Perfect Revenue From Perfectly Informed Players

Jing Chen
CSAIL, MIT
Cambridge, MA 02139, USA
jingchen@csail.mit.edu

Avinatan Hassidim
RLE, MIT
Cambridge, MA 02139, USA
avinatanh@gmail.com

Silvio Micali
CSAIL, MIT
Cambridge, MA 02139, USA
silvio@csail.mit.edu

August 12, 2009

Abstract

Maximizing revenue in the presence of perfectly informed players is a well known goal in mechanism design. Yet, all current mechanisms for this goal are extremely vulnerable to equilibrium selection. In this paper we both clarify and rectify this situation by proving that

- No (even weakly) dominant-strategy mechanism can guarantee an arbitrarily small fraction of the maximum possible revenue; while
- Surviving-strategy mechanisms, a new class of "equilibrium-less" mechanisms, can guarantee a fraction arbitrarily close to 1 of the maximum possible revenue.

We thus provide the first separation result between dominant-strategy and surviving-strategy mechanisms.

1 Introduction

1.1 Classical Mechanism Design

Contexts and mechanisms. A *context* C describes the players, the outcomes and the players' preferences over the outcomes. A traditional and general for of context, indeed the one considered in our paper, is that of quasi-linear utilities with non-negative valuations. Roughly, an outcome consists of a state ω (belonging to a finite set Ω)and a price profile P (specifying a real number P_i for each player i). (For instance Ω may consist of all possible ways of allocating a set of items to the players; which subset of 5 faculty candidates to hire; or a set of possible ways to build a bridge across a given river.) Quasi-linearity means that the utility that each player i has for an outcome (ω, P) is the sum of two components, the non-negative value that i has for the state ω and the price P_i he pays.

A *mechanism* M describes the *strategies* available to the players, and how strategies determine outcomes. For instance, in our considered context, a player strategy may consist of a valuation, that is, of a table specifying the value for each possible ω.

Notice that a context C together with a mechanism M defines a *game* G, $G = (C, M)$, and thus each rational player will endeavor to choose his own strategy so as to maximize his own utility.

The general goal of mechanism design. The goal of mechanism design is to find a mechanism M such that, for a given context C, a desired property $\mathbb{P}$ holds for the outcomes obtained by having the players play (C, M). The difficulty is that the designer of M does not known the players' preferences. In the purest form of mechanism design, all such knowledge is with the players themselves: the designer can only assume that the players are rational when he endeavors to design M so as to "make it in the best interest of the players" so satisfy $\mathbb{P}$. In essence, therefore, mechanism design aims at guaranteeing a property $\mathbb{P}$ by leveraging the players' knowledge and the players' rationality in the sense that at a rational play $\mathbb{P}$ holds. But: *what is a rational play and what is the players' knowledge to be leveraged?*

The classical interpretation. Let us provide the classical answer to both questions.

The classical interpretation of a rational play is an *equilibrium*, that is a profile of strategies $\sigma = \sigma_1, \ldots, \sigma_n$ such that no player i has an incentive to deviate from σ_i to an alternative strategy σ_i'. But equilibria are vastly different in their "quality." The weakest form is that of a *Nash equilibrium,* simply stating that i prefers σ_i to any alternative σ_i' only if he *believes that every other player j will stick to his specific σ_j.* That is, Nash equilibrium only guarantees that i prefers $\sigma_1, \ldots, \sigma_i, \ldots, \sigma_n$ to $\sigma_1, \ldots, \sigma_i', \ldots, \sigma_n$. The strongest form of equilibrium is a *dominant-strategy equilibrium,* where σ_i is the best strategy for i no matter what strategies the other players may choose. More precisely, a dominant-strategy equilibrium σ is called strict (respectively, weak) if, for any player i, any alternative strategy σ_i', and any strategy sub-profile τ_{-i} for the other players, i's utility is strictly larger (respectively, larger or equal to) when i plays σ_i than when he plays σ_i'.

The classical interpretation of player knowledge is *self knowledge*, that is classical mechanisms leverage only the knowledge that every player i has about his own preferences. In the context we focus on, this is the knowledge that i has about his own valuation for the possible states.

1.2 The Problem of Equilibrium Selection and Our Impossibility Result

It should be realized that designing a mechanism so as to guarantee a property $\mathbb{P}$ "at a Nash Equilibrium" is a very weak guarantee, or no guarantee at all. First, because there may be several Nash equilibria, while $\mathbb{P}$ holds for just some of them. Moreover, even if $\mathbb{P}$ held for all equilibria, $\mathbb{P}$ may not hold at all in a real play. For instance, assume that there exist two equilibria, σ and τ, and that some players believe that σ will be played out, while others believe that τ will. Then, because a mix-and-match of σ and τ may not be an equilibrium, $\mathbb{P}$ may not hold at all. Of course, this problem worsens as the number of players and/or equilibria grows.

The problem of equilibrium selection is particularly acute for the classical goal of achieving perfect revenue with perfectly informed players.

Consider the context where there are finitely many possible *states*, $\omega_1, \ldots, \omega_k$; each player i has a non-negative value $v_i(\omega_j)$ for each state ω_j; each outcome consists of a state ω together with a price P_i for each player i; and the utility of each player i for such an outcome is $v_i(\omega) - P_i$. Indeed, this is the setting of non-negative valuation with quasi-linear utilities. How to maximize revenue in this setting assuming that each player not only knows his own valuation but also those of the others?

An obvious mechanism is the following.

HOPE-FOR-THE-BEST: Each player reports the valuations of all players (including himself). If all reports are the same, then (1) choose the state ω maximizing the sum of the reported valuations and (2) for each player i, choose the price P_i to consist of his reported value for ω (minus a small discount ϵ to encourage his participation). If not all reports coincide, then choose the "null outcome" (which all players are assumed to value 0) and the price 0 for every player.

It is trivial to see that the strategy profile in which each player reports everyone's true valuations is a Nash equilibrium for HOPE-FOR-THE-BEST, the *truthful equilibrium*, that in this equilibrium the revenue is the maximum possible (except for a negligible $n\epsilon$).

Notice too, however, that HOPE-FOR-THE-BEST also has additional equilibria. A second equilibrium is that in which all players report all valuations known to them divided by 2. In such an equilibrium the utility of each player is much greater, but the revenue collected by the mechanism is roughly one half of that of the truthful equilibrium. A third equilibrium is that in which all players report all valuations known to them divided by 3, whose revenue is only a third of that of the truthful equilibrium. And so on. Thus:

> *If in the truthful equilibrium the designer is happy, but not the players, while in all other equilibria the players are happy, but not the designer, which equilibrium is more likely to be selected?*

Although the goal was revenue, these player-preferred equilibria at least maximize social welfare. But this is not the case for other equilibria of HOPE-FOR-THE-BEST. Indeed, for any state ω' such that the true value of every player i is at least $v' > 0$, the strategy profile in which all players report v' the value for ω' and 0 for all other states is an equilibrium.

Given the multitude of available equilibria and the fact that different equilibria are preferable to different players: *Will a play of* HOPE-FOR-THE-BEST *result in equilibrium (and thus any revenue) at all?*

Notice that the *equilibrium selection* problem disappears only if the mechanism is such that $\mathbb{P}$ holds at a strictly dominant-strategy equilibrium. Some form of equilibrium selection is still a problem for weakly dominant-strategy mechanisms.

1.3 Alternative Mechanism Design

In [1], Chen and Micali put forward an alternative notion of a rational play, *surviving strategies*, and advocate leveraging (and indeed benchmarking against) an additional type of knowledge, *external knowledge.*

Clearly, no player i will play a strategy τ_i that is *strictly dominated* by another strategy σ_i. This means that, for any strategy subprofile of the other players, σ_i gives i strictly greater utility than τ_i. This being the case, a rational player i might as well eliminate τ_i from his strategy set. But if i trusts the rationality of his opponents, then he may as well assume that they will eliminate their strictly dominated strategies. This may further enable all players to iteratively eliminate further strategies. Any strategy that cannot be eliminated is called surviving, and a strategy profile is surviving if each of its strategy is surviving. (A game is said to be fully *rationalizable* if this iterative elimination yields a single surviving strategy σ_i for each player. In this case, the profile $\sigma = \sigma_1, \ldots, \sigma_n$ must be an equilibrium and, under the assumption of common rationality, is a very strong prediction of how the game will be played.) A surviving-strategy mechanism is one that guarantees its desired property $\mathbb{P}$ for every profile σ of surviving strategies. Notice that this is different from guaranteeing $\mathbb{P}$ "at equilibrium" in that such a σ in general is not an equilibrium at all. That is, surviving-strategy mechanisms cannot be used to predict accurately how the resulting game will be played, but guarantee that the desired property holds as long as the players are rational, independent of what beliefs they may have.

Clearly again, a player knows his own preferences, but in many contexts he also has external knowledge, that is, he knows valuable information about the preferences of other players as well. Thus, if mechanism design wants to leverage the players' knowledge, it should consider leveraging their external knowledge as well. A player's external knowledge may vary tremendously. At an extreme, in our context, i may know the exact value $v_j(\omega)$ that every player j has for each state ω (so called perfect knowledge). Alternatively, he may know the distribution from which $v_j(\omega)$ has been drawn. Alternatively yet, he may know a finite interval $[a, b]$ to which $v_j(\omega)$ belongs. At another extreme, he may just know that $v_j(\omega)$ is lower-bounded by a (so called *hard knowledge*).

Chen and Micali prove that surviving strategy mechanisms can leverage hard knowledge so as to guarantee new revenue benchmarks in combinatorial auctions. They did not prove, however that such benchmarks could not be guaranteed by more traditional methods. Thus a natural question arises:

Are surviving-strategy mechanism inherently more powerrful than dominant-strategy ones?

1.4 Our Results

We prove a strong separation result between implementation in surviving strategies and implementation in dominant strategies for contexts of non-negative valuations and quasi-linear utilities. Essentially, we prove that, when the players are perfectly informed about each other, then for any positive ϵ:

1. *There exists a surviving-strategy mechanism guaranteeing a fraction* $1 - \epsilon$ *of the optimal revenue;* while
2. *No dominant-strategy mechanism can guarantee more than a fraction* ϵ *of the optimal revenue.*

Above, by "optimal revenue" we mean revenue equal to the maximum social welfare, that is $\max_\omega \sum_i v_i(\omega)$. (If a mechanism generates more revenue, then either some player must be irrational, since he acted to receive a negative utility, or the mechanism must have the power of "taxing players unfairly.")

We stress that our impossibility result applies even to *weakly* dominant-strategy mechanisms. Taken together, therefore our positive and our negative results demonstrates the power of implementation in surviving strategies. Not only can it often be convenient, but can also be the only way to robustly achieve classical properties in classical contexts.

2 Preliminaries

Contexts. In our paper we work with reasonably general contexts with semi-linear utilities. Namely, our context is defined by the following items:

- N, the finite set of *players:* $N = \{1, \ldots, n\}$
- $\Omega \times \mathbb{R}^n$, the set of possible outcomes, where Ω is finite.
  A member ω of Ω is referred to as a *state* and a member P of $\mathbb{R}^n$ is referred to as a *price profile.* Set Ω is required to include the *empty state*, denoted by $\perp$.
- V is the set of all possible profiles of (non-negative) player *types* or *valuations.*
  Each type is a function from the set of states to the set $\mathbb{N}$ of natural numbers mapping $\perp$ to 0.
  We consistently denote by TV the profile of the *true types* (that is, for each player i, TV_i describes i's actual value for each possible state).
- u_i, for each player i, is i's *utility function*, mapping outcomes to real numbers as follows: $u_i(\omega, P) = TV_i(\omega) - P_i$. That is, i's utility is i's true value for the state minus the price he pays.

Accordingly, to specify a context C, it suffices to specify just its "variable" components, that is, $C = (N, \Omega, TV)$. As usual, each player i knows his own type.

We say that a context is *perfect-knowledge* if the entire true-valuation profile TV is common knowledge to all players. (We stress, however, that the mechanism *designer* has no knowledge about TV! In other words, we adhere to the classic spirit of mechanism design, where all knowledge lies with just the players.)

Strategies and Mechanisms. We now must specify the players'strategies, and how these lead to outcomes. Traditionally, thanks to the *revelation principle*, one can restrict attention to mechanisms in which each player, simultaneously with the others, announces a type for himself (which may or may not coincide with his true valuation function). Thus, without loss of generality, a player's set of strategies consists of the set of all possible valuations.

In our case, however, the players do not only know their own types, but also those of the others. And to leverage this extra knowledge, it is crucial that the players be able to announce types for all players. That is, a player's *strategy* consists of a profile of valuations (in other words, it is a member of V). The *empty strategy* is the one whose valuations map every possible state to 0.

A *mechanism* for a context (N, Ω, TV) consists of a (possibly probabilistic) function $M : V^n \to \Omega \times \mathbb{R}^n$ satisfying the following

Opt-Out Condition: For any strategy profile $v = v_1, \ldots, v_n$, if $M(v) = (\omega, P)$ then $P_i = 0$ whenever v_i is the empty strategy.

(By defining our mechanisms in "normal form" we loose no generality. However, to make our mechanisms more intuitive and "communication efficient" we prefer to describe them in extensive-form.)

Social Welfare, Revenue, and Our Goal. The *social welfare* and the *revenue* of an outcome (ω, p) are respectively defined to be $\sum_i TV_i(\omega)$ and $\sum_i p_i$.

The *maximum rational revenue* for a context $C = (N, \Omega, TV)$ is defined to coincide with the *maximum social welfare* (MSW for short), that is, $\max_\omega \sum_i TV_i(\omega)$.

We are interested in designing mechanisms (essentially) guaranteeing the maximum rational revenue.

3 Two Notions of Implementation

A *play* σ of a mechanism M consists of a profile of strategies. If M is probabilistic, then $M(\sigma)$ is a distribution over outcomes, and $u_i(M(\sigma))$ is the expected utility of player i over such distribution, that is, it is short hand for $\mathbb{E}[u_i(M(\sigma))]$.

3.1 Implementation in Dominant Strategies

Dominant Plays. A play σ of a mechanism M is said to be *weakly* dominant if, for each player i, each possible alternative strategy σ'_i for player i, and each possible strategy sub-profile τ_{-i} for the other players, we respectively have (in expectation if M is probabilistic)

$$u_i(M(\sigma_i \sqcup \tau_{-i})) \geq u_i(M(\sigma'_i \sqcup \tau_{-i})).$$

If the above inequality is always strict, σ is a *strictly* dominant play of M.

Implementation in dominant strategies. Let P be a property defined over the outcomes of a context $C = (N, \Omega, v)$. We say that P is *implementable* in weakly (strictly) dominant strategies if there is a mechanism M for C such that P holds for $M(\sigma)$ for each weakly (strictly) dominant play σ of M.

DST mechanisms. The revelation principle guarantees that, whenever a property P is implementable in strictly/weakly dominant strategies, then P is so implementable by a mechanism M^* for which v^n is a strictly/weakly dominant play. Such an M^* is called *dominant-strategy truthful*, or DST for short. Thus, to prove that a property P is not implementable in dominant strategies, it suffices to show that P is not implementable by any DST mechanism.

(Notice that we have generalized the traditional notion of a DST mechanism to "incorporate the players' knowledge in their types.")

3.2 Implementation in Surviving Strategies

Here we simplify the general notion of [1] which was applicable to settings where some players are collusive.

Distinguishable domination. Let Σ be a set of plays of a mechanism M, $\Sigma = \prod_i \Sigma_i$, where each Σ_i is a set of strategies for player i. We say that σ_i is *distinguishably dominated* (by σ'_i) over Σ, if

1. Both σ_i and σ'_i belong to Σ_i
2. $\exists \tau_{-i} \in \Sigma_{-i}$ such that $M(\sigma_i \sqcup \tau_{-i}) \neq M(\sigma'_i \sqcup \tau_{-i})$
(we refer to such a τ_{-i} as a strategy *distinguishing* σ_i and σ'_i over Σ)
3. $u_i(M(\sigma_i \sqcup \tau_{-i})) < u_i(M(\sigma'_i \sqcup \tau_{-i}))$ for any τ_{-i} distinguishing σ_i and σ'_i over Σ.

Surviving Plays. Let Σ_i^0 be the set of all possible strategies of each player i, and let $\Sigma^k = \prod_i \Sigma_i^k$ for any $k \geq 0$. For each player i, Σ_i^{k+1} is the set of strategies in Σ_i^k that are not distinguishably dominated over Σ^k, and $\Sigma_i^\infty \subseteq \Sigma_i^k$ for any $k \geq 0$ is the set of strategies such that no strategy $\sigma_i \in \Sigma_i^\infty$ is distinguishably dominated over $\Sigma^\infty = \prod_i \Sigma_i^\infty$. A play σ is said to be *surviving* if it belongs to Σ^∞.

Implementation in Surviving Strategies. We say that a property P is *implementable in surviving strategies* if there exists a mechanism M such that P holds for $M(\sigma)$ for each surviving play σ of M. We say that P is *implementable in* Σ^k *plays* if it holds for $M(\sigma)$ for each $\sigma \in \Sigma^k$.

3.3 Remarks

The main result of our paper consists of proving that implementation in surviving strategies has wider applicability than implementation in dominant strategies. But it is useful to make right away a few points enabling some simpler comparisons.

1. *Strict Generalization.* To begin with, it should be realized that if a property is implementable in *strict* dominant strategies, then it is implementable in surviving strategies as well. In fact, it is implementable in Σ^1 plays. (More precisely, Σ^1 will consist of a single play.)

   As we shall soon see, however, the converse is not true in a very strong sense. Namely there are properties implementable in surviving strategies that cannot be implemented even in weakly dominant strategies.

2. *"Strict elimination."* It is well known that iterated elimination of strategies is much more meaningful when one eliminates only strictly dominates strategies rather than weakly dominated strategies. In fact, only in the first case the final set of strategies is *not* affected by the order of elimination. Notice that implementation in surviving strategies does *not* call for the iterative elimination of just strictly dominated strategies. (Doing so would severely limit its applicability.) Rather, the notion calls for iteratively eliminate strategies that are distinguishably weaker, so as to maintain wide applicability while guaranteeing the independence of the elimination order. In essence, the set of surviving strategies is unique up to "renaming indistinguishable strategies."[1]

4 Impossibility Result for DST mechanisms

Let us prove that DST mechanisms are incapable of properly leveraging external knowledge: namely, in a perfect-knowledge context, they cannot guarantee even a minuscule fraction of the maximum rational revenue.

[1]Note that if the mechanism is normal-form and the outcome space coincides with the set of all strategies profiles, then every two strategies of the same player are distinguishable. However, things are different when the outcome space is "more limited" or when the mechanism is of extensive form. For instance, consider the case where the game is full information and consists of two moves: first player 1 chooses between Left and Right, and then player 2 does the same. Then, if choosing Left is strictly dominant for player 1, any two strategies of player 2 that differ only when player 1 chooses Right ought to be considered equivalent from a rationality analysis.

Definition 1. *A DST mechanism M guarantees a fraction ϵ of the maximum rational revenue if for any context $C = (N, \Omega, TV)$ we have*

$$(*) \qquad M(TV, \ldots, TV) = (x, P) \text{ implies } \sum P_i \geq \epsilon \cdot MSW.$$

Note that, in proposition $(*)$, each TV is not just the true valuation of a single player, but the profile of all such valuations, because a player's strategy includes his declaration about the others' valuations as well.

Note too that the mechanism is not required to choose the outcome which maximizes the social welfare. Moreover, when not all the players are telling the truth, there is no requirement on the behavior of the mechanism.

Finally note the following immediate corollary of the opt-out condition. Namely,

Non-negative utility property: if M is a DST mechanism and $M(v^1, \ldots, v^n) = (\omega, P)$, then $P_i \leq v_i^i(\omega)$.

Theorem 1. *For any $\epsilon > 0$ no DST mechanism M guarantees a fraction ϵ of the maximum rational revenue.*

Proof. We actually prove our result even for contexts with just two players and only two possible outcomes. Without loss of generality, consider the context (N, Ω, TV) where $N = \{1, 2\}$ and $\Omega = \{\perp, \omega\}$. In such a context, a valuation v_i of a player i coincides with a single number $v_i(\omega)$ (because $v_i(\perp)$ is bound to be 0), and so a strategy v for i coincides with a pair of numbers, $v = (c_1, c_2)$, where c_1 is the declared value for player 1 and c_2 the declared value for player 2.

Our proof is by contradiction. We start by analyzing the behavior of M when the two players make identical and positive (but not necessarily truthful) declarations. More precisely, we prove the following proposition:

$(\star)$ if $c_1, c_2 > 0$, then $M(\ (c_1, c_2)\ ,\ (c_1, c_2)\ ) = (x, (P_1, P_2))$ where

$\star_1$: $P_1 + P_2 \geq \epsilon \cdot (c_1 + c_2)$

$\star_2$: $x = \omega$

To see that proposition $(\star)$ holds, assume the players bid truthfully; that is assume that $c_1 = TV_1(\omega)$ and $c_2 = TV_2(\omega)$. In this case, according to $(*)$ the mechanism must extract a revenue of at least $\epsilon \cdot MSW = \epsilon \cdot (c_1 + c_2)$, and thus $P_1 + P_2 \geq \epsilon \cdot (c_1 + c_2)$, in agreement with inequality $\star_1$.

Now, the hypothesis $c_1 + c_2 > 0$ implies $P_1 + P_2 > 0$. Thus, in light of the non-negative utility property, the state returned by M cannot be $\perp$. Since ω is the only other state, M has to return ω in agreement with equality $\star_2$.

Consider now the declaration $K = (1, 1)$ and let $M(K, K) = (y, Q)$. Then proposition $(\star)$ guarantees that $y = \omega$ and that $Q_1 + Q_2 \geq 2\epsilon$. This implies that $Q_i \geq \epsilon$ for at least a player i. Thus, without loss of generality, we can assume $Q_1 \geq \epsilon$.

Consider now the strategy $\tilde{K} = (\epsilon/2, \epsilon/2)$, and let us analyze the behavior of $M(\tilde{K}, K)$. Let $M(\tilde{K}, K) = (x, P)$.

We start by proving that $x = \omega$. Assume for contradiction purposes that $x = \perp$. Then, when $TV = K$ (and thus player 1 is not truthful), player 2 has an incentive to lie. Indeed, by being truthful, under the current assumption, his utility is 0. However, if player 2 chose the strategy $\tilde{K}$, then according to $(\star)$, the outcome would be (ω, P_1, P_2). In this case, according to the non-negative utility property, since player 2's self-valuation is $\epsilon/2$, $P_2 \leq \epsilon/2$. Thus player 2's utility would be at least $1 - \epsilon/2$. Since this utility is positive, while his utility of being truthful is 0, player 2 has an incentive to lie when $TV = K$ and player 1's strategy is $\tilde{K}$. Therefore we must have $x \neq \perp$, or equivalently $x = \omega$.

Let us now analyze the possible values for P_1 and derive a contradiction in every case.

1. *Case 1:* $P_1 < \epsilon$. In this case, assume that $TV = K$ and compute player 1's utility under the following two strategy profiles: (K, K) and $(\tilde{K}, K)$. In the first case we already know that $M(K, K) = (\omega, Q)$,

Algorithm 1: The global partitioning algorithm with parameters k and δ

```
(π_1, ..., π_n) := random permutation of vertices;
P := ∅;
for i = 1, ..., n do
    if π_i still in the graph then
        if there exists a (k, δ)-isolated neighborhood of π_i in the remaining graph then
            S := this neighborhood
        else
            S := {π_i}
        P := P ∪ {S};
        remove vertices in S from the graph
```

where $Q_1 \geq \epsilon$. Therefore player 1's utility when being truthful is $1 - Q_1$ which is at most $1 - \epsilon$. On the other hand, under the strategy profile $(\tilde{K}, K)$, player 1's utility is equal to $1 - P_1$ and thus strictly greater than $1-\epsilon$ by hypothesis. Thus, the context $(\{1,2\}, \{\perp, \omega\}, K)$ contradicts the dominant-strategy truthfulness of M.

2. *Case 2:* $P_1 > \epsilon/2$. In this case, since $M(\tilde{K}, K) = (\omega, P)$ and $\tilde{K} = (\epsilon/2, \epsilon/2)$, the non-negative utility property implies that $P_1 \leq \epsilon/2$, and thus a contradiction.

In sum, if M guarantees an ϵ fraction of the maximum possible revenue, no price profile exists for $M(\tilde{K}, K)$ that does not contradict the dominant-strategy truthfulness of M. Since we have not assumed any property of M beyond its being DST, this establishes our theorem. *Q.E.D.*

5 Aggregate Knowledge

ϵ and ϵ_1 are two constants such that $0 < 5n\epsilon_1 < \epsilon < 1/5$;

Each player announces his external knowledge $K_i(j)$.

Let

6 Our Mechanism

In the description of our mechanism,

- B is an upperbound of the players' valuations, that is, $v_i(\omega) < B$ for each player i and each state ω;
- ϵ and ϵ_1 are two constants such that $0 < 5n\epsilon_1 < \epsilon < 1/5$;
- numbered steps are taken by the players, while steps marked by letters are taken by the mechanism.

Mechanism $\mathcal{M}$

a. Set $\omega = \perp$, and $P_i = 0$ for each player i.

(ω and P will respectively be the final state and price profile.)

1. Each player i simultaneously and publicly announces "his alleged knowledge" a valuation profile K^i such that for each player j and each state s, $K^i_j(s) < B$.

b. For each $s \in \Omega$ and each player i, define "the maximum knowledge of i about s as follows" $MK_i(s) = \max_{j \neq i} K_i^j(s)$.

   For each $s \in \Omega$, define "the known revenue about s to be" $R_s = \sum_i MK_i(s)$.

   Publicly select the "provisional" state $\omega^\star$ as follows: with probability ϵ, $\omega^\star$ is uniformly selected in Ω, and with complementary probability $\omega^\star = \operatorname{argmax}_{s \in \Omega} R_s$, with ties broken lexicographically.

2. Each player i simultaneously and publicly "suggests a raise of the price for each player, that is," announces a profile of natural numbers Δ^i such that $MK_j(\omega^\star) + \Delta_j^i < B$ for each player j.

c. For each player i, publicly select bip_i "the best informed player about i" and the "provisional" price $P_i^\star$ as follows.

   If the set of players j other than i for which $\Delta_i^j > 0$ is non-empty, then:

   - with probability ϵ, bip_i is uniformly chosen in this set; and with complementary probability $bip_i = \operatorname{argmax}_{j \neq i} \Delta_i^j$ (with ties broken lexicographically);
   - Set $P_i^\star = MK_i(\omega^\star) + \Delta_i^{bip_i}$.

   Else, if the set of players j other than i for which $K_i^j(\omega^\star) > 0$ is non-empty, then:

   - with probability ϵ, bip_i is uniformly chosen among the players $j \neq i$ such that $K_i^j(\omega^\star) > 0$; and with complementary probability $bip_i = \operatorname{argmax}_{j \neq i} K_i^j(\omega^\star)$ (with ties broken lexicographically);
   - Set $P_i^\star = K_i^{bip_i}(\omega^\star)$.

   Else, bip_i is undefined and $P_i^\star = 0$.

   "We refer to $(\omega^\star, P^\star)$ as the *provisional outcome*."

3. Each player i such that $P_i^\star > 0$ simultaneously and publicly announces YES or NO "to the price $P_i^\star - \epsilon_1$." Each player i such that $P_i^\star = 0$ announces YES "by default."

d. Let Y be the number of players announcing YES. Publicly flip two biased coins, c_1 and c_2, such that $\Pr[c_1 = \text{Heads}] = 1 - \frac{\epsilon_1}{B}$ and $\Pr[c_2 = \text{Heads}] = Y/n$.

e. When $c_1 = $ Heads: if $Y = n$, then reset $\omega := \omega^\star$ and reset $P_i := P_i^\star - \epsilon_1$ for each player i; otherwise, for each player i announcing NO, reset $P_{bip_i} := P_{bip_i} + P_i^\star$ "that is, bip_i is punished due to i announcing NO."

f. When $c_1 = $ Tails, rely on c_2 to reset ω and P as follows: if $c_2 = $ Heads, then reset $\omega := \omega^\star$ and reset $P_i := P_i^\star - \epsilon_1$ for each player i; otherwise, ω and P continue to be $\perp$ and 0^n.

g. For each player i, reset $P_i := P_i - \epsilon_1 \left[2 - \frac{1}{1+\sum_{j\neq i}\left(\frac{\Delta_j^i}{2}+K_j^i(\omega^\star)\right)+P_i^\star} \right]$.

   "That is, each player gets a reward which is at least ϵ_1 and at most $2\epsilon_1$."

7 Analysis of Our Mechanism

7.1 Statements of Our Lemmas

Lemma 1. *For all players i and all strategies $\sigma_i \in \Sigma_i^1$: in Step 3, if $P_i^\star > 0$ then*

1. *i announces YES whenever $v_i(\omega^\star) \geq P_i^\star$, and*
2. *i announces NO whenever $v_i(\omega^\star) < P_i^\star$.*

Lemma 2. *For all players i, all strategies $\sigma_i \in \Sigma_i^2$, and all players $j \neq i$: in Step 2,*

1. *if $MK_j(\omega^\star) \geq v_j(\omega^\star)$, then i announces $\Delta_j^i = 0$;*
2. *if $MK_j(\omega^\star) < v_j(\omega^\star)$, then i announces $\Delta_j^i \leq v_j(\omega^\star) - MK_j(\omega^\star)$.*

Lemma 3. *For all players i, all strategies $\sigma_i \in \Sigma_i^3$, and all players $j \neq i$: in Step 2, if $MK_j(\omega^\star) < v_j(\omega^\star)$, then i announces $\Delta_j^i = v_j(\omega^\star) - MK_j(\omega^\star)$.*

Lemma 4. *For all players i, all strategies $\sigma_i \in \Sigma_i^4$, all players $j \neq i$, and all states s: in Step 1, i announces $K_j^i(s) \leq v_j(s)$.*

The next lemma and our main theorem uses the following notation: let $os = \text{argmax}_{o\in\Omega} \sum_l v_l(o)$ with ties broken lexicographically, and refer to it as the *optimal state*, in the sense that it has the maximum social welfare.

Lemma 5. *For all players i, all strategies $\sigma_i \in \Sigma_i^5$, and all players $j \neq i$: in Step 1, i announces $K_j^i(os) = v_j(os)$.*

7.2 Our Main Theorem

Theorem 2. *For all Σ^5 plays σ, let $(\omega, P) = \mathcal{M}(\sigma)$, we have*
(1) $\mathbb{E}[\sum_i P_i] \geq (1-\epsilon)\sum_i v_i(os) - \epsilon$, and
(2) $\mathbb{E}[\sum_i v_i(\omega)] \geq (1-\epsilon)\sum_i v_i(os)$.

Proof. By Lemma 5, for any players l and $k \neq l$, $K_l^k(os) = v_l(os)$, and thus $MK_l(os) = v_l(os)$. Therefore $R_{os} = \sum_l v_l(os)$. By Lemma 4, for any state $s \neq os$, any players l and $k \neq l$, $K_l^k(s) \leq v_l(s)$, and thus $MK_l(s) \leq v_l(s)$. Therefore

$$R_s \leq \sum_l v_l(s) \leq v_l(os) = R_{os},$$

where the second inequality is by definition of os. Because $os = \text{argmax}_{s\in\Omega} \sum_l v_l(s)$, combining with the inequality above, we also have that $os = \text{argmax}_{s\in\Omega} R_s$, with ties broken lexicographically. Accordingly, in Step b of execution σ, $\Pr[\omega^\star = os] > 1 - \epsilon$.

By Lemma 2, when $\omega^\star = os$, for any players l and $k \neq l$, $\Delta_l^k = 0$ since $MK_l(os) = v_l(os)$, and thus no matter what bip_l is, $P_l^\star = K_l^{bip_l}(os) = v_l(os)$. Accordingly, every player announces YES in Step 3, $(\omega^\star, P^\star)$ is implemented when $c_1 = \text{Heads}$, $Y = n$ when $c_1 = \text{Tails}$, and $\Pr[c_2 = \text{Heads}|c_1 = \text{Tails}] = 1$. Therefore before Step g, the expected revenue that the mechanism gets is

$$\begin{aligned}
\mathbb{E}[\sum_l P_l] &\geq \Pr[\omega^\star = os]\mathbb{E}[\sum_l P_l|\omega^\star = os] \\
&> (1-\epsilon)\left[(1-\frac{\epsilon_1}{B})\sum_l (P_l^\star - \epsilon_1) + \frac{\epsilon_1}{B}\cdot 1 \cdot \sum_l (P_l^\star - \epsilon_1)\right] \\
&= (1-\epsilon)\sum_l (P_l^\star - \epsilon_1) = (1-\epsilon)(\sum_l v_l(os) - n\epsilon_1).
\end{aligned}$$

Since the expected reward each player gets in Step g is at most $2\epsilon_1$, the total expected revenue that the mechanism gets is

$$\mathbb{E}[\sum_l P_l] > (1-\epsilon)(\sum_l v_l(os) - n\epsilon_1) - 2n\epsilon_1 = (1-\epsilon)\sum_l v_l(os) - (3-\epsilon)n\epsilon_1 > (1-\epsilon)\sum_l v_l(os) - \epsilon$$

where the last inequality is because $3n\epsilon_1 < \epsilon$, and the first part of Theorem 2 holds.

In the meanwhile, when $\omega^\star = os$, each player l gets value $v_l(\omega) = v_l(\omega^\star) = v_l(os)$ once $(\omega^\star, P^\star)$ is implemented, either because $c_1 = \text{Heads}$ or because $c_2 = \text{Heads}$. Therefore the total expected social welfare the mechanism gets is

$$\mathbb{E}[\sum_l v_l(\omega)] \geq \Pr[\omega^\star = os]\mathbb{E}[\sum_l v_l(\omega)|\omega^\star = os] > (1-\epsilon)\left[(1-\frac{\epsilon_1}{B})\sum_l v_l(os) + \frac{\epsilon_1}{B}\sum_l v_l(os)\right] = (1-\epsilon)\sum_l v_l(os),$$

and the second part of Theorem 2 holds.
Q.E.D.

References

[1] J. Chen and S. Micali. A New Approach to Auctions and Resilient Mechanism Design. STOC'09, full version available at http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design/.

Appendix

A Proofs of Lemmas 1 to 5

Lemma 1. *For all players i and all strategies $\sigma_i \in \Sigma_i^1$: in Step 3, if $P_i^\star > 0$ then*

1. *i announces YES whenever $v_i(\omega^\star) \geq P_i^\star$, and*
2. *i announces NO whenever $v_i(\omega^\star) < P_i^\star$.*

Proof. We focus ourselves on proving, by contradiction, the first implication (the proof of the second one is totally symmetric). Assume there exist a player i and a strategy profile σ such that (1) $\sigma_i \in \Sigma_i^1$ and $\sigma_{-i} \in \Sigma_{-i}^0$; and (2) in the execution of σ, there exists a sequence of coin tosses of the mechanism before Step 3 for which $P_i^\star > 0$, $v_i(\omega^\star) \geq P_i^\star$, and i announces NO in Step 3. Consider the following alternative strategy $\widehat{\sigma}_i$ for i.

Strategy $\widehat{\sigma}_i$

Step 1. Run σ_i and announce K^i as σ_i does.

Step 2. Continue running σ_i and announce Δ^i as σ_i does.

Step 3. If $P_i^\star > 0$ and $v_i(\omega^\star) \geq P_i^\star$, announce YES.
Otherwise, continue running σ_i and answer whatever σ_i does.

We prove that σ_i is distinguishably dominated by $\widehat{\sigma}_i$ over Σ^0, which implies that $\sigma_i \notin \Sigma_i^1$. By definition, we need to show that there exists a strategy subprofile $\tau_{-i} \in \Sigma_{-i}^0$ such that $\mathcal{M}(\sigma_i \sqcup \tau_{-i}) \neq \mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i})$, and that for any such τ_{-i}, $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$. Notice that if execution $\sigma_i \sqcup \tau_{-i}$ is such that for any sequence of coin tosses of the mechanism used before Step 3, either $P_i^\star = 0$, or $v_i(\omega^\star) < P_i^\star$, or i announces YES in Step 3, then the two executions $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ coincides everywhere by construction of $\widehat{\sigma}_i$, and the two outcome distributions $\mathcal{M}(\sigma_i \sqcup \tau_{-i})$ and $\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i})$ are the same. Therefore such a τ_{-i} can not distinguish σ_i and $\widehat{\sigma}_i$ over Σ^0. Accordingly, it suffices for us to consider all strategy subprofiles τ_{-i} such that in execution $\sigma_i \sqcup \tau_{-i}$ for some sequence of coin tosses used before Step 3, $P_i^\star > 0$, $v_i(\omega^\star) \geq P_i^\star$, and i announces NO. Notice that by hypothesis, σ_{-i} is such a strategy subprofile.

Arbitrarily fix such a τ_{-i}. For similar reasons, notice that for all sequences of coin tosses used before Step 3 such that either $P_i^\star = 0$, or $v_i(\omega^\star) < P_i^\star$, or i announces YES, the two executions coincide, and $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ conditioned on such coin tosses. Therefore it suffices for us to consider all sequences of coin tosses used before Step 3 such that $P_i^\star > 0$, $v_i(\omega^\star) \geq P_i^\star$, and i announces NO. Notice that by hypothesis such a sequence of coin tosses exists.

Arbitrarily fix such a sequence of coin tosses. We show that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ conditioned on those coin tosses. By construction of $\widehat{\sigma}_i$, the two executions $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ coincides everywhere before Step 3. Therefore for variables that do not change after Step 3, we can use the same notations ($\omega^\star$, $P^\star$, bip_j, etc) in both executions, without any ambiguity. It follows immediately that:

(1) in execution $\widehat{\sigma}_i \sqcup \tau_{-i}$ we have $P_i^\star > 0$, $v_i(\omega^\star) \geq P_i^\star$, and i announces YES by construction;

(2) for each player $j \neq i$, j announces the same content in Step 3 in the two executions, and thus $\widehat{Y} = Y + 1$ with Y and $\widehat{Y}$ being the number of players announcing YES in $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ respectively; and

(3) player i gets the same amount of reward in Step g in the two executions, because the reward only depends on variables whose values do not change after Step 3, and thus it suffices to compare $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))]$ and $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ before Step g.

Letting (ω, P) and $(\widehat{\omega}, \widehat{P})$ be the final outcome in execution $\sigma_i \sqcup \tau_{-i}$ and execution $\widehat{\sigma}_i \sqcup \tau_{-i}$ respectively, we now distinguish 3 exhaustive events, according to c_1 and c_2.

Event E_1: $c_1 =$ Heads.

Notice that $\Pr[E_1|\sigma_i \sqcup \tau_{-i}] = \Pr[E_1|\widehat{\sigma}_i \sqcup \tau_{-i}] = 1 - \frac{\epsilon_1}{B} > 0$.

When this event occurs, if all players other than i announce YES in Step 3, then (1) in execution $\sigma_i \sqcup \tau_{-i}$ we have that $\omega = \perp$ (due to i announcing NO), $P_i = 0$ (i is not punished as nobody else announces NO), and thus

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_1] = 0;$$

and (2) in execution $\widehat{\sigma}_i \sqcup \tau_{-i}$ we have that $\widehat{\omega} = \omega^\star$, $\widehat{P}_i = P_i^\star - \epsilon_1$ (again i is not punished), and thus

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_1] = v_i(\widehat{\omega}) - \widehat{P}_i = v_i(\omega^\star) - P_i^\star + \epsilon_1 \geq \epsilon_1 > 0,$$

because $v_i(\omega^\star) - P_i^\star$ by hypothesis.

If there exists a player $j \neq i$ announcing NO, then we have that $\omega = \widehat{\omega} = \perp$ and

$$P_i = \widehat{P}_i = \sum_{j:\ bip_j = i,\ j \text{ announces NO}} P_j^\star,$$

which implies that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_1] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_1]$.

Therefore if this event occurs, then no matter what the other players announce, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_1] \leq \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_1],$$

and thus

$$\Pr[E_1|\sigma_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_1] \leq \Pr[E_1|\widehat{\sigma}_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_1]. \tag{1}$$

Event E_2: $c_1 =$ Tails and $c_2 =$ Heads.

Notice that $\Pr[E_2|\sigma_i \sqcup \tau_{-i}] = \frac{\epsilon_1 Y}{Bn}$, and $\Pr[E_2|\widehat{\sigma}_i \sqcup \tau_{-i}] = \frac{\epsilon_1 \widehat{Y}}{Bn} > \Pr[E_2|\sigma_i \sqcup \tau_{-i}]$, because $Y < Y+1 = \widehat{Y}$.

When this event occurs, we have $\omega = \widehat{\omega} = \omega^\star$ and $P_i = \widehat{P}_i = P_i^\star - \epsilon_1$, and thus

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_2] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_2] = v_i(\omega^\star) - P_i^\star + \epsilon_1 \geq \epsilon_1 > 0.$$

Therefore

$$\Pr[E_2|\sigma_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_2] < \Pr[E_2|\widehat{\sigma}_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_2]. \tag{2}$$

Event E_3: $c_1 =$ Tails and $c_2 =$ Tails.

When this event occurs, we have $\omega = \widehat{\omega} = \perp$ and $P_i = \widehat{P}_i = 0$, and thus

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_3] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_3] = 0.$$

Therefore

$$\Pr[E_3|\sigma_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_3] = \Pr[E_3|\widehat{\sigma}_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_3] = 0. \tag{3}$$

By Equations 1, 2, and 3, we have that before Step g,

$$\begin{aligned}
\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] &= \sum_{k=1}^{3} \Pr[E_k|\sigma_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|E_k] \\
&< \sum_{k=1}^{3} \Pr[E_k|\widehat{\sigma}_i \sqcup \tau_{-i}]\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|E_k] \\
&= \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))],
\end{aligned}$$

as desired.

To prove the second implication in Lemma 1, just need to notice that both $v_i(\omega^\star)$ and $P_i^\star$ are integers, and thus $v_i(\omega^\star) < P_i^\star$ implies $v_i(\omega^\star) < P_i^\star - \epsilon_1$. $\square$

Lemma 2. *For all players i, all strategies $\sigma_i \in \Sigma_i^2$, and all players $j \neq i$: in Step 2,*

1. *if $MK_j(\omega^\star) \geq v_j(\omega^\star)$, then i announces $\Delta_j^i = 0$;*
2. *if $MK_j(\omega^\star) < v_j(\omega^\star)$, then i announces $\Delta_j^i \leq v_j(\omega^\star) - MK_j(\omega^\star)$.*

Proof. We focus on proving, by contradiction, the first implication (the proof of the second one is very similar, and will be briefly mentioned at the end of the whole proof). Assume that there exists a player i, a strategy profile σ, and a player $j \neq i$ such that: (1) $\sigma_i \in \Sigma_i^2$ and $\sigma_{-i} \in \Sigma_{-i}^1$; and (2) there exists a sequence of coin tosses of the mechanism used before Step 2, according to which $MK_j(\omega^\star) \geq v_j(\omega^\star)$ and i announces $\Delta_j^i > 0$. Consider the following alternative strategy $\widehat{\sigma}_i$ for i.

Strategy $\widehat{\sigma}_i$

Step 1. Run σ_i and announce K^i as σ_i does.

Step 2. Continue running σ_i and compute Δ^i as σ_i does.
For each player $l \neq j$, announce $\widehat{\Delta}_l^i = \Delta_l^i$.
If $MK_j(\omega^\star) \geq v_j(\omega^\star)$, then announce $\widehat{\Delta}_j^i = 0$.
If $MK_j(\omega^\star) < v_j(\omega^\star)$, then announce $\widehat{\Delta}_j^i = \Delta_j^i$.

Step 3. If $P_i^\star = 0$, announce nothing.
If $P_i^\star > 0$ and $v_i(\omega^\star) \geq P_i^\star$, announce YES.
Otherwise, announce NO.

We prove that σ_i is distinguishably dominated by $\widehat{\sigma}_i$ over Σ^1, which implies that $\sigma_i \notin \Sigma_i^2$. To do so, we first provide the following observations:

O_1: in Step 3 of σ_i, i announces YES or NO consistently with Lemma 1.

O_2: in Step 3 of $\widehat{\sigma}_i$, i announces YES or NO consistently with Lemma 1.

O_3: if a strategy subprofile $\tau_{-i} \in \Sigma_{-i}^1$ is such that in execution $\sigma_i \sqcup \tau_{-i}$, for any sequence of coin tosses used before Step 2, either $MK_j(\omega^\star) < v_j(\omega^\star)$, or $MK_j(\omega^\star) \geq v_j(\omega^\star)$ and i announces $\Delta_j^i = 0$ in Step 2, then $\mathcal{M}(\sigma_i \sqcup \tau_{-i}) = \mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i})$ and such a τ_{-i} does not distinguish σ_i and $\widehat{\sigma}_i$ over Σ^1.

O_4: for any strategy subprofile $\tau_{-i} \in \Sigma_{-i}^1$ and for each player $l \neq i$, in both executions $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$, l announces YES or NO in Step 3 consistently with Lemma 1.

O_1 is because $\sigma_i \in \Sigma_i^1$ by hypothesis $\sigma_i \in \Sigma_i^2$ and the fact that $\Sigma_i^2 \subseteq \Sigma_i^1$; O_2 is by construction of $\widehat{\sigma}_i$; O_3 is because O_1, O_2, and the fact that σ_i and $\widehat{\sigma}_i$ coincides everywhere before Step 3 for such τ_{-i}; and O_4 is by Lemma 1.

According to O_3, it suffices for us to consider all strategy subprofiles $\tau_{-i} \in \Sigma_{-i}^1$ such that in execution $\sigma_i \sqcup \tau_{-i}$, there exists a sequence of coins tossed before Step 2 according to which $MK_j(\omega^\star) \geq v_j(\omega^\star)$ and i announces $\Delta_j^i > 0$ in Step 2. By hypothesis σ_{-i} is such a strategy subprofile.

Arbitrarily fix such a τ_{-i}. Similar to O_3, if a sequence of coins tossed before Step 2 is such that in execution $\sigma_i \sqcup \tau_{-i}$, either $MK_j(\omega^\star) < v_j(\omega^\star)$, or $MK_j(\omega^\star) \geq v_j(\omega^\star)$ and i announces $\Delta_j^i = 0$, then we have that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ conditioned on such coins, because the two executions coincide everywhere. Thus it suffices for us to consider all sequences of coins tossed before Step 2 such that in execution $\sigma_i \sqcup \tau_{-i}$, $MK_j(\omega^\star) \geq v_j(\omega^\star)$ and i announces $\Delta_j^i > 0$. Notice that by hypothesis such a sequence of coins exists.

Arbitrarily fixing such a sequence of coins, we show that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ conditioned on them. Because the two executions coincide before Step 2, for each variable whose value does not change in or after Step 2, we use the same notation in both executions —K^l, $\omega^\star$, $MK_j(\omega^\star)$, etc—, without any ambiguity. For the other variables, we use different notations in the two executions —Δ^i and $\widehat{\Delta}^i$, bip_j and $\widehat{bip}_j$, $P^\star$ and $\widehat{P}^\star$, etc, for $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ respectively—, and it should be clear from the context which execution a notation belongs to.

Because in both executions, the reward that player i gets in Step g is always in $[\epsilon_1, 2\epsilon_1)$, we have that

$$\mathbb{E}\left[\epsilon_1\left(2-\frac{1}{1+\sum_{l\neq i}\left(\frac{\widehat{\Delta}_l^i}{2}+K_l^i(\omega^\star)\right)+\widehat{P}_i^\star}\right)\right]-\mathbb{E}\left[\epsilon_1\left(2-\frac{1}{1+\sum_{l\neq i}\left(\frac{\Delta_l^i}{2}+K_l^i(\omega^\star)\right)+P_i^\star}\right)\right]>-\epsilon_1, \quad (4)$$

and thus it suffices for us to show that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] > \epsilon_1$ before Step g.

To do so, we further notice that for each player $l \neq j$, for each sequence of coins tossed in Step c to set bip_l and $P_l^\star$ in execution $\sigma_i \sqcup \tau_{-i}$, we have that with the same sequence of coins, $\widehat{bip}_l = bip_l$, $\widehat{P}_l^\star = P_l^\star$, and l announces the same content in Step 3 in the two executions. Therefore to simplify the analysis, we assume that all coins tossed in Step c for players $l \neq j$ have been fixed, and show that for any such coins, conditioned on them we have that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] > \epsilon_1$ before Step g.

Recall that after the coins above have been fixed, for each player $l \neq j$ (in particular for $l = i$), $P_l^\star = \widehat{P}_l^\star$ and l announces the same content in Step 3 in the two executions. We define the following variables:

- $u_i^1 = v_i(\omega^\star) - P_i^\star + \epsilon_1$.
  This is the utility that i gets in both executions when $(\omega^\star, P^\star)$ is implemented, either because $c_1 =$ Heads and everybody announces YES, or because $c_1 =$ Tails and $c_2 =$ Heads. Notice that
  $$|u_i^1| < B, \quad (5)$$
  since $0 \leq v_i(\omega^\star) < B$ and $0 \leq P_i^\star < B$, and they are both integers.
- $p_i^1 = \sum_{l\neq j:\ bip_l = i,\ l \text{ announces NO}} P_l^\star$.
  This is the punishment that i pays to the mechanism in both executions due to players other than j announcing NO, when $c_1 =$ Heads. Notice that
  $$p_i^1 \geq 0. \quad (6)$$
- $Y = |\{l : l \neq j,\ l \text{ announces YES}\}|$.
  This is the number of players other than j who announces YES. Notice that
  $$\frac{Y}{n} \leq \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}] \leq \frac{Y+1}{n} \quad \text{and} \quad \frac{Y}{n} \leq \Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] \leq \frac{Y+1}{n},$$
  because the only player whose announcement in Step 3 has not been fixed yet is player j. Therefore
  $$|\Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] - \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}]| \leq \frac{1}{n}. \quad (7)$$
- $p_i^2 = \Pr[bip_j = i|\sigma_i \sqcup \tau_{-i}] \cdot (1 - \frac{\epsilon_1}{B}) \cdot (MK_j(\omega^\star) + \Delta_j^i)$.
  This is the expected punishment that i pays to the mechanism due to j announcing NO in execution $\sigma_i \sqcup \tau_{-i}$ — since $MK_j(\omega^\star) \geq v_j(\omega^\star)$ and $\Delta_j^i > 0$, j announces NO whenever $bip_j = i$. Notice that $\Pr[bip_j = i|\sigma_i \sqcup \tau_{-i}] \geq \frac{\epsilon}{n-1}$ and $MK_j(\omega^\star) + \Delta_j^i \geq 1$. Therefore
  $$p_i^2 \geq \frac{\epsilon}{n-1}(1 - \frac{\epsilon_1}{B}). \quad (8)$$

We now distinguish the following two cases, and show that in each case, $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] > \epsilon_1$, before Step g, and conditioned on all previously mentioned coin tosses being fixed.

Case 1. $\Delta_j^l = 0$ for all $l \neq i, j$.

In this case, in execution $\sigma_i \sqcup \tau_{-i}$, $\Pr[bip_j = i | \sigma_i \sqcup \tau_{-i}] = 1$ by construction of $\mathcal{M}$, and thus j announces NO in Step 3. Therefore $p_i^2 = (1 - \frac{\epsilon_1}{B}) \cdot (MK_j(\omega^\star) + \Delta_j^i)$, and $\omega = \perp$ when $c_1 = \text{Heads}$. Accordingly, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = -(1 - \frac{\epsilon_1}{B})p_i^1 - p_i^2 + \frac{\epsilon_1}{B} \cdot \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}] \cdot u_i^1. \tag{9}$$

In execution $\widehat{\sigma}_i \sqcup \tau_{-i}$, we have that (1) $\Pr[\widehat{bip}_j = i | \widehat{\sigma}_i \sqcup \tau_{-i}] \leq 1$; (2) when $c_1 = \text{Heads}$ and $\widehat{bip}_j = i$, i is punished by at most $K_j^i(\omega^\star)$; and (3) i answers YES if and only if $u_i^1 > 0$ by Lemma 1, and thus does not lose money except the possible punishments he is asked to pay. Accordingly, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] \geq -(1 - \frac{\epsilon_1}{B})p_i^1 - (1 - \frac{\epsilon_1}{B})K_j^i(\omega^\star) + \frac{\epsilon_1}{B} \cdot \Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] \cdot u_i^1. \tag{10}$$

Subtracting Equation 9 from Equation 10, we have that

$$\begin{aligned}
& \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \\
\geq\ & (1 - \frac{\epsilon_1}{B})(MK_j(\omega^\star) + \Delta_j^i - K_j^i(\omega^\star)) + \frac{\epsilon_1 u_i^1}{B} \cdot (\Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] - \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}]) \\
>\ & 1 - \frac{\epsilon_1}{B} - \frac{\epsilon_1 |u_i^1|}{B} \cdot |\Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] - \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}]| \\
>\ & 1 - \frac{\epsilon_1}{B} - \epsilon_1 \cdot \frac{1}{n} > 3n\epsilon_1 - 2\epsilon_1 \geq \epsilon_1,
\end{aligned}$$

where the first inequality is because $p_i^2 = (1 - \frac{\epsilon_1}{B}) \cdot (MK_j(\omega^\star) + \Delta_j^i)$; the second one is because $MK_j(\omega^\star) \geq K_j^i(\omega^\star)$, $\Delta_j^i \geq 1$, and $x \geq -|x|$ with $x = \frac{\epsilon_1 u_i^1}{B} \cdot (\Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] - \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}])$; the third one is because Equations 5 and 7; the last two are because $3n\epsilon_1 < 1$, $B \geq 1$, and $n \geq 1$.

Case 2. There exists $l \neq i, j$ such that $\Delta_j^l > 0$.

In this case, $\widehat{bip}_j \neq i$ with probability 1, and thus i is not punished due to j's announcement in Step 3 in execution $\widehat{\sigma}_i \sqcup \tau_{-i}$. But in both executions, no matter what bip_j and $\widehat{bip}_j$ are, since $MK_j(\omega^\star) \geq v_j(\omega^\star)$, $\Delta_j^{bip_j} > 0$ and $\widehat{\Delta}_j^{\widehat{bip}_j} > 0$, we have that $P_j^\star = MK_j(\omega^\star) + \Delta_j^{bip_j} > v_j(\omega^\star)$ and $\widehat{P}_j^\star = MK_j(\omega^\star) + \widehat{\Delta}_j^{\widehat{bip}_j} > v_j(\omega^\star)$ with probability 1, and thus j always announces NO in both executions. Therefore in execution $\sigma_i \sqcup \tau_{-i}$, it is still true that $\Pr[bip_j = i | \sigma_i \sqcup \tau_{-i}] \geq \frac{\epsilon}{n-1}$ even conditioned on j announcing NO, and Equation 8 still holds.

Accordingly, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = -(1 - \frac{\epsilon_1}{B})p_i^1 + \frac{\epsilon_1}{B} \cdot \Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] \cdot u_i^1,$$

and

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = -(1 - \frac{\epsilon_1}{B})p_i^1 - p_i^2 + \frac{\epsilon_1}{B} \cdot \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}] \cdot u_i^1.$$

Thus

$$\begin{aligned}
& \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \\
=\ & p_i^2 + \frac{\epsilon_1 u_i^1}{B} \cdot (\Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] - \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}]) \\
\geq\ & p_i^2 - \frac{\epsilon_1 |u_i^1|}{B} \cdot |\Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}] - \Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}]| \\
>\ & \frac{\epsilon}{n-1}(1 - \frac{\epsilon_1}{B}) - \epsilon_1 \cdot \frac{1}{n} > \frac{3n\epsilon_1}{n-1}(1 - \frac{\epsilon}{3n}) - \frac{\epsilon_1}{n} \\
=\ & 3\epsilon_1 + \epsilon_1(\frac{3 - \epsilon}{n-1} - \frac{1}{n}) > \epsilon_1,
\end{aligned}$$

where the second inequality is because Equations 5, 7, and 8; and the third one is because $\epsilon > 3n\epsilon_1$ and thus $\frac{\epsilon_1}{B} \leq \epsilon_1 < \frac{\epsilon}{3n}$.

In sum, we have that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] > \epsilon_1$ before Step g. Combining with Equation 4, we have that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] > \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))]$, as desired.

To prove the second implication of Lemma 2, that is, if $MK_j(\omega^\star) < v_j(\omega^\star)$ then i announces $\Delta^i_j \leq v_j(\omega^\star) - MK_j(\omega^\star)$, the alternative strategy $\widehat{\sigma}_i$ becomes to announce $\widehat{\Delta}^i_j = v_j(\omega^\star) - MK_j(\omega^\star)$ in Step 2 when $MK_j(\omega^\star) < v_j(\omega^\star)$, and coincide with σ_i everywhere else. The analysis is similar to the proof of the first implication, except that when $\Delta^l_j > 0$ for some $l \neq i, j$, it is not true that j always announces NO in both executions, and a more careful case analysis is needed. $\square$

Lemma 3. *For all players i, all strategies $\sigma_i \in \Sigma^3_i$, and all players $j \neq i$: in Step 2, if $MK_j(\omega^\star) < v_j(\omega^\star)$, then i announces $\Delta^i_j = v_j(\omega^\star) - MK_j(\omega^\star)$.*

Proof. We proceed by contradiction. Assume that there exist a player i, a strategy profile σ, and a player $j \neq i$ such that: (1) $\sigma_i \in \Sigma^3_i$ and $\sigma_{-i} \in \Sigma^2_{-i}$; (2) there exists a sequence of coins tossed before Step 2, according to which $MK_j(\omega^\star) < v_j(\omega^\star)$, and i announces $\Delta^i_j < v_j(\omega^\star) - MK_j(\omega^\star)$ in Step 2 (by Lemma 2, $\Delta^i_j \leq v_j(\omega^\star) - MK_j(\omega^\star)$ always, and thus we only need to consider the strictly-less case). Consider the following alternative strategy $\widehat{\sigma}_i$ for i.

Strategy $\widehat{\sigma}_i$

Step 1. Run σ_i and announce K^i as σ_i does.

Step 2. Continue running σ_i and compute Δ^i as σ_i does.
For each player $l \neq j$, announce $\widehat{\Delta}^i_l = \Delta^i_l$.
If $MK_j(\omega^\star) \geq v_j(\omega^\star)$, then announce $\widehat{\Delta}^i_j = \Delta^i_j$.
If $MK_j(\omega^\star) < v_j(\omega^\star)$, then announce $\widehat{\Delta}^i_j = v_j(\omega^\star) - MK_j(\omega^\star)$.

Step 3. If $P^\star_i = 0$, announce nothing.
If $P^\star_i > 0$ and $v_i(\omega^\star) \geq P^\star_i$, announce YES.
Otherwise, announce NO.

We prove that σ_i is distinguishably dominated over Σ^2, which implies that $\sigma_i \notin \Sigma^3_i$. To do so, similarly as in Lemma 2, it suffices for us to consider all strategy subprofiles $\tau_{-i} \in \Sigma^2_{-i}$ such that in execution $\sigma_i \sqcup \tau_{-i}$, there exists a sequence of coins tossed before Step 2, according to which $MK_j(\omega^\star) < v_j(\omega^\star)$ and i announces $\Delta^i_j < v_j(\omega^\star) - MK_j(\omega^\star)$ in Step 2. Notice that by hypothesis, such a τ_{-i} and a sequence of coins exist. Arbitrarily fix such a τ_{-i} and such a sequence of coins. We show that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ conditioned on these coins. We adopt the same rule to use variables in the two executions as in Lemma 2.

Because for each player $l \neq i$, l uses the same strategy τ_l in the two executions, and because σ_i and $\widehat{\sigma}_i$ coincide in Step 1, we have that $\Delta^l_i = \widehat{\Delta}^l_i$, and thus the value $P^\star_i$ and $\widehat{P}^\star_i$ have the same distribution — that is, for any non-negative integer $m < B$, $\Pr[P^\star_i = m] = \Pr[\widehat{P}^\star_i = m]$. Further because $\widehat{\Delta}^i_j > \Delta^i_j$ and $\widehat{\Delta}^i_l = \Delta^i_l$ for each $l \neq i, j$, we have that in Step g

$$\mathbb{E}\left[\epsilon_1\left(2 - \frac{1}{1 + \sum_{l \neq i}\left(\frac{\Delta^i_l}{2} + K^i_l(\omega^\star)\right) + P^\star_i}\right)\right] < \mathbb{E}\left[\epsilon_1\left(2 - \frac{1}{1 + \sum_{l \neq i}\left(\frac{\widehat{\Delta}^i_l}{2} + K^i_l(\omega^\star)\right) + \widehat{P}^\star_i}\right)\right], \quad (11)$$

that is, i receives in expectation less reward in Step g in execution $\sigma_i \sqcup \tau_{-i}$ than in $\widehat{\sigma}_i \sqcup \tau_{-i}$. Therefore it suffices to show that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ before Step g, with the coins tossed before Step 2 fixed as mentioned above. Further, as in Lemma 2, it suffices for us to fix the coins tossed for all players other than j in Step 2, and show that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ before Step g, conditioned on these additional coins also.

After all above-mentioned coins being fixed, for each player $l \neq j$, we have that $bip_l = \widehat{bip}_l$, $P^\star_l = \widehat{P}^\star_l$, and l announces the same content in Step 3 in the two executions. Therefore the only difference between the two executions comes from the values of bip_j, $\widehat{bip}_j$, $P^\star_j$, and $\widehat{P}^\star_j$, as well as the announcements of j in Step 3 in the two executions. However, because both $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ belong to Σ^2, by Lemma 2, no matter what bip_j and $\widehat{bip}_j$ are, their announced values for j in Step 2 are at most $v_j(\omega^\star) - MK_j(\omega^\star)$. Therefore $P^\star_j \leq v_j(\omega^\star)$ and $\widehat{P}^\star_j \leq v_j(\omega^\star)$, and by Lemma 1, j announces YES in Step 3 in both executions.

Accordingly, when $c_1 = \text{Heads}$, $\omega = \omega^\star$ if and only if $\widehat{\omega} = \omega^\star$, and thus

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|c_1 = \text{Heads}] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|c_1 = \text{Heads}],$$

since i is asked to pay the same price and the same punishment. When $c_1 = \text{Tails}$, because $Y = \widehat{Y}$, we have that

$$\Pr[c_2 = \text{Heads}|\sigma_i \sqcup \tau_{-i}, c_1 = \text{Tails}] = \Pr[c_2 = \text{Heads}|\widehat{\sigma}_i \sqcup \tau_{-i}, c_1 = \text{Tails}],$$

and

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|c_1 = \text{Tails}, c_2 = \text{Heads}] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|c_1 = \text{Tails}, c_2 = \text{Heads}].$$

Combining the above three equalities, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))],$$

before Step g and conditioned on the fixed coins. Combining with Equation 11, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))],$$

as desired. $\square$

Lemma 4. *For all players i, all strategies $\sigma_i \in \Sigma^4_i$, all players $j \neq i$, and all states o: in Step 1, i announces $K^i_j(o) \leq v_j(o)$.*

Proof. We again proceed by contradiction. Assume there exists a player i, a strategy profile σ, a player $j \neq i$, and a state s such that: (1) $\sigma_i \in \Sigma^4_i$ and $\sigma_{-i} \in \Sigma^3_{-i}$; and (2) in Step 1, i announces $K^i_j(s) > v_j(s)$. Consider the following alternative strategy $\widehat{\sigma}_i$ for i.

Strategy $\widehat{\sigma}_i$

Step 1. Run σ_i and compute K^i as σ_i does.
For each state o and each player j, if $K^i_j(o) \leq v_j(o)$, then announce $\widehat{K}^i_j(o) = K^i_j(o)$;
if $K^i_j(o) > v_j(o)$, then announce $\widehat{K}^i_j(o) = v_j(o)$.

Step 2. For each player l, if $MK_l(\omega^\star) \geq v_l(\omega^\star)$, then announce $\widehat{\Delta}^i_l = 0$;
if $MK_l(\omega^\star) < v_l(\omega^\star)$, then announce $\widehat{\Delta}^i_l = v_l(\omega^\star) - MK_l(\omega^\star)$.

Step 3. If $P^\star_i = 0$, announce nothing.
If $P^\star_i > 0$ and $v_i(\omega^\star) \geq P^\star_i$, announce YES.
Otherwise, announce NO.

We prove that σ_i is distinguishably dominated by $\widehat{\sigma}_i$ over Σ^3, which implies that $\sigma_i \notin \Sigma^4_i$. To do so, we show that for any strategy subprofile $\tau_{-i} \in \Sigma^3_{-i}$, $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$, which actually implies that σ_i is *strictly dominated* by $\widehat{\sigma}_i$ over Σ^3. Because the two executions $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ now differs from the very beginning, for each variable in the mechanism, we use different notations to refer to it in the two executions (K^l and $\widehat{K}^l$, $\omega^\star$ and $\widehat{\omega}^\star$, $P^\star$ and $\widehat{P}^\star$, etc). It should be clear from the context to which execution a notation belongs. Of course, for each player $l \neq i$, we have that $K^l = \widehat{K}^l$, since l uses the same strategy τ_l in the two executions. We have the following two observations:

O_1: in both executions, in Step 3, every player l announces YES or NO consistently with Lemma 1, that is, l announces YES if and only if $v_l(\omega^\star) \geq P^\star_l$ in $\sigma_i \sqcup \tau_{-i}$, and announces YES if and only if $v_l(\widehat{\omega}^\star) \geq \widehat{P}^\star_l$ in $\widehat{\sigma}_i \sqcup \tau_{-i}$.

O_2: every player l announces Δ^l and $\widehat{\Delta}^l$ in Step 2 of the two executions respectively consistently with Lemmas 2 and 3. That is, l announces Δ^l_k (respectively, $\widehat{\Delta}^l_k$) for player k to be 0 if $MK_k(\omega^\star) \geq v_k(\omega^\star)$ (respectively, if $\widehat{MK}_k(\widehat{\omega}^\star) \geq v_k(\widehat{\omega}^\star)$), and $v_k(\omega^\star) - MK_k(\omega^\star)$ (respectively, $v_k(\widehat{\omega}^\star) - \widehat{MK}_k(\widehat{\omega}^\star)$) otherwise.

O_3: $R_o \geq \widehat{R}_o$ for each state o.

Here O_1 and O_2 are because both $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ belong to Σ^3, and thus belong to Σ^2 and Σ^1 as well; and O_3 is because for each player $l \neq i$, l announces the same valuation profile in both executions, while the values announced by player i for each player k and each state o only decreases from $\sigma_i \sqcup \tau_{-i}$ to $\widehat{\sigma}_i \sqcup \tau_{-i}$.

In execution $\sigma_i \sqcup \tau_{-i}$, let state om be the lexicographically first state o' such that $R_{o'} = \max_o R_o$. We distinguish two cases, according to om.

Case 1. $K^i_l(om) \leq v_l(om)$ for each player l. That is, i does not overbid about l's value on om.

In this case, in execution $\widehat{\sigma}_i \sqcup \tau_{-i}$, om is also the lexicographically first state o' such that $\widehat{R}_{o'} = \max_o \widehat{R}_o$. Because on one hand, by construction of $\widehat{\sigma}_i$, $\widehat{K}^i_l(om) = K^i_l(om)$ for each player l, and thus $MK_l(om) = \widehat{MK}_l(om)$ for each l, which implies $R_{om} = \widehat{R}_{om}$; and on the other hand, $R_o \geq \widehat{R}_o$ for each state $o \neq om$, by O_3.

Accordingly, in execution $\sigma_i \sqcup \tau_{-i}$ (respectively, $\widehat{\sigma}_i \sqcup \tau_{-i}$), with probability $1-\epsilon$, $\omega^\star$ (respectively, $\widehat{\omega}^\star$) is set to be om; and with probability ϵ, $\omega^\star$ (respectively, $\widehat{\omega}^\star$) is set to be a random state in Ω. That is,

$$\Pr[\omega^\star = o] = \Pr[\widehat{\omega}^\star = o] > 0 \tag{12}$$

for each state o. Since

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \sum_{o \in \Omega} \Pr[\omega^\star = o] \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|\omega^\star = o]$$

and

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = \sum_{o \in \Omega} \Pr[\widehat{\omega}^\star = o] \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|\widehat{\omega}^\star = o],$$

combining with Equation 12, we have that

$$\begin{aligned} & \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \\ = & \sum_{o \in \Omega} \Pr[\omega^\star = o] (\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|\widehat{\omega}^\star = o] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|\omega^\star = o]). \end{aligned} \tag{13}$$

Therefore to prove $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] > \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))]$, it suffices for us to prove the following two claims:

Claim 1. for each state o such that $K^i_l(o) \leq v_l(o)$ for all players $l \neq i$, we have that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|\widehat{\omega}^\star = o] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|\omega^\star = o] = 0$; and

Claim 2. for each state o such that $K^i_l(o) > v_l(o)$ for some player $l \neq i$, we have that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|\widehat{\omega}^\star = o] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|\omega^\star = o] > 0$.

By our assumption at the very beginning of the proof of this lemma, there exists a state and a player, namely state s and player j, satisfying the hypothesis in Claim 2. It is easy to see that the validity of Claims 1 and 2, together with the existence of state s and player j, and with Equations 12 and 13, implies $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] > \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))]$.

The validity of Claim 1 is easy to verify. Notice that for each state o satisfying its hypothesis, it must be the case that $K^i_l(o) = \widehat{K}^i_l(o)$ for each player $l \neq i$. Since $K^l = \widehat{K}^l$ for each $l \neq i$, we have that $\widehat{\Delta}^l = \Delta^l$ for each player l (because l's announcement in Step 2 must be consistent with Lemmas 2 and 3, and there is only one way to do so). Therefore the two executions coincide everywhere after Step b, including the reward that each player receives in Step g, and we have that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))]$.

To prove Claim 2, we focus on proving $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|\widehat{\omega}^\star = s] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|\omega^\star = s] > 0$, since for any other state o satisfying the hypothesis of Claim 2, the proof is exactly the same. Similar to Equation 4 in the proof of Lemma 2, we have that in Step g

$$\begin{aligned}&\mathbb{E}\left[\epsilon_1\left(2-\frac{1}{1+\sum_{l\neq i}\left(\frac{\widehat{\Delta}_l^i}{2}+\widehat{K}_l^i(s)\right)+\widehat{P}_i^\star}\right)\middle|\widehat{\omega}^\star=s\right]\\&-\mathbb{E}\left[\epsilon_1\left(2-\frac{1}{1+\sum_{l\neq i}\left(\frac{\Delta_l^i}{2}+K_l^i(s)\right)+P_i^\star}\right)\middle|\omega^\star=s\right]>-\epsilon_1.\end{aligned}\tag{14}$$

Letting $U_i = u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))$ and $\widehat{U}_i = u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))$, it suffices for us to show that

$$\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s] - \mathbb{E}[U_i|\omega^\star = s] > \epsilon_1 \tag{15}$$

before Step g. From now on, unless explicitly specified, all expected utilities talked about below are computed before Step g. Because

$$\begin{aligned}\mathbb{E}[U_i|\omega^\star = s] &= \Pr[c_1 = \text{Heads}|\sigma_i \sqcup \tau_{-i}]\mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Heads}]\\&\quad + \Pr[c_1 = \text{Tails}|\sigma_i \sqcup \tau_{-i}]\mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Tails}]\\&= (1-\frac{\epsilon_1}{B})\mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Heads}] + \frac{\epsilon_1}{B}\cdot\mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Tails}]\end{aligned}$$

and similarly

$$\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s] = (1-\frac{\epsilon_1}{B})\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Heads}] + \frac{\epsilon_1}{B}\cdot\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Tails}],$$

to prove Equation 15, it suffices to show that

$$\frac{\epsilon_1}{B}\left(\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Tails}] - \mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Tails}]\right) > -\epsilon_1, \tag{16}$$

and

$$(1-\frac{\epsilon_1}{B})\left(\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Heads}] - \mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Heads}]\right) > 2\epsilon_1. \tag{17}$$

Notice that in execution $\sigma_i \sqcup \tau_{-i}$, for each player $l \neq i$, l announces Δ_i^l according to Lemmas 2 and 3, therefore no matter what bip_i is, we have that $0 \leq v_i(s) \leq P_i^\star < B$, and thus

$$-(B-1) \leq \mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Tails}] \leq \epsilon_1.$$

Similarly, $-(B-1) \leq \mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Tails}] \leq \epsilon_1$. Combining these two inequalities, we have that

$$\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Tails}] - \mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Tails}] \geq -(B-1) - \epsilon_1 > -B,$$

and Equation 16 holds.

We now prove Equation 17. Notice that in execution $\widehat{\sigma}_i \sqcup \tau_{-i}$, since player i does not overbid on any other player, i is never punished due to the others announcing NO in Step 3. Thus

$$\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Heads}] \geq 0. \tag{18}$$

On the other hand, in execution $\sigma_i \sqcup \tau_{-i}$, $K_j^i(s) > v_j(s)$ by assumption, therefore $MK_j(s) > v_j(s)$, and every player $l \neq j$ (including i himself) announces $\Delta_j^l = 0$ in Step 2. Accordingly, $\Pr[bip_j = i|\omega^\star = s] \geq \frac{\epsilon}{n-1}$, and when $bip_j = i$ and $c_1 = \text{Heads}$, i is punished by at least $P_j^\star = K_j^i(s) \geq 1$. Since $v_i(s) \leq P_i^\star$

as said when proving Equation 16, $v_i(s) \le P_i^\star$ which implies that $v_i(s) - (P_i^\star - \epsilon_1) \le \epsilon_1$ and that when $bip_j \ne i$ the utility i can get when $c_1 = \text{Heads}$ is at most ϵ_1. In sum,

$$\mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Heads}] \le (1 - \frac{\epsilon}{n-1})\epsilon_1 - \frac{\epsilon}{n-1} < \epsilon_1 - \frac{\epsilon}{n-1}. \tag{19}$$

Combining Equations 18 and 19, we have that

$$\begin{aligned} & (1 - \frac{\epsilon_1}{B}) \left(\mathbb{E}[\widehat{U}_i|\widehat{\omega}^\star = s, c_1 = \text{Heads}] - \mathbb{E}[U_i|\omega^\star = s, c_1 = \text{Heads}]\right) \ge (1 - \frac{\epsilon_1}{B})(\frac{\epsilon}{n-1} - \epsilon_1) \\ > \; & \frac{15}{16}(4\epsilon_1 - \epsilon_1) = \frac{45\epsilon_1}{16} > 2\epsilon_1, \end{aligned}$$

where the second inequality is because $4n\epsilon_1 < \epsilon < 1/4$ and $B \ge 1$. That is, Equation 17 holds. Combining Equations 16 and 17, we know that Equation 15 holds. Further combining with Equation 14, we know that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))|\widehat{\omega}^\star = s] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))|\omega^\star = s] > 0$ with Step g included, and Claim 2 holds.

In sum, we have that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ in Case 1.

Case 2. There exists some player $k \ne i$ such that $K_k^i(om) > v_k(om)$. That is, i overbids about $v_k(om)$.

In this case, similar to Equation 14, we have that in Step g,

$$\begin{aligned} & \mathbb{E}\left[\epsilon_1\left(2 - \frac{1}{1 + \sum_{l\ne i}\left(\frac{\widehat{\Delta}_l^i}{2} + \widehat{K}_l^i(\widehat{\omega}^\star)\right) + \widehat{P}_i^\star}\right)\right] \\ & -\mathbb{E}\left[\epsilon_1\left(2 - \frac{1}{1 + \sum_{l\ne i}\left(\frac{\Delta_l^i}{2} + K_l^i(\omega^\star)\right) + P_i^\star}\right)\right] > -\epsilon_1. \end{aligned} \tag{20}$$

Similar to Equation 16, we have that before Step g,

$$\frac{\epsilon_1}{B}\left(\mathbb{E}[\widehat{U}_i|c_1 = \text{Tails}] - \mathbb{E}[U_i|c_1 = \text{Tails}]\right) > -\epsilon_1. \tag{21}$$

Therefore to prove $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] > 0$, it suffices for us to prove the following equation similar to Equation 17:

$$(1 - \frac{\epsilon_1}{B})\left(\mathbb{E}[\widehat{U}_i|c_1 = \text{Heads}] - \mathbb{E}[U_i|c_1 = \text{Heads}]\right) > 2\epsilon_1. \tag{22}$$

Similar to Equation 18, we have that

$$\mathbb{E}[\widehat{U}_i|c_1 = \text{Heads}] \ge 0. \tag{23}$$

Now we upperbound $\mathbb{E}[U_i|c_1 = \text{Heads}]$. In execution $\sigma_i \sqcup \tau_{-i}$, because om is the lexicographically first state o' such that $R_{o'} = \max_o R_o$,

$$\Pr[\omega^\star = om|\sigma_i \sqcup \tau_{-i}] > 1 - \epsilon.$$

Given that $\omega^\star = om$, since $K_k^i(om) > v_k(om)$, by Lemma 2, every player $l \ne k$ (including i himself) announces $\Delta_k^l = 0$ in Step 2. Thus

$$\Pr[bip_k = i|\omega^\star = om] > \frac{\epsilon}{n-1}.$$

Since k announces NO in Step 3 when $bip_k = i$, i is punished by at least $K_k^i(om) \geq 1$ when $bip_k = i$ and $c_1 =$ Heads. Similar to Case 1, since every player $l \neq i$ announces Δ_i^l consistently with Lemmas 2 and 3, when $\omega^\star \neq om$ or $bip_k \neq i$, the utility i can get when $c_i =$ Heads is at most ϵ_1. In sum,

$$\mathbb{E}[U_i | c_1 = \text{Heads}] < \epsilon_1 - (1-\epsilon) \cdot \frac{\epsilon}{n-1}. \tag{24}$$

Combining Equations 23 and 24, we have that

$$\begin{aligned} & (1 - \frac{\epsilon_1}{B}) \left(\mathbb{E}[\widehat{U}_i | c_1 = \text{Heads}] - \mathbb{E}[U_i | c_1 = \text{Heads}] \right) \\ > \;& (1 - \frac{\epsilon_1}{B}) \left((1-\epsilon) \cdot \frac{\epsilon}{n-1} - \epsilon_1 \right) > \frac{24}{25} \cdot (\frac{4}{5} \cdot 5\epsilon_1 - \epsilon_1) = \frac{72\epsilon_1}{25} > 2\epsilon_1, \end{aligned}$$

where the second inequality is because $5n\epsilon_1 < \epsilon < 1/5$ and $B \geq 1$. That is, Equation 22 holds. Combining Equations 20, 21, and 22, we have that $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] - \mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] > 0$, i.e., $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ in Case 2.

Summarizing Cases 1 and 2, $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ for any $\tau_{-i} \in \Sigma_{-i}^3$, as desired. $\square$

The next lemma uses the following notation: let $os = \text{argmax}_{o \in \Omega} \sum_l v_l(o)$ with ties broken lexicographically, os is called the *optimal state*, in the sense that it has the maximum social welfare.

Lemma 5. *For all players i, all strategies $\sigma_i \in \Sigma_i^5$, and all players $j \neq i$: in Step 1, i announces $K_j^i(os) = v_j(os)$.*

Proof. We again proceed by contradiction. Assume there exists a player i, a strategy $\sigma_i \in \Sigma_i^5$, and a player $j \neq i$, such that i announces $K_j^i(os) < v_j(os)$ in Step 1 (by Lemma 4, $K_j^i(os) \leq v_j(os)$ always). Consider the following alternative strategy $\widehat{\sigma}_i$ for i.

Strategy $\widehat{\sigma}_i$

Step 1. Run σ_i and compute K^i as σ_i does.
For each state $o \neq os$ and each player j, announce $\widehat{K}_j^i(o) = K_j^i(o)$.
For each player j, announce $\widehat{K}_j^i(os) = v_j(os)$.

Step 2. For each player l, if $MK_l(\omega^\star) \geq v_l(\omega^\star)$, then announce $\widehat{\Delta}_l^i = 0$;
if $MK_l(\omega^\star) < v_l(\omega^\star)$, then announce $\widehat{\Delta}_l^i = v_l(\omega^\star) - MK_l(\omega^\star)$.

Step 3. If $P_i^\star = 0$, announce nothing.
If $P_i^\star > 0$ and $v_i(\omega^\star) \geq P_i^\star$, announce YES.
Otherwise, announce NO.

We prove that σ_i is distinguishably dominated by $\widehat{\sigma}_i$ over Σ^4, which implies that $\sigma_i \notin \Sigma_i^5$. To do so, for each strategy subprofile $\tau_{-i} \in \Sigma_{-i}^4$, letting $U_i = u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))$ and $\widehat{U}_i = u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))$, we show that $\mathbb{E}[U_i] < \mathbb{E}[\widehat{U}_i]$, which actually implies that σ_i is *strictly dominated* by $\widehat{\sigma}_i$ over Σ^4.

Arbitrarily fix such a τ_{-i}. Similar to the proof of Lemma 4, for each variable in the mechanism, we refer to it using different notations in the two executions $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ (K^l and $\widehat{K}^l$, $\omega^\star$ and $\widehat{\omega}^\star$, $P^\star$ and $\widehat{P}^\star$, etc). It should be clear from the context to which execution a notation belongs. We have the following observations:

O_1: for each player $l \neq i$, the announcements of l in Step 1 are the same in the two executions, that is, $K^l = \widehat{K}^l$;

O_2: for each state $o \neq os$ and each player l, $\widehat{K}_l^i(o) = K_l^i(o)$; and for each player l, if $K_l^i(os) = v_l(os)$ then $K_l^i(os) = \widehat{K}_l^i(os)$, otherwise $K_l^i(os) < \widehat{K}_l^i(os)$.

O_3: for each state $o \neq os$, $R_o = \widehat{R}_o$; and $R_{os} \leq \widehat{R}_{os}$.

O_4: in execution $\sigma_i \sqcup \tau_{-i}$ (respectively, $\widehat{\sigma}_i \sqcup \tau_{-i}$), for each player l, each state o, and each player $k \neq l$, $K^l_k(o) \leq v_k(o)$ (respectively, $\widehat{K}^l_k(o) \leq v_k(o)$) in Step 1;

O_5: in execution $\sigma_i \sqcup \tau_{-i}$ (respectively, $\widehat{\sigma}_i \sqcup \tau_{-i}$), for each player l and each player $k \neq l$, l announces Δ^l_k (respectively, $\widehat{\Delta}^l_k$) in Step 2 such that $MK_k(\omega^\star) + \Delta^l_k = v_l(\omega^\star)$ (respectively, $\widehat{MK}_k(\widehat{\omega}^\star) + \widehat{\Delta}^l_k = v_l(\widehat{\omega}^\star)$).

O_6: in execution $\sigma_i \sqcup \tau_{-i}$ (respectively, $\widehat{\sigma}_i \sqcup \tau_{-i}$), for each player l, $P^\star_l = v_l(\omega^\star)$ (respectively, $\widehat{P}^\star_l = v_l(\widehat{\omega}^\star)$), and l always announces YES in Step 3, no matter what bip_l (respectively, $\widehat{bip_l}$) is.

Here O_1 is because l uses the same strategy τ_l in the two executions; O_2 is by construction of $\widehat{\sigma}_i$; O_3 is by construction of $\widehat{\sigma}_i$ and by definition of R_o for each state o; O_4 is because Lemma 4 and because both $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$ belong to Σ^4; O_5 is because O_4 and Lemmas 3 and 4; O_6 is because O_5 and Lemma 1.

By O_6, we have that: (1) in execution $\sigma_i \sqcup \tau_{-i}$ (respectively, $\widehat{\sigma}_i \sqcup \tau_{-i}$), when $c_1 =$ Heads, $(\omega^\star, P^\star)$ (respectively, $(\widehat{\omega}^\star, \widehat{P}^\star)$) is implemented, and each player gets utility ϵ_1 in Step e; and (2) $\Pr[c_2 = \text{Heads}|c_1 = \text{Tails}, \sigma_i \sqcup \tau_{-i}] = \Pr[c_2 = \text{Heads}|c_1 = \text{Tails}, \widehat{\sigma}_i \sqcup \tau_{-i}] = 1$, and each player gets utility ϵ_1 in Step f. Thus before Step g, we have that

$$\mathbb{E}[U_i] = \mathbb{E}[\widehat{U}_i] = (1 - \frac{\epsilon_1}{B})\epsilon_1 + \frac{\epsilon_1}{B} \cdot 1 \cdot \epsilon_1 = \epsilon_1. \tag{25}$$

Let $r_i = \epsilon_1 \left(2 - \frac{1}{1 + \sum_{l \neq i}\left(\frac{\Delta^i_l}{2} + K^i_l(\omega^\star)\right) + P^\star_i}\right)$ and $\widehat{r}_i = \epsilon_1 \left(2 - \frac{1}{1 + \sum_{l \neq i}\left(\frac{\widehat{\Delta}^i_l}{2} + \widehat{K}^i_l(\widehat{\omega}^\star)\right) + \widehat{P}^\star_i}\right)$ be the reward i gets in Step g in the two executions. By Equation 25, to show that $\mathbb{E}[U_i] < \mathbb{E}[\widehat{U}_i]$, it suffices for us to show $\mathbb{E}[r_i] < \mathbb{E}[\widehat{r}_i]$, which is equivalent to show

$$\mathbb{E}\left[\sum_{l \neq i}\left(\frac{\Delta^i_l}{2} + K^i_l(\omega^\star)\right) + P^\star_i\right] < \mathbb{E}\left[\sum_{l \neq i}\left(\frac{\widehat{\Delta}^i_l}{2} + \widehat{K}^i_l(\widehat{\omega}^\star)\right) + \widehat{P}^\star_i\right].$$

By O_6, $P^\star_i = v_i(\omega^\star)$ and $\widehat{P}^\star_i = v_i(\widehat{\omega}^\star)$. Accordingly, the above inequality is equivalent to

$$\mathbb{E}\left[\sum_{l \neq i}\left(\frac{\Delta^i_l}{2} + K^i_l(\omega^\star)\right) + v_i(\omega^\star)\right] < \mathbb{E}\left[\sum_{l \neq i}\left(\frac{\widehat{\Delta}^i_l}{2} + \widehat{K}^i_l(\widehat{\omega}^\star)\right) + v_i(\widehat{\omega}^\star)\right]. \tag{26}$$

Let $om = \text{argmax}_{o \in \Omega} R_o$ and $\widehat{om} = \text{argmax}_{o \in \Omega} \widehat{R}_o$, with ties broken lexicographically. We now distinguish four cases, according to om and $\widehat{om}$.

Case 1. $om = os$ and $\widehat{om} = os$.

In this case, for each state o,

$$\Pr[\omega^\star = o] = \Pr[\widehat{\omega}^\star = o] > 0, \tag{27}$$

because both equals os with probability $1 - \epsilon$, and equals a random state in Ω with probability ϵ.

For each state $o \neq os$, we have that

$$\mathbb{E}\left[\sum_{l \neq i}\left(\frac{\Delta^i_l}{2} + K^i_l(o)\right) + v_i(o) \middle| \omega^\star = o\right] = \mathbb{E}\left[\sum_{l \neq i}\left(\frac{\widehat{\Delta}^i_l}{2} + \widehat{K}^i_l(o)\right) + v_i(o) \middle| \widehat{\omega}^\star = o\right], \tag{28}$$

because (1) $K^k_l(o) = \widehat{K}^k_l(o)$ for each player k and each player l, by O_1 and O_2; (2) $MK_l(o) = \widehat{MK}_l(o)$ for each player l, by (1) and the definition of $MK_l(o)$; and (3) $\Delta^i_l = \widehat{\Delta}^i_l$ for each player $l \neq i$, by (2) and Lemmas 2 and 3.

For state os, by construction of $\widehat{\sigma}_i$, $\widehat{K}^i_l(os) = v_l(os)$ for each player $l \neq i$. By Lemma 4, no player overbids on the others' values on os, and thus $\widehat{MK}_l(os) = v_l(os)$ also, which implies $\widehat{\Delta}^i_l = 0$. Accordingly, we have that

$$\mathbb{E}\left[\sum_{l \neq i}\left(\frac{\widehat{\Delta}^i_l}{2} + \widehat{K}^i_l(os)\right) + v_i(os) \middle| \widehat{\omega}^\star = os\right] = \mathbb{E}\left[\sum_{l \neq i} v_l(os) + v_i(os)\right] = \sum_l v_l(os). \tag{29}$$

On the other hand, when $\omega^\star = os$, we have that (1) for each player $l \neq i, j$, $\frac{\Delta_l^i}{2} + K_l^i(os) \leq \Delta_l^i + MK_l(os) = v_l(os)$; and (2) $\frac{\Delta_j^i}{2} + K_j^i(os) < \Delta_j^i + MK_j(os) = v_j(os)$, since $K_j^i(os) < v_j(os)$ implies that either $K_j^i(os) < MK_j(os)$, or $0 < \frac{\Delta_j^i}{2} < \Delta_j^i$. Therefore

$$\mathbb{E}\left[\sum_{l\neq i}\left(\frac{\Delta_l^i}{2} + K_l^i(os)\right) + v_i(os)\middle| \omega^\star = os\right] < \sum_{l\neq i,j} v_l(os) + v_j(os) + v_i(os) = \sum_l v_l(os). \tag{30}$$

Combining Equations 29 and 30, we have that

$$\mathbb{E}\left[\sum_{l\neq i}\left(\frac{\Delta_l^i}{2} + K_l^i(os)\right) + v_i(os)\middle| \omega^\star = os\right] < \mathbb{E}\left[\sum_{l\neq i}\left(\frac{\widehat{\Delta}_l^i}{2} + \widehat{K}_l^i(os)\right) + v_i(os)\middle| \widehat{\omega}^\star = os\right]. \tag{31}$$

Combining Equations 27, 28, and 31, we conclude that Equation 26 holds, which implies $\mathbb{E}[r_i] < \mathbb{E}[\widehat{r}_i]$. Further combining with Equation 25, we have that $\mathbb{E}[U_i] < \mathbb{E}[\widehat{U}_i]$ in Case 1.

Case 2. $om \neq os$ and *$\widehat{om} \neq os$.*

In this case, we have that $om = \widehat{om}$, by O_3. Therefore similar to Case 1, $\Pr[\omega^\star = o] = \Pr[\widehat{\omega}^\star = o] > 0$. By the same reasons as in Case 1, Equations 28 and 31 also hold here, and we have that $\mathbb{E}[U_i] < \mathbb{E}[\widehat{U}_i]$ in Case 2.

Case 3. $om \neq os$ and *$\widehat{om} = os$.*

In this case, notice that Equations 28 and 31 still hold, by the same reasons as in Case 1. In addition, we have that

$$\begin{aligned}
& \mathbb{E}\left[\sum_{l\neq i}\left(\frac{\Delta_l^i}{2} + K_l^i(om)\right) + v_i(om)\middle| \omega^\star = om\right] \\
\leq\ & \mathbb{E}\left[\sum_{l\neq i}(\Delta_l^i + MK_l(om)) + v_i(om)|\omega^\star = om\right] = \sum_l v_l(om) \\
\leq\ & \sum_l v_l(os) = \mathbb{E}\left[\sum_{l\neq i}\left(\frac{\widehat{\Delta}_l^i}{2} + \widehat{K}_l^i(\widehat{om})\right) + v_i(\widehat{om})\middle| \widehat{\omega}^\star = \widehat{om}\right].
\end{aligned}$$

where the second inequality is by definition of os.

Because $\Pr[\omega^\star = o|\omega^\star$ is uniformly chosen$] = \Pr[\widehat{\omega}^\star = o|\widehat{\omega}^\star$ is uniformly chosen$]$ for each state o, we have that Equation 26 also holds here, and $\mathbb{E}[U_i] < \mathbb{E}[\widehat{U}_i]$ in Case 3.

Case 4. $om = os$ and *$\widehat{om} \neq os$.*

Fortunately, this case can never happen according to O_3.

Summarizing all cases, we have that $\mathbb{E}[U_i] < \mathbb{E}[\widehat{U}_i]$. □