

Robust Perfect Revenue From Perfectly Informed Players

Jing Chen¹ Avinatan Hassidim² Silvio Micali¹

¹ CSAIL, MIT, Cambridge, MA 02139, USA

² RLE, MIT, Cambridge, MA 02139, USA

jingchen@csail.mit.edu avinatanh@gmail.com silvio@csail.mit.edu

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 vulnerable to equilibrium selection, collusion, privacy and complexity problems, and therefore far from guaranteeing that maximum revenue will be obtained. In this paper we both clarify and rectify this situation by

- Proving that no *weakly dominant-strategy* mechanism (traditionally considered immune to equilibrium selection) guarantees an arbitrarily small fraction of the maximum possible revenue;
- and, more importantly,
- Constructing a new mechanism, of extensive-form and with a *unique* sub-game-perfect equilibrium, which
 - (a) guarantees a fraction arbitrarily close to 1 of the maximum possible revenue;
 - (b) is provably robust against equilibrium selection, collusion, complexity, and privacy problems; and
 - (c) works for any number of players $n > 1$, and without relying on special conditions for the players utilities.

Keywords: perfect revenue, privacy, collusion resiliency; perfect player-knowledge; and robust 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 *mechanism* M describes the *strategies* available to the players, and how strategies determine outcomes. Together, a context C and a mechanism M define a *game* G , $G = (C, M)$, in which each rational player will endeavor to choose his own strategy so as to maximize his own utility.

MECHANISM DESIGN. Mechanism design aims at finding a mechanism M such that, for any context C (or any C in a given class), a desired property \mathbb{P} holds for the outcomes of the game (C, M) , when rationally played. The difficulty is that the designer does not exactly know the players' preferences, while \mathbb{P} typically depends on such preferences. In the purest form of mechanism design, *all knowledge about the players lies with the players themselves*. The designer can count only on the players' rationality. And based solely on this fact, he must design M so that it becomes "in the best interest of the players" to satisfy \mathbb{P} . That is, he must ensure that \mathbb{P} holds in a rational play of M . But: *What is a rational play?*

THE CLASSICAL INTERPRETATION OF A RATIONAL PLAY. The classical interpretation of a rational play is an *equilibrium*, that is a profile of strategies $\sigma = \sigma_1, \dots, \sigma_n$ such that no player i has an incentive to

deviate from his specified strategy 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 specified σ_j* . That is, Nash equilibrium only guarantees that i prefers $\sigma_1, \dots, \sigma_i, \dots, \sigma_n$ to $\sigma_1, \dots, \sigma'_i, \dots, \sigma_n$. If σ is a *dominant-strategy equilibrium*, the strongest form of equilibrium, then, for any player i , σ_i is i 's best strategy 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 when playing σ_i is strictly larger than (respectively, larger than or equal to) his utility when playing σ'_i .

1.2 Our Goal

This paper focuses on a classical context: *quasi-linear utilities with non-negative valuations*. Namely, there are finitely many possible *states*, $\omega_1, \dots, \omega_k$, including the *null state*, which every player values 0; each player i has 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 $v_i(\omega) - P_i$. (The *revenue* of an outcome (ω, P) consists of $\sum_i P_i$. The function v_i is i 's *valuation*.)

Such context models a great deal of situations. For

instance, in an auction of multiple goods, a state ω represents which player wins which items. Accordingly, the utility of player i in an outcome (ω, P) naturally is his value for the items he gets in ω , minus the price he pays. In another example, each state ω represents one of finitely many ways of building a bridge across a given river. Accordingly, and naturally too, each player has different values for each possible bridge. (For instance, a player's value for a given potential bridge may depend on how distant it would be from his house.) The list of examples could go on and on. In all of them, however, no matter what the mechanisms may be, it is also natural for different subsets of the players to collude—that is, to coordinate their strategies—so as to improve their utilities.

In such a classical context, our goal is equally classical: getting an outcome of maximum revenue when the players have *perfect knowledge*. That is, when each player knows the valuations of all players (as well as who colludes with whom, if collusion exists among the players).

When the players' knowledge is best possible, it is natural to ask whether the best possible revenue can be obtained. Note that, without the ability of imposing arbitrary prices, the best possible revenue that a mechanism can hope to get from rational players is the *maximum social welfare*, that is, $\max_{\omega} \sum_i v_i(\omega)$. Thus:

Can a mechanism guarantee perfect revenue from perfectly informed players?

1.3 Four Main Obstacles

Plenty of mechanisms have been proposed for our goal. Yet, *none of them achieves it in a robust way*. Four main obstacles stand on their way. Let us explain.

Equilibrium Selection

It should be realized that designing a mechanism so as to guarantee a property \mathbb{P} "at a Nash Equilibrium" is a weak guarantee. 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, 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.

Collusion

The problem of collusion in mechanism design is well recognized. The problem occurs for obvious 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, by "guaranteeing" a property \mathbb{P} at equilibrium, a classical mechanism is typically vulnerable to collusion. 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 two sufficiently informed players can totally destroy the economic efficiency of the famous VCG mechanism [9, 12, 26], although it too is dominant-strategy.

Complexity

Traditional mechanism design disregards the "complexity of a mechanism." A mechanism of normal-form may require the players to simultaneously announce exponentially long strings. And a mechanism of extensive form may require the players to act over exponentially many rounds. In both cases, therefore, such mechanisms in practice fail to reach their objectives, no matter what their theoretical claims, unless their contexts are extremely "tiny."

Privacy

Privacy has been traditionally neglected in mechanism design, and considered a *quite separate* desideratum: nice to have perhaps, but not central for an incentive analysis. Yet, as especially argued by [13], it has a great potential to distort incentives, and thus to derail classical mechanisms from achieving their desired properties. A mechanism typically neglects privacy by requiring the players to reveal a lot of information about themselves. But if the players value privacy (which by definition implies that divulging their secret information causes them to receive a *negative utility*), then the mechanism gives them both positive and negative incentives, and it is no longer clear how these opposing forces will balance out.

1.4 Prior Mechanisms

Let us now discuss the most relevant mechanisms for our design problem. We start with the traditional "at equilibrium" approach, and then proceed to more sophisticated ones.

The Generic Mechanism

The following may be the first mechanism that comes to mind for our context.

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 be his reported value for ω (possibly minus a small discount ϵ to encourage i 's participation). If not all reports coincide, then choose the "null outcome" (which all players are assumed to value 0) and price 0 for every player.*

Unfortunately, HOPE-FOR-THE-BEST is extremely vulnerable to equilibrium selection. It is trivial to see that the strategy profile in which each player reports all true valuations is a Nash equilibrium for HOPE-FOR-THE-BEST, indeed, it is the *truthful equilibrium*. It is also trivial to see that in this equilibrium the revenue is the maximum possible (disregarding the negligible quantity $n\epsilon$). Notice too, however, that HOPE-FOR-THE-BEST also has *additional* equilibria, E_2, E_3, \dots , 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,

- In the truthful equilibrium E_1 the designer is "happy", but the players are "sad", while
- in all other equilibria E_x the players are "happy" and the designer is "sad."

This being the case: *which equilibrium E_x is more likely to be selected?* Further, while each E_x at least maximizes social welfare, in plenty of other equilibria both revenue and social welfare are quite poor.¹ 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 be an equilibrium and generate any revenue at all?*

The JPS Mechanism

Jackson, Palfrey, and Srivastava [15] provided a quite different mechanism. Again, their mechanism yields optimal revenue only at the truthful equilibrium τ . But this time τ is a much more meaningful equilibrium: *it is the only Nash equilibrium composed*

¹Let ω be any state such that $v_i(\omega) \gg c > 0$ for all players i . And let σ be the strategy profile, where each σ_j consists of reporting that all players have the following valuation v : $v(\omega) = c$ and $v(x) = 0$ for any state $x \neq \omega$. Then, it is easy to see that σ is an equilibrium. Moreover, the revenue of σ is cn , and the social welfare of σ is $\sum_i v_i(\omega)$.

of weakly undominated strategies. Somewhat counterintuitively, however, their solution too is vulnerable to equilibrium-selection. The point is that, as in HOPE-FOR-THE-BEST, there are plenty of equilibria σ that generate smaller revenue while being more attractive to *all* players. Again too, 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 . Accordingly, *the JPS mechanism too is very vulnerable to equilibrium selection.*

What has happened? At a first glance, the JPS mechanism looks very "robust," because "no one should want to play a weakly dominated strategy." But the problem is that the process of eliminating all *weakly* dominated strategies for yourself and the other players is not well defined. Unlike the iterated elimination of *strictly* dominated strategies, the iterated elimination of weakly dominated strategies depends on the order of elimination. For example, if one eliminates first "suicidal strategies" (in fact, if one eliminates first "suicide" for just another one of the players), then all equilibria become equally reasonable, and the attractive ones from the players' point of view are those generating less revenue.

In addition, *the JPS mechanism is totally vulnerable to collusion.* Indeed, it enables some pairs of players (i, j) to jointly deviate from the truthful equilibrium so as to improve the utility of i without hurting that of j . And when they so deviate its revenue cannot be maximum.

Finally, *the JPS mechanism is totally vulnerable to privacy,* because it relies on the players revealing all their knowledge.

The AM (and GP) Mechanisms

Assuming that there are at least 3 players and that some more technical conditions hold, Abreu and Matsushima [1] present a general normal-form mechanism that guarantees that essentially any desired property (including ours) is satisfied in a perfect-knowledge context. Their mechanism is robust against equilibrium selection, because after the iterated elimination

of *strictly* dominated strategies, each player is left with a single (and truthful) strategy, and thus the resulting game has a single equilibrium. However, the *AM mechanism is highly vulnerable to collusion, complexity, and privacy problems*. That is, (1) it no longer guarantees its desired property when any two players jointly deviate from their equilibrium strategies; (2) it requires the players to announce a doubly-exponential number of bits even when there are a constant number of outcomes, a constant number of player, and each player can have one of two types; and (3) it relies on the players revealing all the knowledge in their possession.

A variant of the AM mechanism was put forward by Glazer and Perry [11]. The GP mechanism is of extensive form: informally this means that the players act one at a time, over several rounds. The GP mechanism did not suffer from any equilibrium selection problems either, because its corresponding game admits a *unique* subgame-perfect equilibrium. Essentially, this means that, at each decision node of the game tree, every acting player has a single best action available to him. However, the GP mechanism continues to be vulnerable to collusion, complexity and privacy problems. The vulnerability to collusion and privacy is essentially identical to that of the AM mechanism. Complexity wise the GP mechanisms actually requires exponentially many rounds of communication. A justification of this fact is presented in the appendix.

1.5 Our Results

Our Impossibility Result

The problem of equilibrium selection fully disappears when a mechanism achieves its desired property \mathbb{P} at a strictly dominant-strategy equilibrium, while still “lurks around” for weakly dominant-strategy equilibria. Unfortunately, we prove that neither strong nor weakly dominant-strategy mechanisms exist that can guarantee perfect revenue from perfectly informed players. Worse, our impossibility result holds even even if the mechanism designer were content to generate an *arbitrarily small fraction* of the optimal revenue. In sum, we prove the following.

Thm 1: *No weakly dominant-strategy mechanism guarantees a fraction ϵ of the optimal revenue.*

Our impossibility theorem shows that, in order to guarantee perfect revenue, we must adopt a different solution concept.

Our Possibility Result

To enable the design of mechanisms reasonably resilient against equilibrium selection, collusion, com-

plexity, and privacy, the first and third author have developed *rational robustness*, a new solution concept and implementation notion [7]. (Essentially, rational robustness is based on *iterated elimination of distinguishably dominated strategies*, a new notion “in between” iterated elimination of *strictly* and *weakly* dominated strategies.) But in our perfect knowledge context we are able to achieve perfect revenue and perfect resiliency, by means of an extensive-form mechanism, under the classical solution concept of unique subgame-perfect equilibrium (which indeed is a special case of rationally robust implementation).

Informal Thm 2: *There exists an extensive-form mechanism \mathcal{M} , guaranteeing a fraction $1 - \epsilon$ of the optimal revenue, that*

- *Works for any number of players $n > 1$;*
- *Has a unique subgame perfect equilibrium;*
- *Enjoys perfect collusion resilience;*
- *Enjoys perfect privacy;*
- *Does not require any trust from the players; and*
- *Has perfect communication complexity and $n + 1$ rounds.*

Remarks

- Note that \mathcal{M} works with just two players. By contrast, the mechanism of Jackson, Palfrey and Srivasta, that of Abreu and Matsushima, that of Glazer and Perry, and the generic mechanism all require more than 2 players. In essence, because this requirement facilitates (or makes it possible and meaningful) to identify which player deviated from his equilibrium strategies.
- We stress the word “unique” because a game with multiple subgame-perfect equilibria can still suffer from equilibrium-selection problems.
- Perhaps interestingly, in the presence of collusive players, our mechanism \mathcal{M} has multiple ways to be truthful, but *only one* of them is a subgame-perfect equilibrium.
- By saying that \mathcal{M} is perfectly resilient to collusion we mean two things. First, \mathcal{M} guarantees perfect revenue as long as not all players belong to the same coalition, and each coalition acts rationally. (In our setting, a rational coalition maximizes the sum of the individual utilities of its members. Only when the players have imperfect knowledge about each other, one may want to consider weaker models of coalition rationality.) Second, \mathcal{M} achieves perfect revenue no matter how well players belonging to the same collusive set, if any, may coordinate

their actions. In particular, \mathcal{M} works even when such players are free to make side-payments to each other and/or to enter into binding contracts with each other.

- By saying that \mathcal{M} is perfectly private we essentially mean that, in any rational play, nothing can be learned about the players' valuations, by the mechanism designer or any observer of the play, except for what is deducible from a perfect-revenue outcome. Of course, our \mathcal{M} can be so "perfect" only because we are dealing with perfectly informed players (so that the only privacy concern is with respect to the "outside world"). But this *is* our setting, and thus one has both the right to demand and the obligation to deliver as a perfect solution as possible.
- By saying that \mathcal{M} does not require any trust from the players we mean several things. First, \mathcal{M} is not a *mediated* mechanism. (Indeed, privacy would be easy to achieve if the players could confide their strategies or their preferences to mediator trusted to announce the right outcome and never to reveal to anyone else any information received from the players). Second, \mathcal{M} does not rely on any complexity assumptions, as needed for running a cryptographically secure protocol. (After all, at least some of such assumptions may turn out to be false.) Third, \mathcal{M} does not rely on the security of some underlying communication channels. Fourth, \mathcal{M} does not rely on the "honesty" of even some of the players.

Indeed, in our mechanism \mathcal{M} any action of the players becomes public as soon as it is taken.

- By saying that \mathcal{M} has perfect communication complexity we mean that \mathcal{M} 's players need to announce, altogether, the same number of bits necessary to describe the desired outcome.

In addition to "mechanism complexity", one may consider also "player complexity," that is the time required to a player to figure out and thus choose a rational strategy. Here it is worth pointing out that, to play rationally \mathcal{M} , a player performs a computation linear in the number of states —and the number of players. This is essentially optimal given the generality of our setting. (Only with respect to a specific choice of states, one can consider whether there exists a compact representation of the states together with a compact representation of the players' utilities for them.)

Comparison with other work

- Note that our notion of collusion resiliency is stronger than that offered by other mechanisms. In particular, *group —or coalition— strategyproofness* [3, 16, 19, 21, 25] rules out collusion, but only under the assumption that the players are not able to make side payments to each other. Without restricting how players might cooperate, *t-truthful mechanisms* [10] offer protection against coalitions of at most t players, but only for single-value games. (In such games, a player i values some outcomes 0, and all other outcomes a fixed value v_i .) Again without restricting cooperation abilities, *collusion neutralization* [6, 20] offers collusion protection in more general games, but their notion too is weaker than the one considered in our paper. (Protection against the coalition of all players has also been considered and achieved, but only in Bayesian settings, where the distributions of player preferences are known to everyone, including the mechanism designer [4, 5, 17, 18].) Finally, a different approach altogether, *collusion leveraging*, has been submitted to this same conference [8].
- Some work on privacy preserving mechanisms has already started. However, the privacy is either limited or gained by adding an additional layer to the mechanism —such as one or more mediators, envelopes, or encryption— [13, 22, 23]. By contrast, our mechanism \mathcal{M} achieves perfect privacy without relying on any additional infrastructure. Indeed, \mathcal{M} works by asking the players to take only *public actions*.

Easy Variants and Forthcoming Work

By hindsight, it is easy to modify our mechanism in various ways —or even the original Abreu-Matsushima or Glazer-Perry mechanisms— so as to keep our perfect robustness against collusion, complexity, and privacy, while gaining some perceived additional advantage, such as reducing the number of steps, achieving *additive* revenue approximation, etc.

There are, however, quite important variants that we would like to point out now, and develop in later versions of this paper. Namely, we can achieve perfect revenue not from just perfectly informed players (as in this paper), but also from players with perfect *distributed* knowledge. For instance, it suffices that for each fact about a player i (i.e., for each $TV_i(\omega)$) there exists an additional player who knows this fact.

2 Preliminaries

Our Contexts

We work with reasonably general contexts with quasi-linear utilities, where the players can be collusive and are perfectly informed about each other. More formally,

Definition 1. A perfect-knowledge context C is identified by four quantities, $C = (N, \Omega, TV, \mathbb{C})$, where

- N is the finite set of players, $N = \{1, \dots, n\}$
- Ω is the finite set of states, which includes the empty state, \perp .

The set Ω defines the set of outcomes: namely, $\Omega \times \mathbb{R}^n$. It also defines the set V of valuation profiles v : namely, each valuation v_i is a function from Ω to non-negative reals such that $v_i(\perp) = 0$.

- TV is the profile of true valuations (or types): namely, each valuation TV_i describes player i 's actual value for each possible state.

Each TV_i defines the utility function u_i of player i as follows: for each outcome (ω, P) , $u_i(\omega, P) = TV_i(\omega) - P_i$. That is, i 's utility is his true value for the state minus the price he pays. The profile P is referred to as a price profile.

- \mathbb{C} is the collusion structure: namely, a partition of N .

If S is a subset in \mathbb{C} , then S is the maximal subset of players colluding with each other. 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, and collusive otherwise.

Each independent player tries to maximize his own utility function, and each collusive set tries to maximize the sum of the utilities of its members.

The sets N and Ω are common knowledge to everybody; while the profile TV and the partition \mathbb{C} are only common knowledge to the players.

We stress that the mechanism designer has no knowledge about TV (or \mathbb{C})! In other words, we adhere to the classic spirit of mechanism design, where all knowledge lies with just the players.

Our Mechanisms

Recall that each mechanism M must specify the players' strategies, including the *opt-out strategy* (for each player i this strategy is always denoted by OUT_i), and the outcome (or distribution over outcomes if M is probabilistic), $M(\sigma)$, associated to each possible strategy profile σ . In addition, we insist that each mechanism M must satisfy the following

OPT-OUT CONDITION: $M(\sigma) = (\perp, (0, \dots, 0))$ whenever $\sigma_i = \text{OUT}_i$ for some player i .

For conciseness, we refer to a profile of strategies as a *play*. The expected utility of player i in a play σ is $\mathbb{E}[u_i(M(\sigma))]$. As announced, our mechanisms are finite, of extensive form, and *public-action*, that is, each of our M specifies a game tree, where exactly one player acts at each node, knowing all actions played so far.

Definition 2. (Unique Subgame Perfect Equilibrium) Let $G = (C, M)$ be a game where M is an extensive-form public-action mechanism. G has a unique subgame perfect equilibrium if there is only one strategy profile σ satisfying the following property:

- 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 .

Note that essentially, G has a unique subgame perfect equilibrium if at each decision node of G 's game tree, every acting player has a single best action available to him, given that all players will choose their best action in every decision node thereafter.

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 $C = (N, \Omega, TV, \mathbb{C})$ is defined to coincide with the *maximum social welfare* (MSW for short), that is, $\max_{\omega} \sum_i TV_i(\omega)$.

3 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.

Definition 3. A DST mechanism M guarantees a fraction ϵ of the maximum rational revenue if for any context $C = (N, \Omega, TV)$ we have

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

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 $M(v^1, \dots, v^n) = (\omega, P)$, 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 (\star) 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 $M(K, K) = (\omega, Q)$, 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.*

4 Our Mechanism

Notation

In the mechanism below,

- ϵ and ϵ_j^i , for $i \in \{2, \dots, n\}$ and $j \in \{1, \dots, n\}$, are constants such that $\frac{1}{5n} > \epsilon > \epsilon_1^2 > \dots > \epsilon_n^2 > \epsilon_1^3 > \dots > \epsilon_n^3 > \dots > \epsilon_1^n > \dots > \epsilon_n^n > 0$.
- Numbered steps are taken by the players, while steps marked by letters are taken by the mechanism.
- Sentences between quotation marks are comments, and could be excised if no clarification is needed.

- We denote by n_r the number of outcomes (ω, P) with revenue r . For all such outcomes, we denote by $0 \leq f_r(\omega, P) < n_r$ the rank of the outcome (ω, P) in the lexicographic order that first considers the state and then the price profile (where P_1, \dots, P_n precedes $P'_1 \dots P'_n$ whenever $P_1 > P'_1$, etc.).

Mechanism \mathcal{M}

- (1) *Player 1 announces a state ω^* and a profile K^1 of natural numbers.*

“(ω^*, K^1) is player 1’s proposed outcome, allegedly an outcome of maximum revenue.”

- (a) *Set $\omega = \perp$, and $P_i = 0 \forall i$. If $\sum_i K_i^1 = 0$, the mechanism ends right now. Otherwise, proceed to Step 2.*

“Whenever the mechanism ends, ω and P will be, respectively, the final state and price profile.”

- (2, \dots , n) *In Step i , $2 \leq i \leq n$, player i publicly announces a profile Δ^i of natural numbers such that $\Delta_i^i = 0$.*

“By so doing i suggests to raise the current price of j , that is $K_j^1 + \sum_{\ell=2}^{i-1} \Delta_j^\ell$, by the amount Δ_j^i .”

- (b) *For each player i , publicly select bip_i and P_i^* as follows. Let $R_i = \{j : \Delta_j^i > 0\}$.*

If $R_i \neq \emptyset$, then bip_i is highest player in R_i , and $P_i^ = K_i^1 + \sum_{\ell=2}^{bip_i} \Delta_i^\ell$. Else, $bip_i = 1$ and $P_i^* = K_i^1$.*

“We refer to bip_i as the best informed player about i , and to P_i^* as the provisional price of i .”

- ($n+1$) *Each player i such that $P_i^* > 0$ simultaneously announces YES or NO.*

By default, each player i such that $P_i^ = 0$ announces YES, and player 1 announces YES if $bip_1 = 1$.*

“Each player i announces YES or NO to ω^* as the final state and to $P_i^* - \epsilon$ as his own price. (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 say YES, the updated proposal (ω^*, P^*) is implemented with probability 1. Else:

- With very high probability the null outcome is chosen, except that the best-informed players of those saying NO are punished.
- With small probability the null outcome is chosen

- With very small probability, proportional to the number of players saying YES, we implement (ω^*, P^*) as if all said YES.

Importantly, as we shall see, all get a small reward at the end for their knowledge.”

- (c) *Let Y be the number of players announcing YES. If $Y = n$, then reset ω to ω^* and each P_i to $P_i^* - \epsilon$, and go to Step g. If $Y < n$, proceed to Step d.*
- (d) *Publicly flip a biased coin c_1 such that $\Pr[c_1 = Heads] = 1 - \epsilon$.*
- (e) *If $c_1 = Heads$, reset P_{bip_i} to $P_{bip_i} + 2P_i^*$ for each player i announcing NO.*
- (f) *If $c_1 = Tails$, letting $B = \sum_{i \text{ announces NO}} P_i^*$, flip a biased coin c_2 such that $\Pr[c_2 = Heads] = \frac{Y}{nB}$. If $c_2 = Heads$, reset ω to ω^* and each P_i to $P_i^* - \epsilon$. If $c_2 = Tails$, ω and P continue to be \perp and 0^n .*
- (g) *Reset P_1 to $P_1 - \epsilon - 2\epsilon \sum_j K_j^1 + \epsilon \frac{f_r(\omega^*)}{n_r}$ and each other P_i to $P_i - \epsilon - \sum_j \epsilon_j^i \Delta_j^i$.*

“Although players’ prices may be negative, we prove that the mechanism never loses money, and that in the unique rational play the utility of every player is non-negative. For clarity, our rewards are proportional to prices and raises.”

5 Analysis of Our Mechanism

Mechanism \mathcal{M} induces a game G whose game tree has height $n+1$, and where only players act at each internal node. (The mechanism tosses all its coins at leaf nodes, that are defined to be of height 0.) At each node of height 1 all players act simultaneously, and at every other internal node only a single player acts. Specifically, at each node of height $h \geq 2$ the only acting player is player

$$i_h \triangleq n - h + 2.$$

For each internal node N , we denote by G^N the subgame rooted at N . Recall that a strategy σ_i of player i in G specifies, for each node at which i acts, which action i chooses among all those available to him. By σ_i^N we denote the restriction of σ_i to subgame G^N . Given a restricted strategy profile σ^N for G^N , the outcome of \mathcal{M} obtained by executing σ^N is denoted by $\mathcal{M}(\sigma^N)$.

For uniformity, we find it sometimes convenient to assume that every player i belongs to a (necessarily unique) collusive set, denoted by \mathcal{C}_i . If i is independent, then $\mathcal{C}_i = \{i\}$.

5.1 Statements of Our Lemmas

Lemma 1. *If N is a node of height 1, then G^N has a unique subgame-perfect equilibrium σ^N , where*

- *If i is independent, then σ_i^N consists of announcing YES if and only if $TV_i(\omega^*) \geq P_i^*$.*
- *If i belongs to a coalition \mathcal{C} , then σ_i^N consists of announcing YES if and only if*

$$bip_i \in \mathcal{C} \quad \text{or} \quad \sum_{j \in \mathcal{C}} TV_j(\omega^*) \geq \sum_{j \in \mathcal{C}} P_j^*.$$

The proof of this lemma is based on the fact that the probability that an outcome is executed is monotone with the number of players who announce YES. Thus, it is strictly dominant to announce YES, if and only if the player has positive utility from this outcome and price.

Lemma 2. *Let N be a node of height $h \in [2, n]$, $i = i_h$, and $\mathcal{C} = \mathcal{C}_i$. Then G^N has a unique subgame-perfect equilibrium where i acts as follows at node N : For each collusive set $\mathcal{D} \neq \mathcal{C}$,*

1. *if*

$$\sum_{j \in \mathcal{D}} \left(K_j^1 + \sum_{\ell=2}^{i-1} \Delta_j^\ell \right) \geq \sum_{j \in \mathcal{D}} TV_j(\omega^*)$$

then i announces $\Delta_j^i = 0$ for all $j \in \mathcal{D}$;

2. *if*

$$\sum_{j \in \mathcal{D}} \left(K_j^1 + \sum_{\ell=2}^{i-1} \Delta_j^\ell \right) < \sum_{j \in \mathcal{D}} TV_j(\omega^*)$$

then letting k be the minimal player in \mathcal{D} , i announces $\Delta_j^i = 0$ for all $j \in \mathcal{D} \setminus \{k\}$ and

$$\Delta_k^i = \sum_{j \in \mathcal{D}} \left(TV_j(\omega^*) - K_j^1 - \sum_{\ell=2}^{i-1} \Delta_j^\ell \right).$$

For his own collusive set \mathcal{C} ,

1. *if*

$$\sum_{j \in \mathcal{C}} \left(K_j^1 + \sum_{\ell=2}^{i-1} \Delta_j^\ell \right) \geq \sum_{j \in \mathcal{C}} TV_j(\omega^*) \quad \text{or}$$

it is the case that $k \in \mathcal{C}$ for all $k > i$,

then i announces $\Delta_j^i = 0$ for all $j \in \mathcal{C}$;

2. *if*

$$\sum_{j \in \mathcal{C}} \left(K_j^1 + \sum_{\ell=2}^{i-1} \Delta_j^\ell \right) < \sum_{j \in \mathcal{C}} TV_j(\omega^*) \quad \text{and}$$

there exists player $j > i$ such that $j \notin \mathcal{C}$,

then letting k be the minimal player in $\mathcal{C} \setminus \{i\}$, i announces

$$\Delta_k^i = \sum_{j \in \mathcal{C}} \left(TV_j(\omega^*) - K_j^1 - \sum_{\ell=2}^{i-1} \Delta_j^\ell \right).$$

This lemma is technically involved, but conceptually simple. First, we show that a player i never wants to “overbid,” that is raise the price of another player j to more than j ’s true valuation for the proposed state ω^* . When j is independent, this holds because we know that j will announce NO to any price above his true valuation, and thus no player after i will want to further raise j ’s price. Therefore, overbidding on j will cause i to be punished. Care must still be taken to verify the Step- g rewards of i and j will not change this simple analysis. (For example j will not accept a higher price in order to get more reward for volunteering his knowledge about others.) For coalitions, the argument is more subtle.

After ruling out overbidding, we also show that a player i never wants to “underbid,” that is not raise the price of a player j when it is below j ’s true valuation for the proposed state. Again, this is easier to argue for independent players. Arguing this point for coalitions is the only time that requires exploiting the n^2 reward values $\epsilon_{i,j}$.

Lemma 3. *Let N be the root of the tree (so that $G^N = G$), then G has a unique subgame-perfect equilibrium where player 1 acts as follows at node N :*

1. *player 1 announces ω^* , the lexicographically first state ω such that $\sum_\ell TV_\ell(\omega) = MSW$;*
2. *for each collusive set \mathcal{D} , letting i be the minimal player in \mathcal{D} , player 1 announces $K_i^1 = \sum_{j \in \mathcal{D}} TV_j(\omega^*)$, and $K_j^1 = 0$ for each $j \in \mathcal{D} \setminus \{i\}$.*

The proof of this lemma is also done in two stages. First, given Lemma 2, we prove that it is dominant for player 1 to set the prices correctly (although not exactly truthfully in the case of a coalition). Finally, as the prices are set correctly, choosing the outcome which maximizes the total welfare dominates any other course of action.

Detailed proofs of our lemmas are in the full version of this paper, available at http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design/.

5.2 Our Main Theorem

Theorem 2. *Let σ be the unique subgame perfect equilibrium of G , and let $(\omega, P) = \mathcal{M}(\sigma)$. Then:*

- (1) $\sum_i TV_i(\omega) = MSW$, and

$$(2) \sum_i P_i \geq (1 - 4\epsilon n)MSW.$$

Proof. In execution σ , by Lemma 3, player 1 announces ω^* such that $\sum_\ell TV_\ell(\omega^*) = MSW$ and, for each coalition \mathcal{D} , also announces $K_\ell^1 = \sum_{j \in \mathcal{D}} TV_j(\omega^*)$, where ℓ is the minimal player in \mathcal{D} . Thus $\sum_i K_i^1 = MSW$.

If $MSW = 0$, then $\sum_i K_i^1 = 0$ and \mathcal{M} ends at Step a, with $\omega = \perp$ and $P_i = 0$ for each player i . Therefore $\sum_i TV_i(\omega) = \sum_i TV_i(\perp) = 0 = MSW$ and $\sum_i P_i = 0 = MSW$.

If $MSW > 0$, then $\sum_i K_i^1 > 0$ and \mathcal{M} ends at Step g. By Lemma 2, for each player $i \neq 1$, i announces $\Delta_k^i = 0$ for each k . Therefore for each player i , $bip_i = 1$. Furthermore, the total price for each coalition \mathcal{D} equals \mathcal{D} 's total true valuation for ω^* : that is, $\sum_{\ell \in \mathcal{D}} P_\ell^* = \sum_{\ell \in \mathcal{D}} K_\ell^1 = \sum_{\ell \in \mathcal{D}} TV_\ell(\omega^*)$. By Lemma 1, every player in \mathcal{D} announces YES in Step $n + 1$. This implies that, at the end of Step c we have $Y = n$, $\omega = \omega^*$, and, for each coalition \mathcal{D} , $\sum_{\ell \in \mathcal{D}} P_\ell = \sum_{\ell \in \mathcal{D}} P_\ell^* - |\mathcal{D}|\epsilon = \sum_{\ell \in \mathcal{D}} TV_\ell(\omega^*) - |\mathcal{D}|\epsilon$. Because $Y = n$, the execution of \mathcal{M} will then proceed directly to Step g, which does not reset the current state. Thus we have that

$$\sum_i TV_i(\omega) = \sum_i TV_i(\omega^*) = MSW.$$

Because the reward given to each player $i > 1$ in Step g is ϵ , and player 1 gets at most $\epsilon + 2\epsilon MSW$, then the final revenue of the mechanism is

$$\begin{aligned} \sum_i P_i &> \left(\sum_i TV_i(\omega^*) - n\epsilon \right) - (n-1)\epsilon \\ &\quad - \epsilon - 2\epsilon MSW \\ &> (1 - 4\epsilon n)MSW, \end{aligned}$$

where we parenthesized the prices after step c, and used that MSW is integer and thus $MSW \geq 1$. *Q.E.D.*

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] S. Barbera and M.O. Jackson. Strategy-Proof Exchange. *Econometrica*, Vol.63, No.1, pages 51-87, Jan., 1995.
- [4] Y. Che and J. Kim. Robustly Collusion-Proof Implementation. *Econometrica*, Vol.74, No.4, pages 1063-1107, Jul., 2006.
- [5] Y. Che and J. Kim. Optimal collusion-proof auctions. *Journal of Economic Theory*, Vol.144, Issue 2, pages 565-603, Mar. 2009.
- [6] J. Chen and S. Micali. A New Approach to Auctions and Resilient Mechanism Design. *Symposium on Theory of Computing*, pages 503-512, 2009. Full version available at [http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design/](http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design/).
- [7] J. Chen and S. Micali. Rational Robustness for Mechanism Design. Submitted to STOC'10, full version available at [http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design/](http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design/).
- [8] J. Chen S. Micali, and P. Valiant. Leveraging Collusion in Combinatorial Auctions. Submitted to ICS2010.
- [9] E.H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, Vol.11, No.1, pages 17-33, Sep., 1971.
- [10] A. Goldberg and J. Hartline. Collusion-Resistant Mechanisms for Single-Parameter Agents. *Symposium on Discrete Algorithms*, pages 620-629, 2005.
- [11] J. Glazer and M. Perry. Virtual Implementation in Backwards Induction. *Games and Economic Behavior*, Vo.15, pages 27-32, 1996.
- [12] T. Groves. Incentives in Teams. *Econometrica*, Vol. 41, No. 4, pages 617-631, Jul., 1973.
- [13] 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".)
- [14] M. Jackson. Implementation in Undominated Strategies: a Look at Bounded Mechanisms. *Review of Economic Studies*, Vol. 59, No. 201, pages 757-775, Oct., 1992.
- [15] M.O. Jackson, T. Palfrey and S. Srivastava. Undominated Nash Implementation in Bounded Mechanisms. *Games and Economic Behavior*, 1994.
- [16] K. Jain and V. Vazirani. Applications of Approximation Algorithms to Cooperative Games. *Symposium on Theory of Computing*, pages 364-372, 2001.
- [17] J.J. Laffont and D. Martimort. Collusion under Asymmetric Information. *Econometrica*, Vol.65, No.4, pages 875-911, Jul. 1997.
- [18] J.J. Laffont and D. Martimort. Mechanism Design with Collusion and Correlation. *Econometrica*, Vol. 68, Issue 2, pages 309-342, Mar. 2000.
- [19] H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice Welfare*, Vol.16, No.2, pages 279-320, Feb. 1999.
- [20] S. Micali and P. Valiant. Resilient Mechanisms for Unrestricted Combinatorial Auctions, submitted to FOCS 2007; SODA 2008; and STOC 2009. Available at [http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design](http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design)
- [21] H. Moulin and S. Shenker. Strategyproof Sharing

- of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, Vol.18, No.3, pages 511-533, 2001.
- [22] F. McSherry and K. Talwar. Mechanism Design via Differential Privacy. *FOCS07*, 2007.
- [23] M. Naor, B. Pinkas and R. Sumner. Privacy Preserving Auctions and Mechanism Design, *EC99*, 1999.
- [24] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*, MIT Press, 1994.
- [25] M. Pal and E. Tardo, Group Strategyproof Mechanisms via Primal-Dual Algorithms. *44th Annual IEEE Symposium on Foundations of Computer Science (FOCS'03)*, pages 584-593, 2003.
- [26] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, Vol. 16, No. 1, pages 8-37, Mar., 1961.

A The Complexity of the GP Mechanism

In this appendix we give an example in which the GP mechanism requires an exponential number of rounds, although the number outcomes is two, there are five players, each player can have at most two types, and the mechanism is only required to succeed with constant probability (The same example also requires a double-exponential amount of communication in the AM mechanism).

Context Sketch

We give an informal description of the context of the AM [1] and GP [11, 24] mechanisms. Let n denote the number of players. Each player p_i has a set T_i of all possible types for this player. The set T_i is finite, and is known to the designer. In the example we present $|T_i| = 2$ for every i . There is also a set of outcomes A , and for $t_i \in T_i$, and $\alpha \in A$ we denote by $t_i(\alpha)$ how much player i values the outcome α if his type is t_i . The GP mechanism requires that if $s, t \in T_i$, then there exists an outcome α such that $s(\alpha) \neq t(\alpha)$. In addition, each player has a true type $\Theta_i \in T_i$, which is unknown to the designer.

The designer is also given as input a *Social Choice Function*, which is a function from the type of all the players to outcomes

$$SCC : T_1 \times T_2 \times \dots \times T_n \rightarrow A.$$

The goal of the designer is to guarantee that with probability at least $1 - \epsilon$ the output of the mechanism is $SCC(\Theta_1, \dots, \Theta_n)$.

Sketch of the GP Mechanism

Following Abreu and Matsushima, Glazer and Perry design rely on a function f from types of the players to probability distribution over outcomes,

such that for every two types $t_i, \tilde{t}_i \in T_i$, we have that if the player's true valuation is t_i , then he strictly prefers $f(t_i)$ to $f(\tilde{t}_i)$. The GP mechanism then proceeds by performing $k+1$ rounds, where in each round all the players speak, one after the other. In the first k rounds, every player declares the types of all the players, while in the last round each player declares only his own type. Finally, the mechanism chooses the outcome in the following manner:

1. Choosing the outcome:
 - (a) With probability ϵ : round $k+1$ is chosen. Pick a random player i , and let t_i denote his declaration in round $k+1$. The outcome is $f(t_i)$ up to fines (which will be described later)
 - (b) With probability $1 - \epsilon$: pick a random round $1 \leq j \leq k$. If at least $n - 1$ players declared the same vector of types (t_1, \dots, t_n) in this round, select the outcome $SCC((t_1, \dots, t_n))$. Else, pick an arbitrary outcome O .
2. Choosing the fines:
 - (a) Let t_1, \dots, t_n be the vector of types declared by the players in the final round.
 - (b) Let p be the last player not to declare t_1, \dots, t_n , in the previous k rounds. That player pays a small fine δ .

Analysis of their Mechanism

In order for this scheme to work, one must be careful about the choice of parameters. To define these parameters, for $t_i, \tilde{t}_i \in T_i$, let $L_f(t_i, t_j) = U(t_i, f(t_i)) - U(t_i, f(t_j))$ be the advantage that a player whose true type is t_i gains from reporting t_i and not t_j , where in computing A_{dv} we take the minimum over all players i and pairs of types. In addition, let $A_{dv} = \min_{\Theta_i, \Theta_j} L_f(\Theta_i, \Theta_j)$ be the minimum such advantage. Also, let B denote the maximal difference in valuation between the default outcome O , and the outcome dictated by the social choice $SCC(\Theta_1, \dots, \Theta_n)$. We require

$$\epsilon A_{adv}/n \geq \delta \geq B \frac{1}{k}$$

Given these constraints, the proof of Glazer Perry uses a backward induction argument to prove the existence of a unique subgame perfect equilibrium. In the last round, it is better for a player i to report his true valuation regardless of anything which happened in the previous rounds, as the small probability that the result will be $f(\Theta_i)$ and not $f(t_i)$ for some $t_i \neq \Theta_i$ is enough to overcome any fine which the player will need to bear. Given that we assume that all the players report the truth in the last round, a backwards

induction shows that no player has incentive to deviate from the truth. Indeed, let i denote the last player to deviate from the truth, and suppose he does so in round j . In this case, consider the case in which the player does not deviate from the truth in that round. By telling the truth, player i can decrease the fine he pays, but also changes the output of the mechanism if round j is chosen in step 2. While decreasing the fine is always good for player i , he may either gain or lose by the new outcome. To guarantee that he doesn't lose too much by the new outcome, the GP mechanism relies on the low probability of this event, which is proportional to $1/k$. This requires k to be large.

Our Example

Finally, we can present our example. There are two possible outcomes, α, β . We have $T_1 = T_2 = T_3 = T_4$, and $T_5 = \{t_\alpha, t_\beta\}$, where

$$t_\alpha(\alpha) = 1, \quad t_\alpha(\beta) = 0$$

and similarly

$$t_\beta(\alpha) = 0, \quad t_\beta(\beta) = 1.$$

That is, each one of the first four players has a slight preference for one of the alternatives. The fifth player is similar, except that he has a strong preference, that is $T_5 = \{\hat{t}_\alpha, \hat{t}_\beta\}$, where

$$\hat{t}_\alpha(\alpha) = 2^n, \quad \hat{t}_\alpha(\beta) = 0$$

and similarly

$$\hat{t}_\beta(\alpha) = 0, \quad \hat{t}_\beta(\beta) = 2^n.$$

Note that this preference only requires $O(n)$ bits to describe; still the number of rounds will be exponential.

Finally, the required *SCC* picks the alternative which is worse for most of the players. That is, if a majority of the players have type t_α or \hat{t}_α the *SCC* requires to output β , and vice versa.

Analysis of Our Example

We begin by showing that the fine has to be less than ϵ . We do so by considering player 1. Let f_1 denote the function the mechanism uses to decide on the outcome in the $k+1$ round, given that player 1 was picked at random in the last round. Any such function must give $A_{dv} \leq 1$ (this can be obtained by $f_1(t_\alpha) = \alpha$, and $f_1(t_\beta) = \beta$, which is essentially giving player 1 his preferred outcome). However, since the proof requires that the player is truthful at the last round even if this makes him pay a fine, we must

require that the fine $\delta < \epsilon/5$, his expected utility gain from being truthful.

Consider O , the default outcome outputted by the mechanism in case that round j is chosen at random, and there are no $n-1=4$ players who declared the same vector of valuations in that round. Since the types of the players and the *SCC* function are symmetric, we can assume wlog that $O = \alpha$. Suppose that the true types of the players are $\Theta_1 = \Theta_2 = \Theta_3 = \Theta_4 = t_\alpha$, and $\Theta_5 = \hat{t}_\alpha$. Following the induction proof, we can safely assume that since the fine is small, all the players will be truthful in round $k+1$. But what will happen in round k ?

Consider the case in which the first three players declared the truth, and the fourth player lied. We want to argue that in this case the fifth player should declare the truth. If the fifth player declares the truth, and the k 'th round is chosen, the outcome of the mechanism will be β . In this case, player 5 will not be punished, and player 4 will be.

On the other hand, if he chooses to lie, and the mechanism chooses round k , the outcome will be α . However, whether round k is being picked or not, player 5 will have to pay a fine of $\delta < \epsilon/5$. The expected difference in the expected utility between these options is $\frac{2^n}{k} - \delta$. To make this negative (so that the truthful option is chosen), we require $\delta > \frac{2^n}{k}$, which gives $k > 5 \cdot 2^n / \epsilon > 2^n$, as required.