , is a set of valuation profiles, such that $t_i = \theta_i$ for all $t \in \mathcal{B}_i$.

In an augmented combinatorial auction context, \mathcal{B}_i is the set of candidate valuation profiles in i 's mind for θ . The restriction that $t_i = \theta_i$ for all $t \in \mathcal{B}_i$ corresponds to the fact that player i knows his/her own valuation. Player i 's belief is correct if $\theta \in \mathcal{B}_i$ and incorrect otherwise. As we shall see, our result holds whether or not the players' beliefs are correct. From now on, since we do not consider any other type of auctions, we use the terms "augmented" and "combinatorial" for emphasis only.

A revenue benchmark f is a function that maps each auction context C to a real number $f(C)$, denoting the amount of revenue that is desired under this context.

Definition 2. The best-belief revenue benchmark, BB , is defined as follows. For each auction context $C = (n, m, \theta, \mathcal{B})$,

$$BB(C) \triangleq \max_i BB_i,$$

where for each player i ,

$$BB_i \triangleq \max_{(A,P) \in \Omega: A_i = \emptyset \text{ and } P_j \leq t_j(A_j) \forall j \neq i, t \in \mathcal{B}_i} REV(A, P).$$

Note that BB_i represents the maximum revenue that player i would be sure to obtain if he/she were to sell the goods to his/her opponents via take-it-or-leave-it offers, which is also the maximum social welfare player i can guarantee, based on his/her beliefs, by assigning the goods to his/her opponents. As an example, consider a combinatorial auction with two items and three players. Player 1 only wants Item 1, and $\theta_1(\{1\}) = 100$; Player 2 only wants Item 2, and $\theta_2(\{2\}) = 100$; and Player 3 only wants the two items together, and $\theta_3(\{1,2\}) = 50$. All the unspecified values are zero. Moreover, Player 1 believes that Player 2's value for Item 2 is at least 25, and Player 3's value for the two items together is at least 10: that is, $\mathcal{B}_1 = \{v \mid v_1 = \theta_1, v_2(\{2\}) \geq 25, v_3(\{1,2\}) \geq 10\}$. Accordingly, $BB_1 = 25$: the best Player 1 can do in selling to others is to offer Item 2 to Player 2 at price 25. Furthermore, $\mathcal{B}_2 = \{v \mid v_2 = \theta_2, v_1(\{1\}) \geq 80, v_3(\{2\}) \geq 20\}$, which implies $BB_2 = 100$, achieved by offering Item 1 to Player 1 at price 80 and Item 2 to Player 3 at price 20. Finally, $\mathcal{B}_3 = \{v \mid v_3 = \theta_3, v_1(\{1\}) \geq 80, v_2(\{2\}) \geq 70\}$, which implies $BB_3 = 150$, achieved by offering Item 1 to Player 1 at price 80 and Item 2 to Player 2 at price 70. Therefore, $BB = 150$ in this example. Note that Player 1's and Player 3's beliefs are correct, but Player 2's beliefs are incorrect because $\theta_3(\{2\}) = 0$.

Furthermore, note that, if there is really a common prior from which the players' valuations are drawn, then the players' possibilistic beliefs consist of the support of the distribution. In this case, it is expected that the optimal Bayesian mechanism generates more revenue than the best-belief benchmark. However, this is a totally different ball game, because Bayesian mechanisms assume that the seller has much knowledge about the players. Besides, little is known in the literature about the structure of the optimal Bayesian mechanism for unrestricted combinatorial auctions or even a good approximation to it.

Finally, the best-belief benchmark is measured based on the players' beliefs about each other, not on their true valuations. If the players all know nothing about each other and believe that the others' values can be anything from close to zero to close to infinity (or a huge finite number), then the benchmark is low. The power of the benchmark comes from the class of contexts where the players know each other well (e.g., as long-time competitors in the same market) and can effectively narrow down the range of the others' values. In this case, our mechanism generates good revenue without assuming a common prior.

4. A Normal-Form Mechanism

As a warm up, in this section, we construct a normal-form mechanism that implements the best-belief revenue benchmark within a factor of two, under two-step elimination of weakly-dominated strategies. Indeed, weakly-dominant/dominated strategies have been widely used in analyzing combinatorial auctions where the players only report their valuations: that is, it is weakly dominant for each player to report his/her true valuation. When each player reports both his/her own valuation and his/her beliefs about the other players, it is intuitive that a player i first reasons about what the other players report for their valuations and then reasons about what to report for his/her beliefs about them given their reported valuations: that is, an iterated elimination of dominated strategies. However, in our mechanism, there is no need to go all the way to the end of the iterated procedure, and two steps are sufficient.

Roughly speaking, all players first simultaneously remove all of their weakly-dominated strategies; and then, each player further removes all of his/her strategies that now become weakly dominated, based on all players' surviving strategies. However, care must be taken when defining this solution concept in our setting. Indeed, since a player does not know the other players' true valuations, he/she cannot compute their strategies surviving the first round of elimination, which are needed for

him to carry out his/her second round of elimination. To be “on the safe side”, we require that the players eliminate their strategies conservatively: that is, a player eliminates a strategy in the second round only if it is dominated by the same alternative strategy with respect to all valuation profiles that he/she believes to be possible. This notion of elimination is the same as the one used by Aumann in [15], except that in the latter, it is strict instead of weak domination. In [33], the authors provide an epistemic characterization for iterated elimination based on the notion of [15].

More precisely, given a normal-form auction mechanism M , let S_i be the set of strategies of each player i and $S = S_1 \times \dots \times S_n$. For any strategy profile s , $M(s)$ is the outcome when each player i uses strategy s_i . If $T = T_i \times T_{-i}$ is a subset of strategy profiles, $t_i \in \Theta_i$, $s_i \in T_i$, and $\sigma_i \in \Delta(T_i)$ ⁵, then we say that s_i is weakly dominated by σ_i with respect to t_i and T , in symbols $s_i \leq_T^{t_i} \sigma_i$, if:

- $u_i(t_i, M(s_i, s_{-i})) \leq u_i(t_i, M(\sigma_i, s_{-i}))$ for all $s_{-i} \in T_{-i}$ and
- $u_i(t_i, M(s_i, s_{-i})) < u_i(t_i, M(\sigma_i, s_{-i}))$ for some $s_{-i} \in T_{-i}$.

That is, s_i is weakly dominated by σ_i when the valuation of player i is t_i and the set of strategy sub-profiles of the other players is T_{-i} . The set of strategies in T_i that are not weakly dominated with respect to t_i and T is denoted by $U_i(t_i, T)$. For simplicity, we use U_i to denote $U_i(\theta_i, S)$, the set of undominated strategies of player i .

Definition 3. Given an auction context $C = (n, m, \theta, \mathcal{B})$ and a mechanism M , the set of conservatively weakly-rational strategies of player i is:

$$C_i \triangleq U_i \setminus \{s_i : \exists \sigma_i \in \Delta(U_i) \text{ s.t. } \forall t \in \mathcal{B}_i, s_i \leq_{U(t)}^{\theta_i} \sigma_i\},$$

where $U(t) \triangleq \times_j U_j(t_j, S)$ for any $t \in \Theta$. The set of conservatively weakly-rational strategy profiles is $\mathcal{C} = C_1 \times \dots \times C_n$.

Mechanism M conservatively weakly implements a revenue benchmark f if, for any auction context C and any strategy profile $s \in \mathcal{C}$,

$$REV(M(s)) \geq f(C).$$

Now, we provide and analyze our normal-form mechanism M_{Normal} . Intuitively, the players compete for the right to sell to others, and the mechanism generates revenue by delegating this right to the player who offers the most revenue. Besides the number of players n and the number of goods m , the mechanism takes as input a constant $\epsilon \in (0, 1]$. The players act only in Step 1, and Steps a through f are “steps taken by the mechanism”. The expression “ $X := x$ ” sets or resets variable X to value x . Moreover, $[m] = \{1, 2, \dots, m\}$.

Mechanism M_{Normal} :

- 1: Each player i , publicly and simultaneously with the other players, announces:
 - a valuation v_i and
 - an outcome $\omega^i = (\alpha^i, \pi^i)$, such that: $\alpha_j^i = \emptyset$, and for each player j , π_j^i is zero whenever $\alpha_j^i = \emptyset$; and is a positive integer otherwise.

After the players simultaneously execute Step 1, the mechanism chooses the outcome (A, P) by means of the following six steps.

- a: Set $A_i := \emptyset$, and $P_i := 0$ for each player i .

⁵ As usual, for a set T , $\Delta(T)$ is the set of probability distributions over T .

- b: Set $R_i := REV(\omega^i)$ for each player i , and $w := \operatorname{argmax}_i R_i$ with ties broken lexicographically.
- c: Publicly flip a fair coin and denote the result by r .
- d: If $r = Heads$, then $A_w := \operatorname{argmax}_{a \subseteq [m]} v_w(a)$, with ties broken lexicographically, and halt.
- e: (Note that $r = Tails$ when this step is reached.)

For each player i , such that $\alpha_i^w \neq \emptyset$:

- If $v_i(\alpha_i^w) < \pi_i^w$, then $P_w := P_w + \pi_i^w$.
- Otherwise, $A_i := \alpha_i^w$ and $P_i := \pi_i^w - \frac{\epsilon}{n}$.

- f: For each player i , $P_i := P_i - \delta_i$ with $\delta_i = \frac{\epsilon}{n} \cdot \frac{R_i}{1+R_i}$.

The final outcome is (A, P) .

In the analysis, we refer to player w as the winner and each δ_i as player i 's reward. Furthermore, given a context $(n, m, \theta, \mathcal{B})$ and an outcome ω , for succinctness, we use $u_i(\omega)$ instead of $u_i(\theta_i, \omega)$ for player i 's utility under ω . We have the following.

Theorem 1. For any context $(n, m, \theta, \mathcal{B})$ and constant $\epsilon \in (0, 1]$, mechanism M_{Normal} conservatively weakly implements the revenue benchmark $\frac{BB}{2} - \epsilon$.

As we will see in the proof of Theorem 1, the mechanism incentivizes each player i to report his/her true valuation and an outcome whose revenue is at least BB_i . In particular, the latter is achieved by the fair coin toss: when $r = Heads$, the winner is given his/her favorite subset of goods for free, which is better than any offer he/she can possibly get if somebody else becomes the winner. Moreover, the rewards are strictly increasing with the revenue of the reported outcomes. Accordingly, the players do not have incentives to underbid; that is, to report an outcome whose revenue is lower than the corresponding BB_i . Thus, the winner's reported outcome has a revenue of at least $\max_i BB_i$. When $r = Tails$, the mechanism tries to sell the goods as suggested by the winner to the other players, as a take-it-or-leave-it offer. If a player accepts the offer, then he/she pays the suggested price; otherwise, this price is charged to the winner as a fine. Accordingly, with probability 1/2 (that is, when $r = Tails$), the mechanism generates revenue $\max_i BB_i$. Formally, we show the following two lemmas.

Lemma 1. For any context $(n, m, \theta, \mathcal{B})$, constant ϵ , player i and strategy $s_i = (v_i, \omega^i)$, if $s_i \in U_i$, then $v_i = \theta_i$.

Proof. Notice that v_i is used in two places in the mechanism: to select player i 's "favorite subset" in Step d when he/she is the winner and to decide whether he/she gets the set α_i^w in Step e when he/she is not the winner. Intuitively, it is i 's best strategy to announce his/her true valuation so as to select his/her "truly favorite subset" and to take the allocated set if and only if its price is less than or equal to his/her true value for it.

More precisely, arbitrarily fix a strategy $s_i = (v_i, \omega^i)$ with $v_i \neq \theta_i$, and let $s'_i = (\theta_i, \omega^i)$. We show that $s_i \leq_S^{\theta_i} s'_i$, where S is the set of all strategy profiles of M_{Normal} . To do so, arbitrarily fix a strategy sub-profile s_{-i} of the other players; let (A, P) be the outcome of $s = (s_i, s_{-i})$, and let (A', P') be the outcome of $s' = (s'_i, s_{-i})$. Since s_i and s'_i announce the same outcome ω^i , i is the winner under s if and only if he/she is the winner under s' . We discuss these two cases separately.

Case 1: i is the winner under both s and s' .

In this case, conditioned on $r = Heads$, we have $A_i = \operatorname{argmax}_{a \subseteq [m]} v_i(a)$, $A'_i = \operatorname{argmax}_{a \subseteq [m]} \theta_i(a)$ and $P_i = P'_i = 0$. Accordingly, $\theta_i(A'_i | r = Heads) \geq \theta_i(A_i | r = Heads)$ and $u_i(A', P' | r = Heads) \geq u_i(A, P | r = Heads)$.

Conditioned on $r = Tails$, we have $A_i = A'_i = \emptyset$ and:

$$P_i = P'_i = \sum_{j: \alpha_j^i \neq \emptyset \text{ and } v_j(\alpha_j^i) < \pi_j^i} \pi_j^i - \delta_i,$$

where $\delta_i = \frac{\epsilon}{n} \cdot \frac{REV(\omega^i)}{1+REV(\omega^i)}$ is player i 's reward under both strategy profiles. Accordingly, $u_i(A', P' | r = Tails) = u_i(A, P | r = Tails)$.

In sum, $u_i(A, P) \leq u_i(A', P')$ in Case 1.

Case 2: i is the winner under neither s nor s' .

In this case, the winner w is the same under both s and s' . Conditioned on $r = Heads$, we have $A_i = A'_i = \emptyset$ and $P_i = P'_i = 0$; thus, $u_i(A, P | r = Heads) = u_i(A', P' | r = Heads)$.

Conditioned on $r = Tails$, if $v_i(\alpha_i^w) < \pi_i^w$ and $\theta_i(\alpha_i^w) < \pi_i^w$, or if both inequalities are reversed, then $(A_i, P_i) = (A'_i, P'_i)$ and $u_i(A, P | r = Tails) = u_i(A', P' | r = Tails)$. Otherwise, if $v_i(\alpha_i^w) < \pi_i^w$ and $\theta_i(\alpha_i^w) \geq \pi_i^w$, then:

$$u_i(A', P' | r = Tails) = \theta_i(\alpha_i^w) - \pi_i^w + \frac{\epsilon}{n} + \delta_i > \delta_i = u_i(A, P | r = Tails), \tag{1}$$

where again δ_i is i 's reward under both strategy profiles. Otherwise, we have $v_i(\alpha_i^w) \geq \pi_i^w$ and $\theta_i(\alpha_i^w) < \pi_i^w$; thus:

$$u_i(A, P | r = Tails) = \theta_i(\alpha_i^w) - \pi_i^w + \frac{\epsilon}{n} + \delta_i \leq -1 + \frac{\epsilon}{n} + \delta_i < \delta_i = u_i(A', P' | r = Tails). \tag{2}$$

In sum, $u_i(A, P) \leq u_i(A', P')$ in Case 2, as well.

It remains to show there exists a strategy sub-profile s_{-i} , such that $u_i(A, P) < u_i(A', P')$, and such an s_{-i} has actually appeared in Case 2 above. Indeed, since $v_i \neq \theta_i$, there exists $a \subseteq [m]$, such that $v_i(a) \neq \theta_i(a)$. When $v_i(a) < \theta_i(a)$, arbitrarily fix a player $j \neq i$, and choose strategy s_j , such that:

$$\alpha_i^j = a, \pi_i^j = \theta_i(a), \text{ and } REV(\omega^j) > \max\{\pi_i^j, REV(\omega^i)\}.$$

Notice that such a strategy exists in S_j : player j can set π_k^j to be arbitrarily high for any player $k \notin \{i, j\}$. Moreover, for any player $k \notin \{i, j\}$, choose s_k to be such that $REV(\omega^k) = 0$. By construction, $w = j$ under both s and s' , $v_i(\alpha_i^w) < \pi_i^w$ and $\theta_i(\alpha_i^w) \geq \pi_i^w$. Following Case 2 above, $u_i(A, P | r = Heads) = u_i(A', P' | r = Heads)$ and $u_i(A, P | r = Tails) < u_i(A', P' | r = Tails)$ by Inequality 1. Thus, $u_i(A, P) < u_i(A', P')$.

When $v_i(a) > \theta_i(a)$, similarly, choose strategy s_j , such that:

$$\alpha_i^j = a, \pi_i^j = v_i(a), \text{ and } REV(\omega^j) > \max\{\pi_i^j, REV(\omega^i)\},$$

and choose strategy s_k the same as above for any $k \notin \{i, j\}$. The analysis again follows from Case 2 above (in particular, Inequality 2); thus, $u_i(A, P) < u_i(A', P')$.

Combining everything together, $s_i \leq_{S'} s'_i$, and Lemma 1 holds. \square

Lemma 2. For any context $(n, m, \theta, \mathcal{B})$, constant ϵ , player i and strategy $s_i = (v_i, \omega^i)$, if $s_i \in \mathcal{C}_i$, then $REV(\omega^i) \geq BB_i$.

Proof. By Lemma 1, we only need to consider strategies, such that $v_i = \theta_i$. Arbitrarily fix a strategy $s_i = (\theta_i, \omega^i) \in U_i$ with $REV(\omega^i) < BB_i$. Consider a strategy $\hat{s}_i = (\theta_i, \hat{\omega}^i)$, such that $\hat{\omega}^i = (\hat{\alpha}^i, \hat{\pi}^i)$ satisfies the following conditions:

$$\hat{\omega}^i \in \underset{(A,P) \in \Omega: A_i = \emptyset \text{ and } P_j \leq t_j(A_j) \ \forall j \neq i, t \in \mathcal{B}_i}{\text{argmax}} REV(A, P)$$

and

$$\hat{\pi}^i_j > 0 \text{ whenever } \hat{\alpha}^i_j \neq \emptyset.$$

Notice that $REV(\hat{\omega}^i) = BB_i > REV(\omega^i)$. We show that for all $t \in \mathcal{B}_i$, $s_i \leq_{U(t)}^{\theta_i} \hat{s}_i$.

To do so, arbitrarily fix a valuation profile $t \in \mathcal{B}_i$ and a strategy sub-profile s_{-i} , such that $s_j \in U_j(t_j, S)$ for each player j . Note that $t_i = \theta_i$ by the definition of \mathcal{B}_i . Moreover, by Lemma 1, each s_j is of the form (t_j, ω^j) : that is, the valuation it announces is t_j . Let (A, P) be the outcome of the strategy profile $s = (s_i, s_{-i})$ and (\hat{A}, \hat{P}) that of the strategy profile $\hat{s} = (\hat{s}_i, s_{-i})$. There are three possibilities for the winners under s and \hat{s} : (1) player i is the winner under both of them; (2) player i is the winner under neither of them; and (3) player i is the winner under \hat{s} , but not under s . Below, we consider them one by one.

Case 1: i is the winner under both s and \hat{s} .

In this case, conditioned on $r = Heads$, $(A_i, P_i) = (\hat{A}_i, \hat{P}_i)$ and $u_i(A, P | r = Heads) = u_i(\hat{A}, \hat{P} | r = Heads)$, since under both s and \hat{s} , player i gets his/her favorite subset for free.

Conditioned on $r = Tails$, $A_i = \hat{A}_i = \emptyset$, $\hat{P}_i = \sum_{j: \hat{\alpha}^i_j \neq \emptyset \text{ and } t_j(\hat{\alpha}^i_j) < \hat{\pi}^i_j} \hat{\pi}^i_j - \delta_i$, and $P_i = \sum_{j: \alpha^i_j \neq \emptyset \text{ and } t_j(\alpha^i_j) < \pi^i_j} \pi^i_j - \delta_i$, where δ_i is player i 's reward under \hat{s} and δ_i is that under s . By the definition of $\hat{\omega}^i$, the set $\{j : \hat{\alpha}^i_j \neq \emptyset \text{ and } t_j(\hat{\alpha}^i_j) < \hat{\pi}^i_j\}$ is empty, so $\hat{P}_i = -\delta_i$. As $REV(\hat{\omega}^i) > REV(\omega^i)$, by definition we have $\hat{\delta}_i > \delta_i$, which implies $\hat{P}_i < -\delta_i \leq P_i$. Accordingly, $u_i(\hat{A}, \hat{P} | r = Tails) > u_i(A, P | r = Tails)$.

In sum, we have $u_i(\hat{A}, \hat{P}) > u_i(A, P)$ in Case 1.

Case 2: i is the winner under neither s nor \hat{s} .

In this case, the winner w is the same under both strategy profiles. Conditioned on $r = Heads$, $A_i = \hat{A}_i = \emptyset$ and $P_i = \hat{P}_i = 0$, thus $u_i(A, P | r = Heads) = u_i(\hat{A}, \hat{P} | r = Heads)$.

Conditioned on $r = Tails$, i gets the set α^i_w under s if and only if he/she gets it under \hat{s} , as he/she announces valuation θ_i under both strategy profiles. That is, $A_i = \hat{A}_i$. Moreover, the only difference in player i 's prices is the rewards he/she gets, and $P_i - \hat{P}_i = -\delta_i + \hat{\delta}_i > 0$. Accordingly, $u_i(\hat{A}, \hat{P} | r = Tails) > u_i(A, P | r = Tails)$.

In sum, we have $u_i(\hat{A}, \hat{P}) > u_i(A, P)$ in Case 2.

Case 3: i is the winner under \hat{s} , but not under s .

In this case, letting w be the winner under s , we have $REV(\omega^i) \leq REV(\omega^w) \leq REV(\hat{\omega}^i)$, and at least one of the inequalities is strict. We compare player i 's utilities under s and \hat{s} , but conditioned on different outcomes of the random coin. More specifically, we use r to denote the outcome of the coin under s and \hat{r} that under \hat{s} .

First, conditioned on $\hat{r} = Heads$, $\hat{A}_i = \text{argmax}_{a \in [m]} \theta_i(a)$ and $\hat{P}_i = 0$; thus:

$$u_i(\hat{A}, \hat{P} | \hat{r} = Heads) = \theta_i(\hat{A}_i).$$

While conditioned on $r = Tails$, we have either $A_i = \emptyset$ and $P_i = -\delta_i$, or $A_i = \alpha_i^w \neq \emptyset$ and $P_i = \pi_i^w - \frac{\epsilon}{n} - \delta_i$; thus:

$$\begin{aligned} u_i(A, P|r = Tails) &\leq \max\{\delta_i, \theta_i(\alpha_i^w) - \pi_i^w + \frac{\epsilon}{n} + \delta_i\} \leq \max\{\delta_i, \theta_i(\hat{A}_i) - 1 + \frac{\epsilon}{n} + \delta_i\} \\ &\leq \theta_i(\hat{A}_i) + \delta_i, \end{aligned}$$

where the second inequality is because $\theta_i(\alpha_i^w) \leq \theta_i(\hat{A}_i)$ and $\pi_i^w \geq 1$, and the third inequality is because both terms in $\max\{\cdot\}$ are less than or equal to $\theta_i(\hat{A}_i) + \delta_i$. Accordingly,

$$u_i(\hat{A}, \hat{P}|\hat{r} = Heads) - u_i(A, P|r = Tails) \geq -\delta_i. \tag{3}$$

Second, conditioned on $\hat{r} = Tails$, $\hat{A}_i = \emptyset$ and $\hat{P}_i = -\hat{\delta}_i$, similar to Case 1 above. Thus:

$$u_i(\hat{A}, \hat{P}|\hat{r} = Tails) = \hat{\delta}_i.$$

While conditioned on $r = Heads$, $A_i = \emptyset$ and $P_i = 0$; thus:

$$u_i(A, P|r = Heads) = 0.$$

Accordingly,

$$u_i(\hat{A}, \hat{P}|\hat{r} = Tails) - u_i(A, P|r = Heads) \geq \hat{\delta}_i. \tag{4}$$

Combining Inequalities 3 and 4 and given that r and \hat{r} are both fair coins, we have:

$$u_i(\hat{A}, \hat{P}) - u_i(A, P) \geq \frac{\hat{\delta}_i - \delta_i}{2} > 0,$$

thus $u_i(\hat{A}, \hat{P}) > u_i(A, P)$ in Case 3, as well.

In sum, $s_i \leq_{U(t)}^{\theta_i} \hat{s}_i$ for all $t \in \mathcal{B}_i$, which implies $s_i \notin \mathcal{C}_i$. Thus, Lemma 2 holds. \square

We now analyze the revenue of M_{Normal} .

Proof of Theorem 1. Arbitrarily fix an auction context $C = (n, m, \theta, \mathcal{B})$ and a strategy profile $s \in \mathcal{C}$. By Lemma 1, we can write $s_i = (\theta_i, \omega^i)$ for each player i . Let (A, P) be the outcome of M_{Normal} under s . By Lemma 2, $REV(\omega^i) \geq BB_i$ for each i , so:

$$R_w = \max_i REV(\omega^i) \geq \max_i BB_i = BB(C).$$

Note that $REV(A, P|r = Heads) = 0$, while:

$$\begin{aligned} REV(A, P|r = Tails) &= \sum_i P_i \\ &= P_w + \sum_{i:\alpha_i^w \neq \emptyset, \theta_i(\alpha_i^w) \geq \pi_i^w} (\pi_i^w - \frac{\epsilon}{n} - \delta_i) + \sum_{i:\alpha_i^w \neq \emptyset, \theta_i(\alpha_i^w) < \pi_i^w} (-\delta_i) + \sum_{i:\alpha_i^w = \emptyset, i \neq w} (-\delta_i) \\ &= \left(\sum_{i:\alpha_i^w \neq \emptyset, \theta_i(\alpha_i^w) < \pi_i^w} \pi_i^w \right) - \delta_w + \sum_{i:\alpha_i^w \neq \emptyset, \theta_i(\alpha_i^w) \geq \pi_i^w} (\pi_i^w - \frac{\epsilon}{n} - \delta_i) \\ &+ \sum_{i:\alpha_i^w \neq \emptyset, \theta_i(\alpha_i^w) < \pi_i^w} (-\delta_i) + \sum_{i:\alpha_i^w = \emptyset, i \neq w} (-\delta_i) \\ &\geq \sum_{i:\alpha_i^w \neq \emptyset} \pi_i^w - \sum_i \frac{\epsilon}{n} - \sum_i \delta_i = R_w - \epsilon - \sum_i \delta_i > R_w - \epsilon - \sum_i \frac{\epsilon}{n} = R_w - 2\epsilon \geq BB(C) - 2\epsilon. \end{aligned}$$

Combining the two cases together, we have $REV(A, P) > \frac{BB(C)}{2} - \epsilon$, and Theorem 1 holds. \square

5. Conservative Distinguishable Implementation

Our main mechanism, together with an auction context, specifies an extensive game with perfect information, chance moves and simultaneous moves [47]⁶. For such a mechanism M , we denote the set of all pure strategy profiles by $S = S_1 \times \dots \times S_n$, the history of a strategy profile s by $H(s)$ and, again, the outcome of s by $M(s)$. If σ is a mixed strategy profile, then $H(\sigma)$ and $M(\sigma)$ are the corresponding distributions.

Even for extensive games of complete information, the literature has several notions of rationality, with different epistemic foundations and predictions about the players' strategies. Since our setting is of incomplete information without Bayesian beliefs, it is important to define a proper solution concept in order to analyze mechanisms in such settings. Iterated eliminations of dominated strategies and their epistemic characterizations have been the focus of many studies in epistemic game theory. In [30], the authors define distinguishable dominance, prove that it is order independent with respect to surviving histories and characterize it with extensive-form rationalizability [31,32,48]. In some sense, distinguishable dominance is the counterpart of strict dominance in extensive-form games. We incorporate this solution concept with the players' possibilistic beliefs.

Definition 4. Let $C = (n, m, \theta, \mathcal{B})$ be an auction context, M an extensive-form mechanism, i a player, t_i a valuation of i and $T = T_i \times T_{-i}$ a set of pure strategy profiles. A strategy $s_i \in T_i$ is distinguishably-dominated by another strategy $\sigma_i \in \Delta(T_i)$ with respect to t_i and T , in symbols $s_i \prec_T^{t_i} \sigma_i$, if:

1. $\exists s_{-i} \in T_{-i}$ distinguishing s_i and σ_i : that is, $H(s_i, s_{-i}) \neq H(\sigma_i, s_{-i})$; and
2. $u_i(t_i, M(s_i, s_{-i})) < u_i(t_i, M(\sigma_i, s_{-i})) \forall s_{-i} \in T_{-i}$ distinguishing s_i and σ_i .

Intuitively, s_i is distinguishably dominated by σ_i if it leads to a smaller utility for i than σ_i , when played against any s_{-i} , except those s_{-i} that produce the same history with s_i and with σ_i : when such an s_{-i} is used, not only player i has the same utility under s_i and σ_i , but also nobody can distinguish whether i is using s_i or σ_i by observing the history of the game.

For each player i , we denote by $DU_i(t_i, T)$ the set of strategies in T_i that are not distinguishably dominated with respect to t_i and T and by DU_i the set $DU_i(\theta_i, S)$. Having seen how to incorporate the iterated elimination of weakly-dominated strategies into our setting, the readers should find the following definition a natural analog.

Definition 5. Let $C = (n, m, \theta, \mathcal{B})$ be an auction context, M a mechanism and i a player. The set of conservatively distinguishably-rational strategies of player i is:

$$CD_i \triangleq DU_i \setminus \{s_i : \exists \sigma_i \in \Delta(DU_i) \text{ s.t. } \forall t \in \mathcal{B}_i, s_i \prec_{DU_i(t)}^{\theta_i} \sigma_i\},$$

where $DU_i(t) \triangleq \times_j DU_j(t_j, S)$ for any $t \in \Theta$. The set of conservatively distinguishably-rational strategy profiles is $CD = CD_1 \times \dots \times CD_n$.

Mechanism M conservatively distinguishably implements a revenue benchmark f if, for any auction context C and any strategy profile $s \in CD$, $REV(M(s)) \geq f(C)$.

A player i may further refine CD_i , but doing so requires more than mutual belief of rationality. We thus do not consider any further refinements.

⁶ Section 6.3 of [47] defines extensive games with perfect information and chance moves, as well as extensive games with perfect information and simultaneous moves. It is easy to combine the two to define extensive games with all three characteristics. Such a game can be described by a "game tree". A decision node is an internal node, where the players take actions or chance moves. A terminal node is a leaf, where an outcome is specified. The history of a strategy profile is the probability distribution over paths from the root to the leaves determined by this profile. The outcome of a strategy profile is the probability distribution over outcomes at the leaves determined by this profile.

6. The Best-Belief Mechanism

Now, we construct and analyze our best-belief mechanism M_{BB} . Similar to the normal-form mechanism, it is parameterized by n, m and a constant $\epsilon \in (0, 1]$. In the description below, Steps 1–3 correspond to decision nodes, while Steps a–e are again “steps taken by the mechanism”.

The best-belief mechanism, M_{BB} :

- a: Set $A_i := \emptyset$ and $P_i := 0$ for each player i .
- 1: Each player i , publicly and simultaneously with the other players, announces:
 - (1) a subset ξ_i of the goods; and
 - (2) an outcome $\omega^i = (\alpha^i, \pi^i)$, such that: $\alpha_j^i = \emptyset$, and for each player j , π_j^i is zero whenever $\alpha_j^i = \emptyset$ and is a positive integer otherwise.
- b: Set $R_i := REV(\omega^i)$ for each player i and $w := \operatorname{argmax}_i R_i$ with ties broken lexicographically.
- 2: Publicly flip a fair coin, and denote the result by r .
- c: If $r = \text{Heads}$, then $A_w := \xi_w$, and halt.
- 3: (Note that $r = \text{Tails}$ when this step is reached.)

Each player i , such that $\alpha_i^w \neq \emptyset$ publicly and simultaneously announces YES or NO.
- d: For each player i announcing NO, $P_w := P_w + \pi_i^w$.

For each player i announcing YES, $A_i := \alpha_i^w$ and $P_i := \pi_i^w - \frac{\epsilon}{n}$.

For each player i , $P_i := P_i - \delta_i$ with $\delta_i = \frac{\epsilon}{n} \cdot \frac{R_i}{1+R_i}$.
- e: The final outcome is (A, P) .

6.1. Analysis of Our Mechanism

As before, given a context $(n, m, \theta, \mathcal{B})$ and an outcome ω , we use $u_i(\omega)$ instead of $u_i(\theta_i, \omega)$ for player i 's utility under ω . We have the following.

Theorem 2. For any context $(n, m, \theta, \mathcal{B})$ and constant $\epsilon \in (0, 1]$, mechanism M_{BB} conservatively distinguishably implements the revenue benchmark $\frac{BB}{2} - \epsilon$.

Different from the normal-form mechanism, here, a player does not report his/her true valuation. Instead, the use of his/her valuation is divided into two parts: a subset of the goods, which will be his/her favorite subset as we will see in the proof; and a simple “yes or no” answer to the take-it-or-leave-it offer suggested by the winner. All of the other information about his/her true valuation is redundant and has been removed from the player’s report. This can be done because the mechanism is extensive and the players give their answers directly after seeing the offers; thus, the seller does not need to deduce their answers from their reported valuations. We again start by proving the following two lemmas. Some ideas are similar to those for Lemmas 1 and 2; thus, the details have been omitted.

Lemma 3. For any context $(n, m, \theta, \mathcal{B})$, constant ϵ , player i and strategy s_i , if $s_i \in DU_i$, then, according to s_i , in Step 3 of \mathcal{M}_{BB} , i announces YES if and only if $\theta_i(\alpha_i^w) \geq \pi_i^w$ ⁷.

⁷ That is, i will announce YES or NO as above at every decision node corresponding to Step 3, which is reachable (with positive probability) by s_i together with some strategy sub-profile s_{-i} , where i is an acting player.

Proof. We only prove the “if” direction, as the “only if” direction is totally symmetric. Assume that, according to s_i , i announces NO at some reachable decision node d of i where $\theta_i(\alpha_i^w) \geq \pi_i^w$. We refer to such a node d as a deviating node. Consider the following strategy s'_i :

- s'_i announces the same ξ_i and ω^i as s_i in Step 1; and
- according to s'_i , in Step 3, i announces YES if and only if $\theta_i(\alpha_i^w) \geq \pi_i^w$.

Below, we show that $s_i \prec_S^{\theta_i} s'_i$, where S is the set of all strategy profiles of M_{BB} .

For any deviating node d , since d is reachable by s_i , there exists a strategy sub-profile $s_{-i} \in S_{-i}$, such that the history $H(s_i, s_{-i})$ reaches d with positive probability. In fact, by the construction of the mechanism, the probability is exactly 1/2: when $r = Tails$. For any such s_{-i} , by the construction of s'_i , the history $H(s'_i, s_{-i})$ also reaches d with probability 1/2. By definition, i announces YES at d under s'_i and NO under s_i ; thus, $H(s_i, s_{-i} | r = Tails) \neq H(s'_i, s_{-i} | r = Tails)$ and s_{-i} distinguishes s_i and s'_i .

Indeed, for any strategy sub-profile s_{-i} , it distinguishes s_i and s'_i if and only if $H(s_i, s_{-i})$ reaches a deviating node d (with probability 1/2). Arbitrarily fixing such an s_{-i} and the corresponding deviating node d , it suffices to show:

$$u_i(M_{BB}(s_i, s_{-i})) < u_i(M_{BB}(s'_i, s_{-i})). \tag{5}$$

Because $i \neq w$ under (s_i, s_{-i}) when $r = Tails$ (that is, when d is reached), $i \neq w$ under (s_i, s_{-i}) when $r = Heads$, as well, since w is the same in the two cases. Moreover, because s'_i announces the same ξ_i and ω^i as s_i in Step 1, we have $H(s_i, s_{-i} | r = Heads) = H(s'_i, s_{-i} | r = Heads)$ and $u_i(M_{BB}(s_i, s_{-i}) | r = Heads) = u_i(M_{BB}(s'_i, s_{-i}) | r = Heads) = 0$.

Similar to Lemma 1, $u_i(M_{BB}(s_i, s_{-i}) | r = Tails) = \delta_i$, as i announces NO at d under s_i . Furthermore, $u_i(M_{BB}(s'_i, s_{-i}) | r = Tails) = \theta_i(\alpha_i^w) - P_i = \theta_i(\alpha_i^w) - \pi_i^w + \frac{\epsilon}{n} + \delta_i \geq \frac{\epsilon}{n} + \delta_i > \delta_i$, as $\theta_i(\alpha_i^w) \geq \pi_i^w$ at d , and i announces YES at d under s'_i . Therefore, $u_i(M_{BB}(s'_i, s_{-i}) | r = Tails) > u_i(M_{BB}(s_i, s_{-i}) | r = Tails)$, which implies Equation (5). Accordingly, $s_i \prec_S^{\theta_i} s'_i$, $s_i \notin DU_i$, and Lemma 3 holds. \square

Lemma 4. For any context $(n, m, \theta, \mathcal{B})$, constant ϵ , player i and strategy s_i , if $s_i \in \mathcal{CD}_i$, then, according to s_i , player i announces ω^i in Step 1 with $REV(\omega^i) \geq BB_i$.

Proof. Arbitrarily fix a strategy $s_i \in DU_i$ according to which, in Step 1, i announces ξ_i , and $\omega^i = (\alpha^i, \pi^i)$ with $REV(\omega^i) < BB_i$. By Lemma 3, according to s_i , in Step 3, i announces YES if and only if $\theta_i(\alpha_i^w) \geq \pi_i^w$. Consider the following strategy \hat{s}_i :

- In Step 1, i announces $\hat{\xi}_i$, and $\hat{\omega}^i = (\hat{\alpha}^i, \hat{\pi}^i)$, such that:
 - $\theta_i(\hat{\xi}_i) = \max_{A \subseteq \{1, \dots, m\}} \theta_i(A)$;
 - $REV(\hat{\omega}^i) = \max_{(A, P) \in \Omega: A_i = \emptyset \text{ and } P_j \leq t_j(A_j) \forall j \neq i, \forall t \in \mathcal{B}_i} REV(A, P)$; and
 - $\hat{\pi}^i_j > 0$ whenever $\hat{\alpha}^i_j \neq \emptyset$.
- In Step 3, i announces YES if and only if $\theta_i(\alpha_i^w) \geq \pi_i^w$.

By definition, $REV(\hat{\omega}^i) = BB_i > REV(\omega^i)$, which implies that \hat{s}_i and s_i differ in Step 1: the root of the game tree. Thus, any strategy sub-profile s_{-i} distinguishes them. We show that for all

$t \in \mathcal{B}_i, s_i \prec_{DU(t)}^{\theta_i} \hat{s}_i$ ⁸. To do so, arbitrarily fixing a valuation profile $t \in \mathcal{B}_i$ and a strategy sub-profile $s_{-i} \in \times_{j \neq i} DU_j(t_j, S)$, it suffices to show:

$$u_i(M_{BB}(s_i, s_{-i})) < u_i(M_{BB}(\hat{s}_i, s_{-i})). \tag{6}$$

Let δ_i and $\hat{\delta}_i$ be the rewards of player i in Step d, under (s_i, s_{-i}) and (\hat{s}_i, s_{-i}) , respectively. Because $REV(\hat{\omega}^i) > REV(\omega^i)$, we have:

$$\delta_i < \hat{\delta}_i. \tag{7}$$

Similar to Lemma 2, we distinguish three cases.

Case 1. i is the winner under both (s_i, s_{-i}) and (\hat{s}_i, s_{-i}) .

In this case, on the one hand,

$$u_i(M_{BB}(s_i, s_{-i}) | r = Heads) = \theta_i(\xi_i) \leq \theta_i(\hat{\xi}_i) = u_i(M_{BB}(\hat{s}_i, s_{-i}) | r = Heads),$$

where the inequality is by the definition of $\hat{\xi}_i$.

On the other hand, $u_i(M_{BB}(s_i, s_{-i}) | r = Tails) = -(\sum_{j:j \text{ announces NO in } (s_i, s_{-i})} \pi_j^i - \delta_i) \leq \delta_i$ and $u_i(M_{BB}(\hat{s}_i, s_{-i}) | r = Tails) = -(\sum_{j:j \text{ announces NO in } (\hat{s}_i, s_{-i})} \hat{\pi}_j^i - \hat{\delta}_i)$. For any player j , such that $\hat{\alpha}_j^i \neq \emptyset$, because $t \in \mathcal{B}_i$, by the construction of $\hat{\omega}^i$, we have $\hat{\pi}_j^i \leq t_j(\hat{\alpha}_j^i)$. Because $\hat{s}_j \in DU_j(t_j, S)$, by Lemma 3, j announces YES in Step 3 under (\hat{s}_i, s_{-i}) . Accordingly, $\sum_{j:j \text{ announces NO in } (\hat{s}_i, s_{-i})} \hat{\pi}_j^i = 0$ and:

$$u_i(M_{BB}(\hat{s}_i, s_{-i}) | r = Tails) = \hat{\delta}_i > \delta_i = u_i(M_{BB}(s_i, s_{-i}) | r = Tails),$$

where the inequality is by Equation (7). In sum, Equation (6) holds in Case 1.

Case 2. i is the winner under neither (s_i, s_{-i}) nor (\hat{s}_i, s_{-i}) .

Letting w be the winner under both strategy profiles, we have $u_i(M_{BB}(s_i, s_{-i}) | r = Heads) = u_i(M_{BB}(\hat{s}_i, s_{-i}) | r = Heads) = 0$. Moreover, conditioned on $r = Tails$, by the construction of \hat{s}_i , player i announces the same thing under (s_i, s_{-i}) and (\hat{s}_i, s_{-i}) . Thus, the only difference between i 's allocation and price under the two strategy profiles is the rewards: one is δ_i , and the other is $\hat{\delta}_i$. Accordingly, $u_i(M_{BB}(s_i, s_{-i}) | r = Tails) - u_i(M_{BB}(\hat{s}_i, s_{-i}) | r = Tails) = \delta_i - \hat{\delta}_i < 0$, where the inequality is by Equation (7). In sum, Equation (6) holds in Case 2.

Case 3. i is the winner under (\hat{s}_i, s_{-i}) , but not under (s_i, s_{-i}) .

In this case, let w be the winner under (s_i, s_{-i}) and r and \hat{r} be the outcomes of the coins under (s_i, s_{-i}) and (\hat{s}_i, s_{-i}) , respectively. Similar to Lemma 2, we have:

$$\begin{aligned} u_i(M_{BB}(s_i, s_{-i}) | r = Tails) &\leq \max\{\delta_i, \theta_i(\alpha_i^w) - \pi_i^w + \frac{\epsilon}{n} + \delta_i\} \leq \theta_i(\alpha_i^w) + \delta_i, \\ u_i(M_{BB}(s_i, s_{-i}) | r = Heads) &= 0, \\ u_i(M_{BB}(\hat{s}_i, s_{-i}) | \hat{r} = Heads) &= \theta_i(\hat{\xi}_i), \end{aligned}$$

⁸ Without loss of generality, we can assume $\hat{s}_i \in \mathcal{CD}_i$. Otherwise, by the well-studied properties of distinguishable dominance [30], there exists $\sigma_i \in \Delta(\mathcal{CD}_i)$, such that $\hat{s}_i \prec_{DU(t)}^{\theta_i} \sigma_i$ for all $t \in \mathcal{B}_i$, and we can prove $s_i \prec_{DU(t)}^{\theta_i} \sigma_i$.

and:

$$u_i(M_{BB}(\hat{s}_i, s_{-i}) | \hat{r} = Tails) = - \left(\sum_{j: j \text{ announces NO in } (\hat{s}_i, s_{-i})} \hat{\pi}_j^i - \hat{\delta}_i \right) = \hat{\delta}_i.$$

Accordingly,

$$u_i(M_{BB}(\hat{s}_i, s_{-i})) = \frac{\theta_i(\hat{\xi}_i) + \hat{\delta}_i}{2} > \frac{\theta_i(\alpha_i^w) + \delta_i}{2} \geq u_i(M_{BB}(s_i, s_{-i})),$$

and Equation (6) holds in Case 3.

Therefore, $s_i \notin \mathcal{CD}_i$, and Lemma 4 holds. \square

Proof of Theorem 2. Given Lemmas 3 and 4, the proof of Theorem 2 is almost the same as that of Theorem 1, except that, rather than distinguishing players with $\theta_i(\alpha_i^w) \geq \pi_i^w$ or $\theta_i(\alpha_i^w) < \pi_i^w$, here, we distinguish players announcing YES or NO in Step 3. The details have been omitted. \square

Note that the revenue guarantee of the mechanism holds no matter whether the players' beliefs about each other are correct or not. If a player i has low values for the goods and believes the others' values to be high and if the others' true values and beliefs are all low, then player i may end up being the winner and getting a negative utility. However, according to player i 's beliefs, his/her utility will always be positive, and it is individually rational for him to participate. This is not too dissimilar to the stock market, where not everybody makes money, but everybody believes he/she will make money when entering. Indeed, the final outcome implemented may not be an ex-post Nash equilibrium and instead is supported by the two-step elimination of dominated strategies.

Furthermore, note that the idea of asking players to report their beliefs about each other has been explored in the Nash implementation literature (see, e.g., [49,50]). However, our mechanism does not assume complete information or common beliefs. Moreover, our mechanism does not try to utilize the winner's true valuations for generating revenue: indeed, the focus here is how to generate revenue by leveraging the players' beliefs. Simply choosing at random this mechanism or the VCG mechanism (or any other mechanism for unrestricted combinatorial auctions that may achieve better revenue in some contexts), one can achieve a good approximation to the best of the two.

Finally, it suffices for the players' values to be numbers within certain precisions, say two decimal digits, so that there is a gap between any two different values. If the values are real numbers, then the rewards in our mechanisms are set to zero, and our results hold under a weaker notion of dominance: that is, the desired strategies are still at least as good as any deviation, but may not be strictly better.

6.2. Privacy, Complexity and Collusion in Our Mechanism

Finally, we discuss the resiliency of our mechanism with respect to privacy, complexity and collusion concerns.

6.2.1. Privacy

Our main mechanism achieves our revenue benchmark by eliciting from the players much less information than they possess. In Step 1, a player does not reveal anything about his/her own valuation except a subset of goods, which is supposed to be his/her favorite subset. Nor does he/she reveal his/her full beliefs about the valuations of his/her opponents: he/she only reveals a maximum guaranteed-revenue outcome, according to his/her beliefs.

This is all of the information that is revealed if the coin flipped by the mechanism ends up as heads. If it ends up as tails, then a player i reveals at most a modest amount of information about his/her own true valuation in Step 3. Namely, only if he/she is offered a subset A of goods for a price p , he/she reveals that his/her true value for that specific subset is $\geq p$ if he/she answers YES and $< p$ otherwise. In particular, therefore, in our mechanism, a player who is not offered any goods does not

reveal any information about his/her own valuation. This is very far from what may happen in many other auction mechanisms: that is, fully revealing your valuation and receiving no goods.

Because privacy is important to many strategic agents, we hope that trying to preserve it will become a standard goal in mechanism design. Achieving this goal will require putting a greater emphasis on extensive mechanisms, where the players and the mechanism may interact over time⁹. The power of “interaction” for privacy preservation is very well documented in cryptography¹⁰. This power extends to mechanism design, as well: as we have seen in our case, even three sequential moves can save a considerable amount of privacy compared with the previous normal-form mechanism.

6.2.2. Computation and Communication Efficiency

Our mechanism is highly efficient in both computation and communication. Essentially, it only needs to sum up the prices in each reported outcome ω^i and figure out which reported outcome has the highest revenue. Moreover, each player only reports a subset of goods and an outcome and perhaps announces YES or NO in Step 3. One might object, however, that our mechanism transfers all of the hard computation to the players themselves. This is indeed true, but our mechanism also gives them the incentives to approximate this hard computation.

As we have recalled in our Introduction, approximation (1) may be necessary to compute a reasonable outcome when finding “the best one” is computationally hard, but (2) may also distort incentives. Our mechanism instead ensures that approximation is aligned with incentives. Indeed, our mechanism entrusts the players to propose outcomes, but ensures, as per Lemma 4, that each player wishes to become the winner. Thus, our mechanism makes it in a player’s own interest to use the best computationally-efficient approximation algorithm he/she knows, in order to propose a high-revenue outcome. Of course, the best algorithm known by a player may not be the best in the literature, in terms of its approximation ratio to the optimal outcome. In this case, the mechanism’s revenue is at least half of the highest revenue the players are capable of computing. To our best knowledge, this is the first mechanism that gives the buyers incentives to perform computationally-efficient approximations. Incentive-compatible and computationally-efficient approximation on the seller’s side has also been studied, but again for valuations of restricted forms, such as single-minded players [1], single-value players [2], auctions of multiple copies of the same good [53], etc. By contrast, we do not impose any such restrictions.

6.2.3. Collusion

Collusion is traditionally prohibited (e.g., by using solution concepts that only consider individual deviations in game theory) and punished (e.g., by laws). However, it continues to exist. We thus wish to point out that our mechanism offers a reasonable form of protection against collusion. Namely, when at least some players are independent, denoting by I the set of independent player, it guarantees at least half of the revenue benchmark $BB' \triangleq \max_{i \in I} BB_i$.

Thus, our mechanism is not responsible for generating any revenue if all players are collusive, but must generate revenue at least half of BB' otherwise. This guarantee holds in a strong collusion model: that is, even when collusive players are capable of making side payments and coordinating their actions via secret and enforceable agreements, and the independent players have no idea that collusion is afoot. The only constraint is that every coalition is rational, that is, its members act so

⁹ It is well known that every extensive mechanism can be transformed to an “equivalent” normal-form game, but this equivalence does not extend to privacy. Indeed, in an extensive mechanism M , a player i reveals information only if a decision node of i is reached, and in an execution of M , only some of these nodes are reached. Transforming M into the normal form instead asks i to reveal how he/she would like to act at any possible decision node involving him.

¹⁰ Interaction is indeed at the base of zero-knowledge proofs [51,52], where a mistrusted prover can convince a skeptical verifier that a theorem statement is true without revealing any additional information.

to maximize the sum of their individual utilities. In this case, an independent player i , reporting in Step 1 an outcome ω offering a player j a subset of the goods X for a price p , need not worry whether j is independent or collusive. If i becomes the winner and the coin toss of the mechanism is tails, then j will answer YES if and only if his/her individual true value for X is greater than or equal to p . Accordingly, i will report in Step 1 an outcome whose revenue is at least BB_i . If an independent player becomes the winner, then the mechanism will generate at least $BB'/2$ revenue. Else, some collusive player has become the winner; but then, such a player must have reported an outcome with revenue $R \geq BB'$, and the mechanism will generate at least $R/2$ revenue.

Let us point out that the BB' benchmark is actually guaranteed under a weaker requirement of coalition rationality¹¹.

6.2.4. Social Welfare

Note that each player i has a “truthful” strategy: to report θ_i and the outcome $\hat{\omega}^i$ as defined in the proof of Lemma 4, whose revenue is exactly BB_i . Since the price suggested by $\hat{\omega}^i$ for each player $i' \neq i$ is no more than the true value of i' for the suggested subset of goods for him/her, the players all say YES when i is the winner, and player i' 's utility is non-negative. Under the truthful strategy profile, the social welfare of the final outcome is at least $\frac{BB}{2}$. When the players overbid and report outcomes whose revenue is higher than the corresponding BB_i 's, the social welfare may be smaller than the revenue.

6.3. Variants of Our Mechanism

Our mechanism sets aside a “budget” of $\epsilon > 0$ for rewarding the players and achieves the benchmark $BB/2 - \epsilon$. We note that our analysis also holds if the mechanism chooses his/her reward budget to be not an absolute value ϵ , but an ϵ fraction of the revenue it collects. In such a case, however, its guaranteed revenue will be $(1 - \epsilon)BB/2$.

Furthermore, for simplicity, we have assumed that the seller/designer knows nothing about the players. However, it is easy to accommodate the case in which he/she too has some beliefs about the players' valuations. For instance, in keeping with our overall approach, let ω' be the highest revenue outcome among all of the outcomes (A, P) for which he/she is sure that $\theta_i(A_i) \geq P_i$ for all i . Then, he/she can use ω' as a “reserve outcome” as follows. If, in Step 1, the revenue of the outcome reported by the winner is at least that of ω' , then he/she keeps on running our mechanism; otherwise, roughly speaking, he/she makes himself the “winner” and continues with ω' being the outcome reported by the winner.

Acknowledgments: We thank David Easley, Shafi Goldwasser, Avinatan Hassidim, Robert Kleinberg, Eric Maskin, Paul Milgrom, Rafael Pass, Ron Rivest, Paul Valiant and Avi Wigderson for comments and encouragement and several anonymous reviewers for helpful suggestions about the presentation of our results. The authors were partially supported by Office of Naval Research (ONR) Grant No. N00014-09-1-0597. The first author is partially supported by NSF CAREER Award No. 1553385.

Author Contributions: Jing Chen and Silvio Micali contributed equally to this work. Jing Chen and Silvio Micali wrote the paper.

Conflicts of Interest: The authors declare no conflict of interest.

¹¹ That is, when the members of a coalition act so as to maximize a different function of their individual utilities. All we need is a mild “monotonicity” condition, informally described as follows. Consider a coalition C and two outcomes ω and ω' , such that (1) a member i of C is offered a set A_i for a price P_i in outcome ω and no goods for price P'_i in ω' ; and (2) every other member j of C is offered the same set of goods A_j for the same price P_j in both outcomes. Then, the only rationality condition that we require from C is that it prefers ω to ω' if and only if $\theta_i(A_i) - P_i \geq -P'_i$. Under this model, in Step 3 of our mechanism, each coalition can delegate the YES or NO decisions to its members as if they were independent. Thus, again, an independent player need not worry whether another player is independent or collusive.

References

1. Lehmann, D.; O'Callaghan, L.; Shoham, Y. Truth revelation in approximately efficient combinatorial auctions. *J. ACM* **2002**, *49*, 577–602.
2. Babaioff, M.; Lavi, R.; Pavlov, E. Single-Value combinatorial auctions and implementation in undominated strategies. In Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2006), Miami, FL, USA, 22–26 January 2006; pp. 1054–1063.
3. Hart, S.; Nisan, N. Approximate revenue maximization with multiple items. In Proceedings of the 13th ACM Conference on Electronic Commerce (EC), Valencia, Spain, 4–8 June 2012; p. 656.
4. Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *J. Financ.* **1961**, *16*, 8–37.
5. Clarke, E. Multipart pricing of public goods. *Public Choice* **1971**, *11*, 17–33.
6. Groves, T. Incentives in teams. *Econometrica* **1973**, *41*, 617–631.
7. Abreu, D.; Matsushima, H. Virtual Implementation in iteratively undominated actions: Complete information. *Econometrica* **1992**, *60*, 993–1008.
8. Glazer, J.; Perry, M. Virtual implementation in backwards induction. *Games Econ. Behav.* **1996**, *15*, 27–32.
9. Chen, J.; Hassidim, A.; Micali, S. Robust perfect revenue from perfectly informed players. In Proceedings of the Innovations in Theoretical Computer Science (ITCS), Beijing, China, 5–7 January 2010; pp. 94–105.
10. Chen, J.; Micali, S. Mechanism design with possibilistic beliefs. *J. Econ. Theory* **2015**, *156*, 77–102.
11. Harsanyi, J. Games with incomplete information played by “Bayesian” players, I–III. *Manag. Sci.* **1967–1968**, *14*, 159–182, 320–334, 486–502.
12. Aumann, R. Agreeing to disagree. *Ann. Stat.* **1976**, *4*, 1236–1239.
13. Myerson, R. Optimal auction design. *Math. Oper. Res.* **1981**, *6*, 58–73.
14. Ausubel, L.; Milgrom, P. The lovely but lonely Vickrey auction. In *Combinatorial Auctions*; Cramton, P., Shoham, Y., Steinberg, R., Eds.; MIT Press: Cambridge, MA, USA, 2006; pp. 17–40.
15. Aumann, R. Backward Induction and Common Knowledge of Rationality. *Games Econ. Behav.* **1995**, *8*, 6–19.
16. Milgrom, P. *Putting Auction Theory to Work*; Cambridge University Press: Cambridge, UK, 2004.
17. Cramton, P.; Shoham, Y.; Steinberg, R. (Eds.) *Combinatorial Auctions*; MIT Press: Cambridge, MA, USA, 2006.
18. Klemperer, P. *Auctions: Theory and Practice*; Princeton University Press: Princeton, NJ, USA, 2004.
19. Milgrom, P.; Ausubel, L.; Levin, J.; Segal, I. *Incentive Auction Rules Option and Discussion*; Technical Report, FCC-12-118A2, Federal Communications Commission: Washington, DC, USA, 2012.
20. Milgrom, P.; Segal, I. Deferred-Acceptance Auctions and Radio Spectrum Reallocation. In Proceedings of the 15th ACM Conference on Economics and Computation (EC), Palo Alto, CA, USA, 8–12 June 2014; pp. 185–186.
21. Kash, I.A.; Murty, R.; Parkes, D.C. Enabling Spectrum Sharing in Secondary Market Auctions. *IEEE Trans. Mob. Comput.* **2014**, *13*, 556–568.
22. Cramton, P.; Lopez, H.; Malec, D.; Sujarittanonta, P. *Design of the Reverse Auction in the Broadcast Incentive Auction*; Working Paper; University of Maryland: College Park, MD, USA, 2015.
23. Nguyen, T.D.; Sandholm, T. Multi-Option Descending Clock Auction. In Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Singapore, 9–13 May 2016; pp. 1461–1462.
24. Milgrom, P.; Weber, R. A Theory of Auctions and Competitive Bidding, II. In *The Economic Theory of Auctions*; Klemperer, P., Ed.; Edward Elgar: Cheltenham, UK, 2000; Volume I, pp. 179–194.
25. Parkes, D. Iterative Combinatorial Auctions. In *Combinatorial Auctions*; Cramton, P., Shoham, Y., Steinberg, R., Eds.; MIT Press: Cambridge, MA, USA, 2006; Chapter 2.
26. Likhodedov, A.; Sandholm, T. Approximating Revenue-Maximizing Combinatorial Auctions. In Proceedings of the 20th National Conference on Artificial Intelligence (AAAI), Pittsburgh, PA, USA, 9–13 July 2005; pp. 267–274.
27. Balcan, M.; Blum, A.; Mansour, Y. *Single Price Mechanisms for Revenue Maximization in Unlimited Supply Combinatorial Auctions*; Technical Report, CMU-CS-07-111; Carnegie Mellon University: Pittsburgh, PA, USA, 2007.
28. Micali, S.; Valiant, P. *Resilient Mechanisms for Unrestricted Combinatorial Auctions*; Technical Report, MIT-CSAIL-TR-2008-067; Massachusetts Institute of Technology: Cambridge, MA, USA, 2008.

29. Jackson, M. Implementation in undominated strategies: A look at bounded mechanisms. *Rev. Econ. Stud.* **1992**, *59*, 757–775.
30. Chen, J.; Micali, S. The order independence of iterated dominance in extensive games. *Theor. Econ.* **2013**, *8*, 125–163.
31. Pearce, D.G. Rationalizable strategic behavior and the problem of perfection. *Econometrica* **1984**, *52*, 1029–1050.
32. Battigalli, P. On rationalizability in extensive games. *J. Econ. Theory* **1997**, *74*, 40–61.
33. Chen, J.; Micali, S.; Pass, R. Tight revenue bounds with possibilistic beliefs and level- k rationality. *Econometrica* **2015**, *83*, 1619–1639.
34. Chen, J.; Micali, S.; Valiant, P. Robustly leveraging collusion in combinatorial auctions. In Proceedings of the Innovations in Theoretical Computer Science (ITCS), Beijing, China, 5–7 January 2010; pp. 81–93.
35. Izmalkov, S.; Lepinski, M.; Micali, S. Perfect Implementation. *Games Econ. Behav.* **2011**, *71*, 121–140.
36. Parkes, D.; Rabin, M.; Shieber, S.; Thorpe, C. Practical secrecy-preserving, verifiably correct and trustworthy auctions. *Electron. Commer. Res. Appl.* **2008**, *7*, 294–312.
37. Dwork, C.; Roth, A. *The Algorithmic Foundations of Differential Privacy*; Foundations and Trends in Theoretical Computer Science; NOW Publishers: Breda, The Netherlands, 2014.
38. Green, J.; Laffont, J. On coalition incentive compatibility. *Rev. Econ. Stud.* **1979**, *46*, 243–254.
39. Schummer, J. Manipulation through bribes. *J. Econ. Theory* **2000**, *91*, 180–198.
40. Chen, J.; Micali, S. Collusive dominant-strategy truthfulness. *J. Econ. Theory* **2012**, *147*, 1300–1312.
41. Moulin, H.; Shenker, S. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Econ. Theory* **2001**, *18*, 511–533.
42. Jain, K.; Vazirani, V. Applications of approximation algorithms to cooperative games. In Proceedings of the 33rd ACM Symposium on Theory of Computing (STOC), Heraklion, Greece, 6–8 July 2001; pp. 364–372.
43. Feigenbaum, J.; Papadimitriou, C.; Shenker, S. Sharing the cost of multicast transmissions. *J. Comput. Syst. Sci.* **2001**, *63*, 21–41.
44. Laffont, J.; Martimort, D. Mechanism design with collusion and correlation. *Econometrica* **2000**, *68*, 309–342.
45. Goldberg, A.; Hartline, J. Collusion-resistant mechanisms for single-parameter agents. In Proceedings of the Symposium on Discrete Algorithms (SODA), Vancouver, BC, Canada, 23–25 January 2005; pp. 620–629.
46. Che, Y.; Kim, J. Robustly collusion-proof implementation. *Econometrica* **2006**, *74*, 1063–1107.
47. Osborne, M.; Rubinstein, A. *A Course in Game Theory*; MIT Press: Cambridge, MA, USA, 1994.
48. Shimoji, M.; Watson, J. Conditional dominance, rationalizability, and game forms. *J. Econ. Theory* **1998**, *83*, 161–195.
49. Maskin, E. Nash equilibrium and welfare optimality. *Rev. Econ. Stud.* **1999**, *66*, 23–38.
50. Jackson, M.; Palfrey, T.; Srivastava, S. Undominated Nash implementation in bounded mechanisms. *Games Econ. Behav.* **1994**, *6*, 474–501.
51. Goldwasser, S.; Micali, S.; Rackoff, C. The knowledge complexity of interactive proof-systems. *SIAM J. Comput.* **1989**, *18*, 186–208.
52. Goldreich, O.; Micali, S.; Wigderson, A. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proofs. *J. ACM* **1991**, *38*, 691–729.
53. Dobzinski, S.; Dughmi, S. On the power of randomization in algorithmic mechanism design. In Proceedings of the 54th Symposium on Foundations of Computer Science (FOCS), Atlanta, GA, USA, 25–27 October 2009; pp. 505–514.

