Resilient and Virtually 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

January 13, 2010

Abstract

We put forward a new extensive-form mechanism that, in a general context with perfectly informed players and quasi-linear utilities,

- Virtually achieves optimal revenue at a unique subgame-perfect equilibrium;
- Is perfectly resilient to the problems of collusion, complexity, and privacy; and
- Works for any number of players $n > 1$.

1 Introduction

1.1 Our Goal

We consider contexts with quasi-linear utilities, non-negative valuations, and perfect knowledge. Informally, these are contexts where there are finitely many *states*; each player i's has a a true type θ_i specifying a value $\theta(\omega) \geq 0$ for each possible state ω; the true-type profile θ is common knowledge among the players; an outcome is a pair (ω, P) where ω is a state and P a profile of prices (reals); and player i's utility in such an outcome is $\theta_i(\omega) - P_i$.

In such contexts, without any information about the players, we wish to find a mechanism yielding an outcome (ω, P) with maximum revenue, that is, such that $\sum_i P_i$ is maximum. If the players are rational and cannot be forced to accept negative utilities, such maximum revenue is upperbounded by the maximum social welfare, that is the value

$$MSW = \max_{\omega \in \Omega} \sum_i \theta_i(\omega).$$

But since the players must be incentivized to participate in the mechanism, we are happy to "virtually" achieve optimal revenue. That is, for any given $\epsilon > 0$, we seek a mechanism producing an outcome (ω, P) such that

$$\sum_i P_i \geq (1 - \epsilon) MSW.$$

This indeed is a classical goal, and several mechanisms have been proposed for achieving it. But we are interested in achieving virtually optimal revenue in a much more **robust** way than considered so far: namely, *by means of mechanisms* resilient *to the problems of equilibrium selection, collusion, complexity, and privacy.* Let us explain.

The Problem of Equilibrium Selection (and Equilibrium Failure) Some mechanisms achieve a desired social-choice property $\mathbb{P}$ "at a Nash Equilibrium." Such mechanisms provide only a weak guarantee that $\mathbb{P}$ might hold for the final outcome of the game. This is so because there may be several Nash equilibria, while $\mathbb{P}$ holds for just some of them. Furthermore, if such mechanisms were of normal form, then the existence of multiple equilibria might actually make it unlikely for the players to reach any equilibrium.[1] Thus, even if a given social-choice property $\mathbb{P}$ holds at each possible equilibrium, $\mathbb{P}$ may not hold in a real play.

The Problem of Collusion In equilibrium-based mechanism design, the problem of collusion often arises for quite natural reasons. Any equilibrium, even a dominant-strategy one, only guarantees that no *single* player has incentive to deviate from his strategy. However, two or more players may have all the incentive in the world to *jointly* deviate from their equilibrium strategies. Accordingly, the actual profile of strategies played out may not be an equilibrium at all.[2]

The Problem of Complexity The execution of any mechanism relies on two main resources: *communication* and *computation.* The first resource relates to the number of bits exchanged, and the latter to the number of elementary computations performed. (Note that "communication lowerbounds computation" because any bit sent must, at least, be read.) Often, however, the required amount of these resources is exponential in some relevant parameters (such as the number of players). And when this is the case, a mechanism will in practice fail to reach its objective, except when played in a "tiny" context. Else, designers and players alike will die way before the final outcome will be computed.

The Problem of Privacy By definition, an outcome resulting from the rational play of a mechanism betrays some information about the players' true types. We regard this "loss of privacy" as *intrinsic.* Typically, however, mechanisms ask the players to reveal *much more* information about themselves than that deducible from the outcomes: very often they ask the players to reveal their entire true types. This *additional* loss of privacy, as argued by [6], has the potential of distorting incentives, and thus of preventing the achievement of the designer's goals. Players who value their privacy de facto receive a *negative utility* when revealing their true types. Thus, when a mechanism asks such players to so reveal themselves, it simultaneously provides them both with positive and negative incentives, so that it may no longer be clear how these opposing forces will balance out.

1.2 Relevant Prior Mechanisms

Mechanisms that, in our perfect-knowledge contexts, achieve our goal at just one equilibrium are easy to design, but also totally vulnerable to the crucial problems of equilibrium selection and equilibrium failure.[3]

When there are at least 3 players perfectly informed about each other, the general normal-form mechanism of Jackson, Palfrey, and Srivastava [7] can achieve any social-choice property, and thus our goal as a special

[1] 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, rather than an equilibrium, a mixture of σ and τ will be played out, so that $\mathbb{P}$ may not hold. Of course, this problem worsens as the number of players and/or equilibria grows.

[2] In a second-price auction, although the mechanism is dominant-strategy, if the players with the highest two valuations for the item on sale collude, then the revenue generated drops from the second-highest to the third-highest valuation. As for a more extreme example, Ausubel and Milgrom [2] show that, in a combinatorial auction, two sufficiently informed players can totally destroy the economic efficiency of the VCG mechanism [**?**, **?**, **?**], although it too is dominant-strategy.

[3] Consider the following normal-form mechanism, which we term HOPE-FOR-THE-BEST: *Each player reports the (alleged) true-type profile. If all reports coincide with the same profile t, then (1) choose the state* $\omega = \text{argmax}_x \sum_i t_i(x)$*, and (2) choose each price* P_i *to be* $t_i(\omega)$*, minus a small discount* ϵ*. Else, choose the "null outcome"* $(\perp, 0^n)$.

Clearly, truthfully reporting the players' types is a Nash equilibrium yielding virtually maximum revenue. However, HOPE-FOR-THE-BEST also has *additional* equilibria E_x, where in E_x all players report all true valuations divided by x. Thus, the truthful equilibrium is E_1, and in each E_x the utility of each player is increased by a factor x, and the money collected is a fraction $1/x$ of the maximum possible revenue. Accordingly, HOPE-FOR-THE-BEST is extremely vulnerable to equilibrium selection: not only it has many equilibria, but the players actually prefer any one of them to the single one achieving maximum revenue.

case. But their mechanism continues to be vulnerable to equilibrium selection, although to a lesser extent.[4] In addition, their mechanism is also vulnerable to collusion and privacy.

Again when there are at least 3 players perfectly informed about each other (as long as some additional technical conditions are satisfied), the classical normal-form mechanism of Abreu and Matsushima [1] can achieve our goal, since it too "virtually" achieves essentially all social-choice properties. The main advantage of their normal-form mechanism is that it is invulnerable to equilibrium selection or failure. This is so because a single equilibrium survives the iterative elimination of strictly dominated strategies. However, their mechanism is very vulnerable to collusion, complexity, and privacy. More precisely,

1. *It is totally vulnerable to collusion.* That is, every pair of players (i, j) could jointly and profitably deviate from their equilibrium strategies. And in our application, when such a deviation occurs the revenue cannot be maximum.
2. *Its communication complexity (and thus its computational complexity) is doubly exponential in k.* Here k denotes the number of bits required to specify the value that a player has for a given state.
3. *It has maximum privacy loss.* That is, it requires the players to reveal their true types in their entirety.

A variant of this mechanism was put forward by Glazer and Perry [5]. Their mechanism is of extensive form and complete information, and virtually achieves essentially all social-choice properties at a unique subgame-perfect equilibrium. Accordingly, their mechanism too is invulnerable to equilibrium selection. However, the Glazer-Perry mechanism continues to be vulnerable to collusion, complexity and privacy problems. More precisely, its vulnerability to collusion and privacy is identical to that of the Abreu-Matsushima mechanism, while its communication complexity is exponential in k.

1.3 Our Results

Informally, like Glazer and Perry, we put forward an extensive-form mechanism $\mathcal{M}$ of complete information, that, in all contexts with quasi-linear utility and non-negative valuations in which the players are perfectly informed about each other, yields a game with a single subgame-perfect equilibrium. At this equilibrium, our mechanism generates virtually maximum revenue. Our M too, therefore, is immune to equilibrium selection problems. In addition, it satisfies the following new properties:

- $\mathcal{M}$ works for any number of players $n > 1$.
- $\mathcal{M}$ has perfect collusion resilience.
- $\mathcal{M}$ has perfect communication complexity.
- $\mathcal{M}$ has perfect privacy.

By saying that $\mathcal{M}$ has perfect collusion resiliency we mean that $\mathcal{M}$ generates virtually optimal revenue, unless *all* players belong to the *same* coalition. In our collusive model, players are free to coordinate their actions as they please (e.g., by entering binding contracts with each other and by making side-payments to each other). Therefore M's collusion resiliency is indeed best possible: when all players belong to the same coalition and coordinate each other so that the strategy profile appears to be consistent with a different true-type profile, no mechanism can guarantee anything. In our collusive model each coalition is assumed to act so as to maximize the sum of the individual utilities of its members, and the players know who colludes with whom —a natural choice, since we are dealing with perfectly informed players. (Notice that modeling

[4]The point is that the JPS mechanism has plenty of other equilibria σ that generate smaller revenue while being more attractive to *all* players. As in the case of HOPE-FOR-THE-BEST of Footnote 2, each such σ consists of reporting all true valuations divided by the same factor x. To be sure, this time each component σ_i is weakly dominated by some other strategy σ'_i. This means that, in all cases (i.e., for all possible subprofiles of strategies for the other players) σ_i provides no more utility to i than σ'_i does, while in at least some cases σ_i provides less utility to i than σ'_i. But in the JPS mechanism this happens in only one case: when *all other players "suicide"* (i.e., when all other players deliberately choose the worst possible strategies for themselves). Thus, as long as a single player does not believe that all others will commit mass suicide, all players prefer σ to the truthful and revenue-maximizing equilibrium τ.

the players as also knowing who colludes with whom would not help the above discussed prior mechanisms to be collusion-resilient.)

By saying that $\mathcal{M}$ has perfect communication complexity we mean that, at $\mathcal{M}$'s unique subgame-perfect equilibrium, the total number of bits exchanged is at most that required to specify an outcome plus 2 bits per player. (Notice that, since our states and types are totally abstract, one cannot analyze the computation requirement of our players. That of the mechanism are instead trivial no matter what.)

By saying that $\mathcal{M}$ has perfect privacy we mean that, at $\mathcal{M}$'s unique subgame-perfect equilibrium, nothing can be learned about the players' types, by the mechanism designer or any observer of the play, except for what is deducible from a perfect-revenue outcome.

2 Preliminaries

Basic Notation

- *Integers and Reals.* We denote by $\mathbb{Z}$ the set of integers, by $\mathbb{Z}^+$ the set of non-negative integers, by $\mathbb{R}$ the set of real numbers, and by $\mathbb{R}^+$ the set of non-negative reals.
- *Subprofiles.* A subprofile is a vector indexed by a subset of the players. We refer to a subprofile $V = \{V_i : i \in S\}$ as a subprofile over S. To emphasize that V is a subprofile over S we may write V_S.
- *Plays.* We may also refer to a strategy profile of a mechanism M as a *play* of M. For a play σ of M, $M(\sigma)$ denotes the corresponding outcome, or distribution over outcomes if M or the σ_i are probabilistic.

Our Contexts A *collusive context C with quasi-linear utilities and non-negative valuations* (a collusive QLU$^+$context for short) is identified by four components, $C = (\mathcal{N}, \Omega, \mathcal{T}, \theta, \mathbb{C})$, where

- $\mathcal{N}$, is a finite set of *players:* $\mathcal{N} = \{1, \ldots, n\}$.
- Ω, is a finite set of *states*;
- $\mathcal{T}$ is a profile, where each $\mathcal{T}_i$ is a set, the set of i's *types.* Each type maps Ω to $\mathbb{Z}^+$, and $\perp$ to 0.
- θ is a profile, where each θ_i, i's true type, is a member of $\mathcal{T}_i$; and
- $\mathbb{C}$ is a non-trivial partition of $\mathcal{N}$.

In C, the components $\mathcal{N}$, Ω, and $\mathcal{T}$ are common knowledge to everyone. By saying that C is *perfect-knowledge* we mean that θ and $\mathbb{C}$ are also common knowledge to the players (but not to the designer).

The set of outcomes of C is $\Omega \times \mathbb{R}^n$. The *null outcome* consists of $(\perp, 0^n)$. If $\mathscr{O} = (\omega, P)$ is an outcome, we refer to ω as *$\mathscr{O}$'s state*; to P as *$\mathscr{O}$'s price profile*; to each P_i as *the price of i;* to $\sum_i \theta_i(\omega)$ as *$\mathscr{O}$'s social welfare;* and to $\sum_i P_i$ as *$\mathscr{O}$'s revenue.* The *maximum social welfare* of C, denoted by $msw(C)$, is defined to be $\max_\omega \sum_i \theta_i(\omega)$.

For each player i, i's utility function u_i in C is so defined: for each outcome (ω, P), $u_i(\omega, P) = \theta_i(\omega) - P_i$. If D is a distribution over outcomes, i's utility for D is denoted by $U_i(D)$ and is defined to be $E_D[u_i(\omega, P)]$.

If S is a subset in $\mathbb{C}$, then S represents the maximal subset of players colluding with each other, and S's members act so as to maximize the sum of their (individual) utilities. A collusive set is a member of $\mathbb{C}$ with cardinality greater than 1. A player i is independent if $\{i\} \in \mathbb{C}$. The context is non-collusive if all players are independent. If C is a non-collusive context, we may more simply write $C = (\mathcal{N}, \Omega, \mathcal{T}, \theta)$.

Our Mechanisms We solely consider mechanisms of extensive form and complete information. Accordingly, each mechanism M specifies

- a finite game tree T;
- an outcome for each leaf (terminal node) of T; and
- for each internal node (decision node) X of T, a player P_X (the acting player) and an action set A_X consisting of the set of children of X.

In M, a (pure) strategy for a player i specifies an action in A_X for each decision node X for which $i = P_X$.

Games A *game* G consists of a context C and a mechanism M: $G = (C, M)$.

Subgame-Perfect Equilibria Since in a perfect-knowledge collusive QLU$^+$context C all players know the utilities of all other players, and all mechanisms M we consider are of extensive form and complete information, the traditional notion of subgame-perfect equilibrium applies to the games $G = (C, M)$. Recall that such an equilibrium is a strategy profile σ such that, for each decision node N in G's game tree, letting G^N be the subgame of G rooted at N, and letting σ^N be the restriction of σ to G^N, then σ^N is a Nash equilibrium of G^N.

3 Our Mechanism

In describing our $\mathcal{M}$, for clarity only, we make use of multiple "steps taken by the mechanism." These are marked by letters and interspersed with the players' steps, which are marked by numbers. (In reality, only the players act until a terminal node is reached.)

Our $\mathcal{M}$ works for any QLU$^+$context $(\mathcal{N}, \Omega, \mathcal{T}, \theta, \mathbb{C})$, if its parameters are chosen as follows.

- ϵ and each ϵ^i_j, for $i \in \{2, \ldots, n\}$ and $j \in \{1, \ldots, n\}$, are constants so that
$\frac{1}{5n} > \epsilon > \epsilon^2_1 > \cdots > \epsilon^2_n > \epsilon^3_1 > \ldots > \epsilon^3_n > \cdots > \epsilon^n_1 > \cdots > \epsilon^n_n > 0.$
- n_r is the cardinality of $\mathcal{O}_r$, the set of outcomes with price profiles in $(\mathbb{Z}^+)^n$ and revenue equal to r. For any $(\omega, P) \in \mathcal{O}_r$, $f_r(\omega, P)$ is the rank of (ω, P) in $\mathcal{O}_r$ —whose outcomes are ordered considering first their states, in lexicographical order, and then their price profiles, in the order in which P precedes P' if (1) $P_j = P'_j$ for all $j \le i$ and (2) $P_{i+1} > P'_{i+1}$. (Therefore, $0 \le f_r(\omega, P) < n_r$.)

Mechanism $\mathcal{M}$

a: *Set $\omega = \perp$, and $P_i = 0$ $\forall i$.*

> Comment. (ω, P) will be the final outcome of $\mathcal{M}$.

1: *Player 1 announces a state $\omega^\star$ and a profile K of non-negative integers.*

> Comment. $(\omega^\star, K)$ is player 1's proposed outcome, allegedly an outcome of maximum revenue. In the unique rational play of $\mathcal{M}$, $\omega^\star$ is guaranteed to be the final state ω if $\sum_j K_j > 0$.

b: *If $\sum_i K_i = 0$, the mechanism ends right now.*

$[2, n]$: *In Step $i \in [2, n]$, player i announces a profile Δ^i of non-negative integers such that $\Delta^i_i = 0$.*

> Comment. i suggests to raise the current price of j, that is $K_j + \sum_{\ell=2}^{i-1} \Delta^\ell_j$, by the amount Δ^i_j.

c: *For each player i, publicly select bip_i and $P^\star_i$ as follows. Let $R_i = \{j : \Delta^j_i > 0\}$. If $R_i \neq \emptyset$, then bip_i is the lexicographically highest player in R_i, and $P^\star_i = K_i + \sum_{\ell=2}^{bip_i} \Delta^\ell_i$. Else, $bip_i = 1$ and $P^\star_i = K_i$.*

> Comment. bip_i is the "best informed player about i", and $P^\star_i$ the "provisional price of i."

$[n+1, 2n]$: *In Step $n + i$, if $P^\star_i > 0$, player i announces YES or NO.*
(By default, i announces YES if $P^\star_i = 0$, and player 1 announces YES if $bip_1 = 1$.)

> Comment. Each player i announces YES or NO assuming that (1) the final state is $\omega^\star$ and (2) his final price is $P^\star_i - \epsilon - \delta_i$, where
>
> - ϵ is our predetermined constant, and thus a fixed, small but positive discount due to any player if this stage of the mechanism is reached, and

- δ_i is a variable, small and non-negative "knowledge discount" offered to player i. This discount depends on the amount of knowledge contributed by i. The knowledge contributed by player 1 consists firstly of the revenue and secondly of the rank of his proposed outcome. The knowledge contributed by any other player consists of all price raises he proposes.

By default, player 1 accepts his own price if no one raises it.

At this point the players are done playing, and the mechanism proceeds as follows. If all players announce YES, the so ϵ, δ-discounted outcome $(\omega^\star, P^\star)$ is implemented with probability 1. Else:

- With high probability the final state is $\perp$, the best-informed players of the players announcing NO are punished with suitably high fines, and all other players pay nothing.
- With very small probability, but proportional to the number of players saying YES, the ϵ, δ-discounted outcome $(\omega^\star, P^\star)$ is implemented as if all players said YES.
- With complementary probability the (only δ-discounted) null outcome is chosen.

Therefore, any independent player i who has a non-negative utility for the outcome $(\omega^\star, P^\star)$ strictly prefers announcing YES, regardless of (1) what the other players announce and (2) whether he is in line for punishment —for having raised the price of some players too much and become their best informed player. In fact, if the ϵ, δ-discounted outcome $(\omega^\star, P^\star)$ were to be implemented in the end, his utility would be $\geq \epsilon + \delta_i$, while his utility for the δ-discounted null outcome $(\perp, 0^n)$ would be exactly δ_i.

d: *If all players announce YES, then reset ω to $\omega^\star$, and each P_i to $P_i^\star - \epsilon$. Then, go to Step* g.

e: *Publicly flip a biased coin c_1 such that* $\Pr[c_1 = \textit{Heads}] = 1 - \epsilon$.

If $c_1 = \textit{Heads}$, reset P_{bip_i} to $P_{bip_i} + 2P_i^\star$ for each player i announcing NO and go to Step g.

f: *Publicly flip a biased coin c_2 such that* $\Pr[c_2 = \textit{Heads}] = \frac{Y}{nB}$, *where Y is the number of players announcing YES, and* $B = \sum_{i \text{ announces NO}} P_i^\star$.

If $c_2 = \textit{Heads}$, reset ω to $\omega^\star$ and each P_i to $P_i^\star - \epsilon$.

Comment. If $c_2 =$Tails, ω and P continue to be $\perp$ and $(0, \ldots, 0)$.

g: *Reset each P_i to $P_i - \delta_i$, where*

$\delta_1 = 2\epsilon \sum_j K_j - \epsilon \frac{f_r(\omega^\star, K)}{n_r}$, *where* $r = \sum_j K_j$, *and, for* $i \neq 1$

$\delta_i = \sum_j \epsilon_j^i \Delta_j^i$.

Comment. Although some players' prices may be negative, the mechanism never loses money, and in the unique rational play the utility of every player is non-negative. Note that, by ending in Step b when $\sum_j P_j = 0$, $\mathcal{M}$ avoids losing money (due to its "reward program") if no player raises prices.

4 Traditional Analysis of Our Mechanism

Let us first analyze our mechanism disregarding collusion, complexity and privacy.

Notation If N is a decision node of an extensive-form game G, and σ a strategy profile of G. Then G^N denotes the subgame rooted at N; σ^N the restriction of σ to the subgame G^N; and $\sigma_i(N)$, if i is the player acting at N, the action that σ_i specifies for i at N.

We define the depth of a node N in $\mathcal{M}$'s game tree to be the number of nodes in the path from the root to N. That is, N's depth is d if there is a sequence of nodes $x_1, \ldots, x_d$ such that x_1 is the root, each x_{i+1} is a child of x_i, and x_d is node N. Accordingly, the depth of the root is 1.

Note that the maximum depth of a decision node in $\mathcal{M}$'s game tree is $2n$. In fact, unless player 1 ends the mechanism right away by proposing an outcome of 0 revenue, each player acts exactly twice. More precisely, at any node of depth $d \in [1, n]$ the acting player is d, and at any node of depth $d \in [n+1, 2n]$ is $d - n$.

Definition 1. *For any game $(C, \mathcal{M})$, where $C = (\mathcal{N}, \Omega, \mathcal{T}, \theta)$ is a perfect-knowledge* non-collusive *QLU^+ context, we denote by $\overline{\theta}$ the strategy profile defined as follows: for each decision node N,*

- *If N has depth $d \in [n+1, 2n]$, then $\overline{\theta}_{d-n}(N)$ consists of announcing YES if $\theta_{d-n}(\omega^\star) \geq P^\star_{d-n}$, and NO otherwise.*
- *If N has depth $d \in [2, n]$, then $\overline{\theta}_d(N)$ consists of announcing Δ^d such that, for $j \neq d$, $\Delta^d_j = \max\left(0, \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{d-1} \Delta^\ell_j\right)$.*
- *If N has depth 1, then $\overline{\theta}_1(N)$ consists of announcing $\omega^\star = \text{argmax}_\omega \sum_i \theta_i(\omega)$ and $K_i = \theta_i(\omega^\star)$ for all i.*

(We purposely chose the symbol "$\overline{\theta}$" to be graphically similar to the symbol "θ". In fact, the latter represents the true-type profile, and former the "truthful strategy profile", where the players truthfully use their knowledge of θ so as to produce an outcome of virtually maximum revenue.)

Since in $\mathcal{M}$ the parameter ϵ can be chosen to be arbitrarily small, to prove that $\mathcal{M}$ generates virtually maximum revenue at a unique subgame-perfect equilibrium (without imposing arbitrarily high prices), it suffices to prove the following.

Theorem 1. *In any game $(C, \mathcal{M})$, where $C = (\mathcal{N}, \Omega, \mathcal{T}, \theta)$ is a perfect-knowledge non-collusive QLU^+ context, the following three properties hold:*

$\mathbb{P}1$. $\overline{\theta}$ is the unique subgame perfect equilibrium.

$\mathbb{P}2$. if $\mathcal{M}(\overline{\theta}) = (\omega, P)$, then ($\mathbb{P}2.1$) $\sum_i \theta_i(\omega) = msw(C)$, and ($\mathbb{P}2.2$) $\sum_i P_i \geq (1 - 3\epsilon n) \cdot msw(C)$.

$\mathbb{P}3$. $U_i(\mathcal{M}(\overline{\theta})) \geq 0$ for each player i.

Proof. Let C be as per our hypothesis, and let $G = (C, \mathcal{M})$. Then, as proven in the next three subsections, the following lemmas hold.

Lemma 1. If N is a node of depth $d \in [n+1, 2n]$, then $\overline{\theta}^N$ is the unique subgame-perfect equilibrium of G^N.

Lemma 2. If N is a node of depth $d \in [2, n]$, then $\overline{\theta}^N$ is the unique subgame-perfect equilibrium of G^N.

Lemma 3. If N is the root, then $\overline{\theta}^N$ is the unique subgame-perfect equilibrium of G.

Notice that Property $\mathbb{P}1$ coincides with the statement of Lemma 3. Here, therefore, we only prove Properties $\mathbb{P}2$ and $\mathbb{P}3$. Set $MSW = msw(C)$. By definition, when player 1 follows strategy $\overline{\theta}_1$ he announces $\omega^\star$ and K such that $\sum_i \overline{\theta}_i(\omega^\star) = MSW$ and $K_i = \overline{\theta}_i(\omega^\star)$ for all i.

As for Property $\mathbb{P}2$, we distinguish two cases: $MSW = 0$ and $MSW > 0$. In the first case, $\sum_i K_i = 0$, and thus $\mathcal{M}$ ends in Step b, with state $\omega = \perp$ and each price $P_i = 0$. Since $\theta_i(\perp) = 0$ for each i, we have that $\sum_i \theta_i(\perp) = 0 = \sum_i P_i$, that is, the revenue equals the maximum social welfare. Therefore Property $\mathbb{P}2$ holds.

In the second case, $\sum_i K_i > 0$ so that $\mathcal{M}$ ends at Step g. Moreover, under strategy profile $\overline{\theta}$ no player i raises any price (i.e., $\Delta^i_j = 0$ for all j) in Steps 2 through n, and all players say YES in Steps $n+1$ through $2n$, so that $\mathcal{M}$ sets the final state to $\omega^\star$ and each price P_i to $K_i - \epsilon - \delta_i$. In the case under consideration, the knowledge reward δ_i for each player $i \neq 1$ is 0, and for player 1 is $\delta_1 = 2\epsilon MSW - \epsilon\frac{f_r(\omega^\star, K)}{n_r}$. Now, since $\delta_1 < 2\epsilon MSW$, the final revenue of the mechanism (including all rewards) is

$$\sum_i P_i = MSW - n\epsilon - \delta_1 > MSW - n\epsilon - 2\epsilon MSW > MSW - n\epsilon MSW - 2\epsilon MSW = (1 - 3\epsilon n)MSW,$$

where the second inequality holds because, types being integral functions, MSW is a positive integer. Therefore Property $\mathbb{P}2$ holds.

As for property $\mathbb{P}3$, if each θ_i is identically 0 over Ω, then $U_i(\mathcal{M}(\overline{\theta})) = U_i(\perp, 0^n) = 0$ for every i. Otherwise, $U_i(\mathcal{M}(\overline{\theta})) = \epsilon > 0$ for every $i \neq 1$ and $U_1(\mathcal{M}(\overline{\theta})) = \epsilon + \delta_1 > 0$. Thus Property $\mathbb{P}3$ holds too. ■

Remark By giving each player i an extra discount ϵ in Step g, $\mathcal{M}$ could guarantee an additional property. Namely, it could guarantee that the utility of each player i choosing $\overline{\theta}_i$ as his strategy to be non-negative whenever all players announces YES or NO rationally, even if some other player j acted irrationally and raised i's price beyond $\theta_i(\omega^\star)$.

This additional property is important when dealing with players of limited rationality. Indeed, for each player, given a state and a price for him to pay, saying YES or NO correctly requires minimum rationality; however, given a reward proportional to the amount of price-raise suggested by him, suggesting a price-raise correctly depends on a correct reasoning about the other players' strategies, and thus requires "more rationality."

4.1 Proof of Lemma 1

Lemma 1. *If N is a node of depth $d \in [n+1, 2n]$, then $\overline{\theta}^N$ is the unique subgame-perfect equilibrium of G^N.*

Proof. To prove our lemma it suffices to show that, in any subgame-perfect equilibrium σ, i announces YES in Step $n+i$ if $\theta_i(\omega^\star) \geq P_i^\star$, and NO otherwise. The proof is by induction on the depth d of N.

Base Case. In this case the depth is $d = 2n$ and the acting player is n. We distinguish between two cases.

Case 1: $P_n^\star \leq \theta_n(\omega^\star)$. In this case, regardless of the number of the other players who announced YES before him, if player n announces YES he increases the chances that the final outcome is the properly ϵ, δ-discounted version of $(\omega^\star, P^\star)$. Now, because in such a final outcome he pays at most $P_n^\star - \epsilon - \delta_n$, and because his true valuation for $\omega^\star$ is at lest $P_n^\star$, player n has utility at least $\epsilon + \delta_n$, and strictly prefers this outcome to the only δ-discounted null outcome $(\perp, 0^n)$, where his utility is exactly δ_n. Therefore, it is strictly dominant for n to announce YES in Step $2n$, as we wanted to show.

Case 2: $P_n^\star > \theta_n(\omega^\star)$. In this case, again regardless of the number of players who announced YES, by announcing YES himself player n strictly increases the probability that the final outcome is the properly ϵ, δ-discounted version of $(\omega^\star, P^\star)$. Now, because both $P_n^\star$ and $\theta_n(\omega^\star)$ are integers, in our case we have $P_n^\star \geq \theta_n(\omega^\star) + 1$. Accordingly, in the envisaged final outcome, because his discount is $\epsilon + \delta_n$, player n has utility at most $\epsilon + \delta_n - 1$, which is less than δ_n. Therefore he strictly prefers the δ-discounted null outcome $(\perp, 0^n)$, where his utility is exactly δ_n, and it is strictly dominant for him to announce NO in Step $2n$, as we wanted to show.

Induction Step Let N be a node of depth d, with $n < d < 2n$. Then the acting player is $i = d - n$, and, by inductive hypothesis, for every action taken by player i in N, denoting by M the corresponding child of N, we have a subgame perfect equilibrium in the subgame G^M. In such an equilibrium each player $j > i$ announces YES if and only if $P_j^\star \leq \theta_j(\omega^\star)$. Therefore, given that all players after i play according to this equilibrium, we can consider two separate cases:

Case 1: $P_i^\star \leq \theta_i(\omega^\star)$. In this case, regardless of the number of the other players who announced YES before him, or will announce YES after him, if player i announces YES he increases the chances that the final outcome is the ϵ, δ-discounted version of $(\omega^\star, P^\star)$. This is so because, by inductive hypothesis, the players acting after him announce YES or NO according to their own valuations, independently of what i announces. Since i's utility in the discounted version of $(\omega^\star, P^\star)$ is $\epsilon + \delta_i$, while his utility in the δ-discounted null outcome $(\perp, 0^n)$ is exactly δ_i, i's single best strategy is to announce YES, as we wanted to show.

Case 2: $P_i^\star > \theta_i(\omega^\star)$. The proof of this case can be easily derived from that of the previous one, taking in consideration that i strictly prefers the δ-discounted null outcome $(\perp, 0^n)$ where his utility is δ_i, to the ϵ, δ-discounted version of $(\omega^\star, P^\star)$ where his utility is at most $\epsilon + \delta_i - 1 < \delta_i$. ■

4.2 Proof of Lemma 2

Lemma 2. *If N is a node of depth $d \in [2, n]$, then $\overline{\theta}^N$ is the unique subgame-perfect equilibrium of G^N.*

Proof. To prove our thesis it suffices to prove that, since d is the acting player at node N, for any subgame-perfect equilibrium σ and every player $j \neq d$ the following two properties hold for σ_d:

1. *No Overbidding:* $\Delta_j^d \le \max\left(0, \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{d-1} \Delta_j^\ell\right)$;
2. *No Underbidding:* if $K_j + \sum_{\ell=2}^{d-1} \Delta_j^\ell < \theta_j(\omega^\star)$, then $\Delta_j^d = \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{d-1} \Delta_j^\ell$.

We prove both properties by induction on d.

Base Step: $d = n$**.**

Proof of No Overbidding. To prove that strategy σ_n satisfies the no-overbidding condition in the base case, we must consider two separate possibilities: namely, $\max\left(0, \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{d-1} \Delta_j^\ell\right) = 0$ and $\max\left(0, \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{d-1} \Delta_j^\ell\right) = \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{d-1} \Delta_j^\ell$. That is, we must prove that σ_n satisfies the following two conditions:

B1. If $K_j + \sum_{\ell=2}^{n-1} \Delta_j^\ell \ge \theta_j(\omega^\star)$ then $\Delta_j^n = 0$;

B2. If $K_j + \sum_{\ell=2}^{n-1} \Delta_j^\ell < \theta_j(\omega^\star)$, then $\Delta_j^n \le \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{n-1} \Delta_j^\ell$.

We prove that B1 holds by contradiction. Suppose σ_n causes n to *overbid* on another player j, that is, to announce $\Delta_j^n > 0$. Therefore, we have $K_j + \sum_{\ell=2}^{n} \Delta_j^\ell > \theta_j(\omega^\star)$. Consider now an alternative strategy $\widehat{\sigma}_n$ for player n, in which he makes the same announcements as σ_n at each of the nodes of G^N in which n acts, except that, at node N, he announces

$$\widehat{\Delta}_j^n = 0.$$

To establish B1 is thus sufficient to prove that $\widehat{\sigma}_n$ gives player n a strictly larger expected utility than σ_n (when all other players stick to their strategies in σ). The difference between i's utility under these two strategies comes from the sum of three terms:

1. *(Punishment)* Under strategy σ_n, $bip_j = n$ because n is the last player who raises player j's price. In light of Lemma 1 and the fact that σ is assumed to be a subgame-perfect equilibrium, player j will announce NO to the price proposed by n. Thus, with probability $1-\epsilon$, player n will be punished with a fine of $2P_j^\star$. By contrast, player n will not be punished under strategy $\widehat{\sigma}_n$. Accordingly, relative to this term alone, the difference between n's (expected) utility under $\widehat{\sigma}_n$ and his utility under σ_n is $(1-\epsilon)2P_j^\star$, which is always $\ge (1-\epsilon)2\Delta_j^n$.

2. *(Reward)* Strategy $\widehat{\sigma}_n$ gives player n a smaller reward than σ_n. However, the difference between n's utilities under $\widehat{\sigma}_n$ and σ_n due to this term is lowerbounded by $-\epsilon_j^n\Delta_j^n$, and thus by $-\epsilon\Delta_j^n$.

3. *(Probability-Change)* By using $\widehat{\sigma}_n$ instead of σ_n, the probability that $c_2 =$ Heads changes. This change can either increase or decrease player n's expected utility, depending on whether he prefers state $\perp$ with price $-\delta_n$, or state $\omega^\star$ with price $P_n^\star - \epsilon - \delta_n$. However, the difference between n's expected utilities under $\widehat{\sigma}_n$ and σ_n due to this term is in any case lowerbounded by $-\Delta_j^n$. This can be seen by analyzing separately the case in which all players but j announce YES, and the case in which at least one player besides j announces NO. This last case requires $\mathcal{M}$ to use two coin flips.

Thus the difference between n's overall expected utility under $\hat{\sigma}_n$ and n's overall expected utility under σ_n is greater than $(1-\epsilon)2\Delta_j^n - \epsilon\Delta_j^n - \Delta_j^n = (1-3\epsilon)\Delta_j^n$ and thus greater than 0 as desired. This proves B1.

The proof of B2 is similar to that of B1 and is therefore omitted.

Proof of No Underbidding. Suppose that strategy σ_n causes player n to underbid on another player j, that is, to announce Δ_j^n such that $K_j + \sum_{\ell=2}^{n-1} \Delta_j^\ell + \Delta_j^n < \theta_j(\omega^\star)$. Consider now an alternative strategy $\widehat{\sigma}_n$

for player n, in which he makes the same announcements as σ_n at each of the nodes of G^N in which n acts, except that, at node N, he announces

$$\widehat{\Delta}_j^n = \theta_j(\omega^\star) - (K_j + \sum_{\ell=2}^{n-1} \Delta_j^\ell).$$

It is easy to see that $\widehat{\sigma}_n$ gives player n a strictly larger expected utility than σ_n. Indeed, the only difference between n's utilities under these two strategies comes from the δ-discount, which is proportional to $\widehat{\Delta}_j^n$ and thus larger under strategy $\widehat{\sigma}_n$. Thus underbidding can not be a subgame-perfect strategy for player n.

This concludes the proof for our base step. □

Together with Lemma 1, the base case implies that there is a unique subgame perfect equilibrium at each node of depth n, in which player

(1) n raises each player j's price exactly up to $\theta_j(\omega^\star)$, and

(2) in each depth-$(n+j)$ node, player j announces YES if and only if $\theta_j(\omega^\star)$ is greater than or equal to $P_j^\star$.

Inductive Step $d \in [2, n-1]$.

Proof of No Overbidding. To prove that σ_d satisfies the no-overbidding property it suffices to show that σ_d satisfies the following two conditions. (Recall: d is the player active at a node of such depth.)

I1. If $K_j + \sum_{\ell=2}^{d-1} \Delta_j^\ell \geq \theta_j(\omega^\star)$ then $\Delta_j^d = 0$;

I2. If $K_j + \sum_{\ell=2}^{d-1} \Delta_j^\ell < \theta_j(\omega^\star)$, then $\Delta_j^d \leq \theta_j(\omega^\star) - K_j - \sum_{\ell=2}^{d-1} \Delta_j^\ell$.

As for the matching condition of our Base Step, we prove that I1 holds by contradiction. Suppose σ_d causes d to *overbid* on another player j, that is, to announce $\Delta_j^d > 0$. Then, we have $K_j + \sum_{\ell=2}^{d-1} \Delta_j^\ell + \Delta_j^d > \theta_j(\omega^\star)$. Consider now an alternative strategy $\widehat{\sigma}_d$ for player d, in which he makes the same announcements as σ_d everywhere except that at node N he announces

$$\widehat{\Delta}_j^d = 0.$$

To establish I1 it thus suffices to prove that $\widehat{\sigma}_d$ gives player d a strictly larger expected utility than σ_d. The difference between d's utilities under these two strategies comes from the sum of the same three terms considered in the Base Step: Punishment, Reward, and Probability Change. The analyses of the Reward and the Probability-Change terms are totally similar to those of their counterparts in the Base Step, and are therefore omitted. As for Punishment, in the Base Step the active player was n and thus, automatically, he was the last player capable of raising prices. Thus, $\Delta_j^n > 0$ straightforwardly implied that $bip_j = n$. In our Inductive Step we have to rely on our inductive hypothesis to prove that $bip_j = d$, but this is not hard to do. Indeed, according to the inductive hypothesis, in the subgame arising after d acts, there is a unique subgame-perfect equilibrium, and in this equilibrium no player $k > i$ will raise j's price. Thus $bip_j = d$. The rest of the analysis of Punishment is totally similar to that of the Base Step. This establishes the proof of I1.

The proof of I2 does not present any additional difficulties and is omitted.

No Underbidding. The proof of the no-underbidding property is totally similar to that of the base case, and is therefore omitted.

This concludes the proof of our Inductive Step, and thus of Lemma 2. ■

4.3 Proof of Lemma 3

Lemma 3. *If N is the root, then $\overline{\theta}^N$ is the unique subgame-perfect equilibrium of G.*

Proof. In light of Lemmas 1 and 2, it suffices to prove that, whenever σ is a subgame-perfect equilibrium of G, according to $\sigma_1(N)$ player 1 announces an outcome $(\omega^\star, K)$ of maximum revenue and first rank. We divide the proof into two parts.

Part 1. In this part we prove that, no matter what state $\omega^\star$ player 1 may announce, σ_1^N causes him to announce K so that $K_i = \theta_i(\omega^\star)$ for each player i. We proceed by contradiction.

Assume first that σ_1 is such that $K_i > \theta_i(\omega^\star)$ for some player i. Define $\widehat{\sigma}_1$ to be the strategy under which player 1 makes the same announcements as under σ_1, except that, at the root and relative to player i, player 1 announces a proposed price $\widehat{K}_i = \theta_i(\omega^\star)$. Consider the difference in the expected utilities of player 1 under $\widehat{\sigma}_1$ and under σ_1 (assuming that all other players stick to their strategies in σ). This difference is again due to three terms: Reward, Punishment and Probability Change. And as in prior analysis the key to see that this difference is positive is realizing that, because as per Lemmas 2 and 3 we know what the other players will do in all subsequent steps, (1) under $\widehat{\sigma}_1$ player 1 is not punished but gets a smaller reward, while (2) under σ_1 he is punished and gets a bigger reward, and (3) the punishment is much bigger than the reward and any possible gain or loss from Probability Change. This of course leads to a contradiction.

A similar analysis shows that a contradiction is also reached assuming that $K_i < \theta_i(\omega^\star)$ for some player i. Thus it must be $K_i = \theta_i(\omega^\star)$ for some player i, as we wanted to show for Part 1. □

Proof of Part 2. In this part we prove that in any subgame-perfect equilibrium σ, Player 1 must choose the state $\omega^\star$ so as to maximize the social welfare, with ties broken lexicographically. Again, the proof is by contradiction. Leaving the breaking of ties aside for a moment, assume that σ_1 causes player 1 to announce a state $\omega^\star$ such that $\sum_j \theta_j(\omega^\star) < msw(C)$. Let $\widehat{\sigma}_1$ be the strategy under which player 1 makes the same announcements as under σ_1, except that, at the root, he announces $\widehat{\omega}^\star$ and $\widehat{K}$, where $\sum_j \theta_j(\widehat{\omega}^\star) = msw(C)$ and $\widehat{K}_j = \theta_i(\widehat{\omega}^\star)$ for each player j. Let all other players stick to their strategies in σ. Then, Lemmas 1 and 2 imply that player 1's expected utility under $\widehat{\sigma}_1$ is

$$\epsilon + 2\epsilon \sum_i \theta_i(\widehat{\omega}^\star) - \epsilon \frac{f_{\widehat{r}}(\widehat{\omega}^\star, \widehat{K})}{n_{\widehat{r}}},$$

where $\widehat{r} = \sum_i \widehat{K}_i$; while his expected utility under σ_1 is

$$\epsilon + 2\epsilon \sum_i \theta_i(\omega^\star) - \epsilon \frac{f_r(\omega^\star, K)}{n_r}.$$

Thus the difference between the two is

$$2\epsilon \left(\sum_i \theta_i(\widehat{\omega}^\star) - \sum_i \theta_i(\omega^\star) \right) - \epsilon \left(\frac{f_{\widehat{r}}(\widehat{\omega}^\star, \widehat{K}^1)}{n_{\widehat{r}}} - \frac{f_r(\omega^\star, K)}{n_r} \right)$$

which is greater than 0, because the first parenthetical quantity is a positive integer, while the second one is a number in the interval $[-1, 1]$.

Essentially the same argument shows that in the case of two or more states with maximum social welfare, player 1 chooses the lexicographically first of them. ■

5 Collusion Analysis of Our Mechanism

Let us now show that our mechanism continues to produce virtually maximum revenue in every collusive context $(\mathcal{N}, \Omega, \mathcal{T}, \theta, \mathbb{C})$. Recall that in such a context $\mathbb{C}$ is a non-trivial partition of $\mathcal{N}$ (i.e., the set of all players does not constitute a single collusive set) and is common knowledge among the players (since in our setting the players have perfectly informed about each other). Recall too that in our context all coalitions $\mathscr{C}$ of players are *rational,* that is, act so as to maximize the sum of the utilities of each members. Indeed, since our players (unknown to the mechanism designer!) are free to coordinate their actions in any way they please, including entering binding contracts with one another, limiting the "amount of money flowing into a coalition $\mathscr{C}$" would be irrational: the members of $\mathscr{C}$ could always and easily "distribute more money among themselves later on in any proportions they want."

In general (i.e., in an extensive-form game of imperfect and incomplete information), to act rationally the members of a coalition $\mathscr{C}$ should adopt a *joint* strategy, and communicate to each other every bit of information they have about the game. That is, there should be a single strategy that, at any node in which a member i of $\mathscr{C}$ acts, chooses i's action based on the information collectively available to $\mathscr{C}$. But as we are in a setting not only of *perfect* information but also of *complete* information (at least for what the players are concerned), a joint strategy of $\mathscr{C}$ simply consists of a subprofile of individual strategies.

Small Difficulties Transforming our non-collusive analysis of $\mathcal{M}$ to a collusive one does not present major difficulties, but requires some additional attention (to notation, of course, and) to "collusive utilities." Let us explain. The utility of a coalition $\mathscr{C}$ is the same as long as the state is the same and the sum of the prices paid by its members, p, is the same. It does not matter to $\mathscr{C}$ whether if the total amount p is paid by just one of its members, or by two or more of its members, in equal or different proportions. Thus, in Step $n+i$ of $\mathcal{M}$, each member i of $\mathscr{C}$ should say YES if and only if $\sum_{j\in\mathscr{C}}\theta_j(\omega^\star) \geq \sum_{j\in\mathscr{C}} P^\star_j$. Accordingly, in Step $k \in [2,n]$, a player k should be mindful, before raising the price of a member j of $\mathscr{C}$ of the total price currently proposed for $\mathscr{C}$. For instance, even if $\theta_j(\omega^\star) > P^\star_j$, it would be irrational for k to raise j's price by $P^\star_j - \theta_j(\omega^\star)$ if the currently proposed, total price for $\mathscr{C}$ is already $\sum_{\ell\in\mathscr{C}}\theta_\ell(\omega^\star)$. In fact, should k so raise the price of j, the utility of $\mathscr{C}$ under the properly ϵ,δ-discounted version of $(\omega^\star, P^\star)$ (assuming that no one else raises the price of any member of $\mathscr{C}$) would be at most

$$\sum_{i\in\mathscr{C}}(\epsilon+\delta_i) - 1 = \sum_{i\in\mathscr{C}}\delta_i + |\mathscr{C}|\epsilon - 1,$$

which is less than $\sum_{i\in\mathscr{C}}\delta_i$, which is the utility of $\mathscr{C}$ under the δ-discounted null outcome $(\perp, 0^n)$. Therefore, if k raises the price of j as above, it will be in $\mathscr{C}$'s best interest to minimize the chance that the so updated outcome $(\omega^\star, P^\star)$ will be implemented by having *all* of its members say NO in the last n steps of $\mathcal{M}$, so that k will end up being punished with very high probability.

Paying attention to collective utilities is however still compatible with several of strategic choices. In order to guarantee a unique subgame-perfect equilibrium, $\mathcal{M}$ needs to introduce incentives so as to artificially generate a strict order of preference among otherwise equivalent strategies. It is here that $\mathcal{M}$'s multiplicity of small constants become handy. As we shall see, these constants incentivize all players to raise or set the price of a coalition $\mathscr{C}$ by raising or setting the price of $\mathscr{C}$'s lexicographically first member.

In light of this discussion, in a collusive setting, the "truthful strategies" of $\mathcal{M}$ are defined as follows.

Definition 2. *For any game $(C,\mathcal{M})$, where $C = (\mathcal{N},\Omega,\mathcal{T},\theta,\mathbb{C})$ is a perfect-knowledge* collusive *QLU^+ context, we denote by $\overline{\Theta}$ the strategy profile defined as follows: for each decision node N*

- *If N has depth $d\in[n+1,2n]$, letting $i = d-n$ and $\mathscr{C}_i$ be the coalition i belongs to, then $\overline{\Theta}_i(N)$ consists of announcing YES if*
$$\sum_{j\in\mathscr{C}_i}\theta_j(\omega^\star) \geq \sum_{j\in\mathscr{C}_i} P^\star_j,$$
and NO otherwise.
- *If N has depth $d\in[2,n]$, then $\overline{\Theta}_d(N)$ consists of announcing Δ^d such that, for any coalition $\mathscr{C}$, letting k be the lexicographically first player in $\mathscr{C}$, we have:*
$$\Delta^d_k = \max\left(0, \sum_{j\in\mathscr{C}}\theta_j(\omega^\star) - \sum_{j\in\mathscr{C}} K_j - \sum_{j\in\mathscr{C}}\sum_{\ell=2}^{d-1}\Delta^\ell_j\right)$$
and $\Delta^d_j = 0$ if j is not the lexicographically first player in a coalition.

- *If N has depth 1, then $\overline{\Theta}_1(N)$ consists of announcing $\omega^\star = \text{argmax}_\omega \sum_i \theta_i(\omega)$ (ties broken lexicographically) and for every coalition $\mathscr{C}$, letting k be the lexicographically first player in the coalition*

$$K_k = \sum_{i \in \mathscr{C}} \theta_i(\omega^\star)$$

and $K_i = 0$ otherwise.

Accordingly, in a collusive setting, our main theorem can be stated as follows.

Theorem 2. *Any game $(C, \mathcal{M})$, where $C = (\mathcal{N}, \Omega, \mathcal{T}, \theta, \mathbb{C})$ is a perfect-knowledge collusive QLU^+ context, has a unique subgame-perfect equilibrium $\mathscr{E}$, and, letting $\mathcal{M}(\mathscr{E}) = (\omega, P)$, we have that*

(1) $\sum_i \theta_i(\omega) = msw(C)$, (2) $\sum_i P_i \geq (1 - 3\epsilon n) \cdot msw(C)$, and (3) $U_{\mathscr{C}}(\mathcal{M}(\mathscr{E})) \geq 0$ for each coalition $\mathscr{C} \in \mathbb{C}$.

The proof of Theorem 2 actually parallels that of Theorem 1, after restating Lemmas 1, 2, and 3 as follows for the collusive setting.

Lemma 1′. *If a node N has depth $d \in [n+1, 2n]$, then $\overline{\Theta}^N$ is the unique subgame-perfect equilibrium of G^N.*

Lemma 2′. *If a node N has depth $d \in [2, n]$, then $\overline{\Theta}^N$ is the unique subgame-perfect equilibrium of G^N.*

Lemma 3′. *If N is the node of depth 1, then $\overline{\Theta}^N$ is the unique subgame-perfect equilibrium of G.*

The proof of Lemma 1′ is trivially derived from that of Lemma 1 after noticing that it now is *strictly* dominant for a coalition to be unanimous, that is, that every player in a coalition $\mathscr{C}$ announces YES if and only if the sum of utilities of $\mathscr{C}$'s members is greater than the sum of their payments.

The proof of Lemma 2′ is trivially derived from that of Lemma 2. As mentioned, one simple difference concerns payments inside a coalition $\mathscr{C}$. Although $\mathscr{C}$ cares only about the sum of the utilities and the sum of the payments of its members, to get a unique subgame-perfect equilibrium, $\mathcal{M}$ properly structures its small rewards by choosing $\epsilon^i_j > \epsilon^i_k$ if and only if $j < k$. A second small difference concerns the case where one player i in a coalition $\mathscr{C}$ can either raise the price of another player j, or let another member of $\mathscr{C}$ raise j's price instead. To guarantee that, whenever a player i has a chance to raise the price of another player, i does so immediately (so as to guarantee the uniqueness of subgame-perfect equilibrium), mechanism $\mathcal{M}$ again properly structure its rewards, this time choosing $\epsilon^i_j > \epsilon^k_j$ if $i < k$.

The proof of Lemma 3′ is very similar to that of Lemma 3. The main difference is that, to guarantee unique subgame-perfect equilibrium, $\mathcal{M}$ uses the reward system to incentivize player 1 to announce the lexicographically first outcome maximizing the revenue.

6 Complexity of Our Mechanism

In mechanism $\mathcal{M}$, each player talks exactly twice. Moreover, the number of bits required to be communicated at the unique subgame-perfect equilibrium is optimal within a constant factor. In fact, no matter what the mechanism may be, the number of bits exchanged must be at least x, where x is the minimum number of bits sufficient to describe an outcome. And at the unique subgame-perfect equilibrium of $\mathcal{M}$, (1) the first player sends the lexicographically first outcome with maximum revenue, (2) no other player raises any prices, and (3) every player announces YES. Now notice that item 1 requires x bits by definition, while items 2 and 3 can be accomplished by sending $2n - 1$ bits altogether, if one reasonably uses the bit 0 to encode "I do not want to raise any prices" in Steps 1 through n, and "YES" in Steps $n+1$ through $2n$. (Such encoding makes more expensive to encode price raises, but at the unique equilibrium such a need does not arise.)

7 Privacy Analysis of Our Mechanism

Consider a deterministic mechanism M designed to obtain a desired outcome $\mathcal{O}$ —in our case, the lexicographically last outcome with maximum revenue— according to a given solution concept. Informally, we say that M is perfectly private if there exists a predetermined algorithm A that, on input just (a description of) the desired outcome $\mathcal{O}$, without any other information about the true-type profile of the players, computes the information divulged in a play of M that, according to the chosen solution concept, yields $\mathcal{O}$.

This is trivially the case of our $\mathcal{M}$. Indeed, given a string s describing the desired lexicographically first outcome with maximum revenue, the "transcript" of $\mathcal{M}$ at the unique subgame-perfect equilibrium consists of $s0^{2n-1}$, that is the encoding of player 1 announcing s and everyone else announcing no raises and YES.

Let us now formalize perfect privacy for the specific case at hand. (The proof that $\mathcal{M}$ is perfectly private is trivial any way.)

Notation For any profile of strategies σ for a mechanism M of extensive form and perfect information, we respectively denote by $M(\sigma)$ and $M[\sigma]$ the outcome and the history of M under σ.[5]

Definition 3. *Let $\mathscr{M}$ be an extensive-form perfect-information mechanism such that, for any perfect-knowledge QLU$^+$ context $C = (\mathcal{N}, \Omega, \mathcal{T}, \theta, \mathbb{C})$, the game $(C, \mathscr{M})$ has a unique subgame-perfect equilibrium.*

Then, we say that $\mathscr{M}$ is perfectly private *if there is an algorithm A such that, for any unique subgame-perfect equilibrium e of $\mathscr{M}$:*

$$A(\mathscr{M}(e)) = \mathscr{M}[e].$$

8 Additional Results

- *No Dominant Strategies.* Of course, another way to avoid equilibrium-selection problems consists of adopting a strictly dominant-strategy mechanism. Unfortunately, even disregarding collusion, complexity and privacy, we have proved that no (even weakly) dominant-strategy mechanism can guarantee any positive fraction of the maximum revenue in non-collusive QLU$^+$contexts [3].

- *Collectively Perfect Knowledge.* Our mechanism can be extended to a setting of *distributively perfect* knowledge. In such a setting, the true type of each player i needs not to be known by each individual player, nor should i's true type be known to at least another player. It suffices that every aspect of a player i is known to at least another player, but different players may know different aspects of i. More precisely, for each player i and each state ω, there exists at least a player $j_{i,\omega}$ who knows $\theta_i(\omega)$.

  In a sense, in a setting of distributively perfect knowledge, "each player is collectively known by the other players." In this more realistic setting, sticking again to the purest spirit of the field, we assume that the designer has no idea of who knows what about whom.

  We note that, even disregarding collusion, complexity, and privacy, the discussed mechanisms of Jackson, Palfrey, and Srivastava; of Abreu and Matsushima; and of Glazer and Perry, do not provide any revenue guarantees in this less demanding setting. (In fact, strictly speaking, they are applicable to such a setting.) By contrast, our mechanism works with a small complexity and privacy loss.

- *Arbitrary Social-Choice Functions.* Beyond revenue maximization, the first and third authors have been able to *virtually* and resiliently achieve any social-choice function.

  In a sense, they have enabled the Abreu-Matsushima and Glazer-Perry mechanisms to work without collusion, complexity and privacy problems.

[5] By history we mean the sequence of nodes —equivalently, the terminal node— of M's game tree reached during the play.

- *Imperfect Knowledge and a General Alternative Approach.* The additional results mentioned so far involve players perfectly informed about each other (whether individually or "collectively"), and have been obtained relying on a classical solution concept: indeed, unique sub-game perfect equilibrium. The first and third authors, have however proved that it is possible to design mechanisms resilient to collusion, complexity, and privacy, even in setting of *imperfect knowledge* (for instance, when the knowledge of the players about each other is *approximate*, or even when *no guarantee* exists on the knowledge (if any) that the players have about each other [4]). Examples of design problems for which such resilient mechanisms have been found currently include provisions of a public good, (unrestricted) combinatorial auctions, auctions of a single good in multiple supply, and sales of arbitrarily divisible goods. These resilient mechanisms, however, rely on a new solution concept, *implementation in resilient strategies*, that perhaps may be of independent interest.

Acknowledgements

We are grateful to Andy Postlewaite for challenging us to investigate dominant-strategies solutions for our problem. After proving their inadequacy, it became natural to start investigate alternative approaches.

We also thank Matthew Jackson for encouragement and suggestions, and Jacob Glazer for guiding us computer scientists through the rich and beautiful literature relevant to this work.

References

[1] D. Abreu and H. Matsushima. Virtual Implementation in Iteratively Undominated Strategies: Complete Information. *Econometrica*, Vol. 60, No. 5, pages 993-1008, Sep., 1992.

[2] L.M. Ausubel and P. Milgrom. The Lovely but Lonely Vickrey Auction. *Combinatorial Auctions*, MIT Press, pages 17-40, 2006.

[3] J. Chen, A. Hassidim, and S. Micali. Robust Perfect Revenue From Perfectly Informed Players. *Innovations in Computer Science*, 2010. (A preliminary version of the present paper, but containing a sketch of the "No-Dominant Strategy Theorem".)

[4] J. Chen and S. Micali. A New Approach to Auctions and Resilient Mechanism Design. *41st ACM Symposium on Theory of Computing* (STOC'09), pages 503-512, 2009. Full version available at http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design/.

[5] J. Glazer and M. Perry. Virtual Implementation in Backwards Induction. *Games and Economic Behavior*, Vo.15, pages 27-32, 1996.

[6] S. Izmalkov, M. Lepinski, and S. Micali. Perfect Implementation of Normal-Form Mechanisms. *MIT-CSAIL-TR-2008-028*, May 2008. (A preliminary version appeared in FOCS'05, pages 585-595, with title "Rational Secure Computation and Ideal Mechanism Design".)

[7] M.O. Jackson, T. Palfrey and S. Srivastava. Undominated Nash Implementation in Bounded Mechanisms. *Games and Economic Behavior*, 1994.