Precluding Ex Ante Collusion in Externality Mechanisms

Olga Gorelkina[*]

March, 2009

preliminary and incomplete

---

[*] Toulouse School of Economics

1 Introduction

The established approach to collusion in mechanisms, put forward by Laffont, Martimort (1997), assumes full obedience of the colluding players to the coalition's joint decisions, such as the manipulation of reports to the grand mechanism. Put differently, after the payoff-relevant information has been disclosed within the "side mechanism", infinite punishment could be applied to anyone who deviates from the mutually agreed manipulation. Participation to such contract has to be however justified ex-interim - that is, each player, upon learning his private information and knowing which actions would be appointed by the contract if he reveals, should be willing to do so.

In this paper, full enforcement of actions is infeasible. Those deviating from coalition utility maximization can, at most, be deprived of any share in its extra surplus. At the preceding stage the players decide whether to form coalitions, which is here equivalent to commitment to private information exchange. If any such commitment occurs, we say that *ex ante collusion* takes place. We assume that players can commit, at no cost, to reveal truthfully to each other: in the sense of Wilson (1978) this implies that the communication structure within each coalition is *full*. Beyond the private information that is shared within the coalition and the composition of that coalition, the players remain agnostic. In particular, they are not aware of any incidence of collusion between other players.

To exemplify ex ante collusion, suppose that each agent has a privately-known code to transform public information into the one relevant for his payoff. That is, when public information comes out, each agent privately learns his type according to his transformation code. Take as example of such code an asset-valuation formula that transforms publicly observed market data into the firm's private value for the asset. In this case collusion between firms implies, in the first place, revealing these privately-known codes to each other - so that at the moment of public information release the coalition members' types are observed simultaneously.

We say that a mechanism is susceptible to collusion, if there are benefits from sharing information within a group and using it to manipulate the reports. Formally, we compare the ex ante payoffs under different partitions of the player set into coalitions and study whether the finest partition (into singletons) provides better payoffs than those arising under a coalition's deviation. A mechanism is said to be robust to collusion, if the finest partition is stable. Note that partition of players into finest coalitions is equivalent to fair play in a mechanism - when private information remains in the possession of just one player. We find that the standard externality mechanisms like Vickrey auction and the expected externality mechanism are non-robust to collusion when coalitions can share surplus via transfers.

We modify standard externality mechanisms so as to make them robust to ex ante collusion . The new mechanisms work in environments where any player partitions can emerge without being observed by the mechanism designer. This natural possibility is not considered in the standard collusion literature, but represents an significant challenge to the design of collusion-proof mechanisms.

2 Example: Vickrey Auction

A well-known example of externality mechanisms is the Vickrey (or sealed-bid second-price) auction. The most remarkable feature of this auction consists in that bidders have a dominant strategy, namely bidding their true valuations. This nice property however breaks down as soon as collusion by any subgroup of bidders is allowed for; besides, the auctioneer's revenue may be reduced when coalitions form.[1] Suppose that bidders within a bidding ring submit only the highest valuation truthfully - and substantially under-report the lower valuations. This way the coalition reduces the expected second price

---

[1] As is well known from Green, Laffont (1979), there exists no balanced mechanism that implements an efficient rule in dominant strategies. The Vickrey auction described here is a case in point of this impossibility.

to be paid by its highest-valuation member at no cost to other coalition members. Whereas in no-transfers environment it is only their weakly-dominant strategy to withdraw from competition - the collusion has more bite if the arising surpluses can be shared. Observe that positive extra surplus is generated if, *ex post,* the second-highest valuation within coalition is higher than the second bid realized in the auction. The probability of this event is almost surely positive *ex interim* and positive *ex ante* - and therefore, there exist positive gains to collusion and losses in terms of seller's revenue.

	Extra surplus
Ex Post	$max[\theta_{(2\|C)} - a_{(1\|N/C)}; 0]$
Interim	$Prob\left(\theta_{(2\|C)} > \tilde{\theta}_{(1\|N/C)}\right) E\left[\theta_{(2\|C)} - \tilde{\theta}_{(1\|N/C)} \| \theta_{(2\|C)} > \tilde{\theta}_{(1\|N/C)}\right]$
Ex Ante	$Prob\,(id(1), id(2) \in C)\, E(\tilde{\theta}_{(2)} - \tilde{\theta}_{(\|C\|+1)})$

In order to prevent collusion ex ante, the seller seeks to induce a profitable deviation from the interim collusive equilibrium. We construct a mechanism, in which an informed insider can deviate and make sure to get no less than his agreed share in the coalition's surplus.[2]

Consider surplus sharing based on each member's contribution ex post. Contribution in this context is the price reduction resulting from bid withdrawal. Clearly, the maximal contribution is equal to the coalition's full extra surplus, and it can be pivoted by the second-highest valuation bidder if he enters the coalition last. It seems reasonable to assume that any bidder's reward in a collusive ring does not exceed this maximal contribution - otherwise, coalition as a whole would have been necessarily better off without him. [3] The direct implication of this assumption is that collusion can be

---

[2] In an environment where the bidders could present hard evidence of collusion, the problem is solved by rewarding a whistle blower at the expense of other members. More intricate mechanisms are needed if the hard evidence is not at hand.

[3] If, for example, the coalition divides its surplus assigning the Shapley values, this assumption is satisfied with a slack - since all possible contributions are weighed by the probability respective orderings arise.

broken down by allowing a coalition member to grab the extra surplus by deviating.

In a framework, where the object cannot be reallocated after sale, consider first the following modification of the second-price auction. A good is sold to the highest-valuation bidder at the second price as soon as the difference between the first and the second bid is higher than the difference between the second and the third bids. In the latter case the object is allocated to the second-highest bidder, who will pay the third bid to the auctioneer. Put differently, the object is allocated to that of the two highest-valuation bidders, who enjoys the larger "revealed surplus"[4] from winning. In this (non-efficient) auction the second-highest valuation member of a coalition ($id(2|C)$), knowing the coalition leader's valuation, can grab the surplus every time it arises. Indeed, by placing his bid at the average of $\theta_{1|C}$ and his true valuation $\theta_{2|C}$ he gets $\max\{\theta_{2|C} - \theta_{1|N/C}; 0\}$, which is precisely the ex-post extra surplus of the coalition.

This auction has the following appealing properties:

- Under fair play (when each bidder stays privately informed) bidding the true valuation remains dominant strategy. Whenever $bid_i > v_i$ there is a positive probability of getting the object at a price higher than v_i, which yields negative expected gain. On the other side, whenever $bid_i < v_i$, the gain conditional on winning the object is positive, so it is worthwhile bidding up in order to increase the probability of winning. The only stationary point is at bid equal to the valuation.

- The rules of the game, as well as the best response of the coalition's number two do not depend on the *a priori* value distribution and thus on players' beliefs.

- The auction provokes a conflict of interest between the first two highest-valuation bidders within the coalition. When the second deviates from

---

[4] difference between his announced valuation and the price he pays

> the collusive equilibrium the first is either deprived of the object he would have won otherwise or pays a price higher than he would have paid in case the second reported truthfully.

The obvious drawback of this mechanism is its inefficiency - with a positive probability the object is allocated to a bidder who does not appraise it most. Nevertheless it can be shown that the ex ante probability of such undesirable event can be driven arbitrarily close to zero, that is, we can achieve ε-efficiency. To do this, assign the object to the second-highest bidder only in case his revealed surplus is at least m times higher than the first bidder's revealed surplus. $(a_{2|N} - a_{3|N} > m * (a_{1|N} - a_{2|N}))$. Increasing m translates directly into the reduced odds of an inefficient allocation.

Provided the ε-efficiency of our modified second-price auction, the decrease in the seller's expected revenue due to collusion prevention can be made arbitrarily small, too. This result follows directly from the revenue equivalence theorem. The ε-efficient version maintains the three desirable properties listed above.

Last thing we need to check is whether the required efficiency level affects the collusion-proofness of our mechanism. The answer is, it does not: the closest follower's expected deviation payoff, conditional on $\theta_{1|C}$ and $\theta_{2|C}$, is invariant in m. By setting his bid at $\frac{m}{m+1}\theta_{1|C} + \frac{1}{m+1}\theta_{2|C}$ he gets the positive extra surplus with $Pr(\theta_{1|N/C} < \theta_{2|C})$ and zero with complementary probability - equivalent to the simplified version.

Next section provides a formalization of collusion proofness.

3 Game-Theoretical Formalization

The discourse of this section follows the two-stage nature of our game. At the first (*ex ante*) stage, when no information has yet been revealed, the players agree to collude. Decisions to collude are taken with regard to the course of play and payoffs arising at the second (*interim*) stage. At that stage, having observed the information, which is common within coalitions, the players take *individual* decisions on their actions in the grand mechanism.

3.1 Ex Ante Stage: Forming a Partition

As we have seen, in the original Vickrey auction the bidders are better off ex-ante committing to share their information. Let us think about such commitments by groups of players as partitioning - for example, partition $\{\{1,2\},\{3\}\}$ of player set $\{1,2,3\}$ would imply that players 1 and 2 commit to have common information, whereas player 3 stays apart with his information bit. (Note that we consider only bilateral information exchange - so that not providing information rules out learning from others). With this formalization, we say that collusion occurs whenever the partition of player set $N=\{1,2,..,n\}$ is different from the finest partition $\pi_N^f=\{\{1\},\{2\},..,\{n\}\}$.

Collusion is precluded in mechanism M, if it generates ex ante payoffs such that no partition of the player set dominates the finest partition. In other words, the finest partition must be stable - or belong to the core, defined relative to a certain concept of dominance. Following Koczy (2007), we will say that a partition is dominated via a coalition, if it can separate from the rest whereby making its members' better off. The decision to separate will be affected by the coalition's beliefs about the outside partition. Hence, different specifications of the coalition's beliefs give rise to different definitions of dominance. Two such definitions are given below.

Definition (Koczy, 2008). Partition π_N is *dominated via the coalition* C *forming partition* π_C if for *all (some)* $\pi_{N/C}$ we have $v_i(\pi_C\cup\pi_{N/C})\geq v_i(\pi_N)$ for all $i\in C$ with at least one strict inequality. Partition π_N is *dominated* if it is dominated via a coalition.

The first version of the definition that involves *all* residual partitions $\pi_{N/C}$ corresponds to "pessimist" coalitions - who refer to the least favorable outside partition. The second version that involves *some* residual partitions $\pi_{N/C}$

corresponds to "optimist" coalitions, that count for payoffs arising in the most favorable case. In designing mechanisms, it is clearly the "optimist" environment that collusion is hardest to preclude in: If separation is not profitable in the best-case scenario, it is never profitable. This further implies that $Core^{opt}(M) \subseteq Core^{pess}(M)$, where $Core(M)$ is the set of undominated partitions in mechanism M, defined respectively.

Going back to the Vickrey example and setting $v(\cdot)$ to represent the ex ante payoffs attainable in the eventual collusive equilibria, we observe that the finest partition is dominated via *any* coalition consisting of 2 or more players. The payoffs of the members of coalition C do not depend on the residual partition $\pi_{N/C}$ - in particular, the payoffs do not depend on whether other coalitions have been formed. This is true since any coalition's highest valuation will appear as bid, and obviously $\max_{S \in \pi_{N/C}} \max_{i \in S}(\theta_i) = \max_{i \in N/C}(\theta_i)$ for any $\pi_{N/C}$. Only this maximal value appears relevant for coalition C's payoff: it is either the first bid - then the leader looses and gets 0, or the second bid - in which case it becomes the price he pays for the object.

Due to the irrelevance of the residual partition, in the case of Vickrey auction π_N^f is dominated in the sense of both definitions given above. Whereas in case of the modified auction we proposed, π_N^f remains undominated and lies in the core - again, in the sense of both definitions. Indeed, in the new auction no coalition can increase its members' ex ante payoffs, with at least one *strict* increase.

Observe that in this paper dominance and core are defined for *partitions*, and not, as usual, *allocations*. Recall that in the latter case we always depart from the grand coalition and study whether a sub-coalition can improve upon the allocation. In our case however, we look at every pair of partitions to establish a dominance relation, if possible. Comparing partitions calls for a proper specification of the payoff function $v : \Pi(N) \rightarrow R^n$, where $\Pi(N)$ is the set of all possible partitions of player set N. We take up this point next.

3.2 Interim Stage: Defining Payoffs in the Partition-Form Game

While discussing the Vickrey example we referred to the ex ante payoffs $v : \Pi(N) \to R^n$ but did not specify their exact origination. In our multi-stage environment these payoffs depend on the outcomes of the interim stage. Provided dominant strategies in the Vickrey auction, the outcome prediction was straightforward. Indeed, playing a dominant strategy whenever available is a simple implication of rationality.

However in general partition-form games with imperfect information predicting outcomes is less evident. First of all, there is simultaneous interaction both *between* and *within* coalitions. Second, if a player only observes the composition of his coalition, his payoff estimation would involve the beliefs about the residual partition, strategies therein, type distributions, as well as all the higher-order beliefs. Here we suggest a way to pin down strategies that are best-response for at least one belief profile, whereby putting few restrictions on the admissible belief set. As there is an obvious parallel with the standard rationalizability notion of Bernheim (1984), Pearce (1984), we call this strategy set partition-rationalizable. The definition follows; it is introduced in several steps.

3.2.1 Best Response

First, consider the best-response correspondence in a game within coalition C that follows private information sharing between its members. The best-response correspondence is presented for an environment, where coalitions form *secretly* from the rest and each of the players only observes who his partners are. We define best response within C for every partition that does not contradict any C-member's observation; call such partition a notional partition for C and denote $\pi_N(C)$. The set of admissible notional partitions is $\Pi_N(C) = \{\pi_N \in \Pi_N | \{C\} \in \pi_N\}$

Example *Take* $N = \{1,2,3,4,5\}$*,* $C = \{1,2\}$*. Within* C *it is common knowledge that* $1 \in C$*,* $2 \in C$*. Hence the set of admissible partitions* $\Pi_N(\{1,2\}) = \{\{12|3|4|5\}, \{12|34|5\}, \{12|3|45\}, \{12|35|4\}, \{12|345\}\}$*, where by e.g.* $\{12|345\}$ *we mean* $\{\{1,2\}, \{3,4,5\}\}$*.*

For all coalitions S from the notional partition $\pi_N(C)$ of coalition C, different from itself, we fix the coalition C's (common) *second*-order belief $\breve{\pi}_N(S) \in \Pi_N(S)$, such that it contains a non-trivial[5] partition of coalition C. The latter requirement reflects the belief that the coalition unites unobservably by outsiders. That is, coalition C assumes that in the mind of any other coalition S it is still split apart. Note that no coherence between outside coalitions' beliefs is required. Formally, we demand every second-order belief to belong to set $\breve{\Pi}_N(C,S)$, where:

$$\breve{\Pi}_N(C,S) = \{\breve{\pi}_N \in \Pi_N | S \in \breve{\pi}_N, \exists \breve{\pi}_C \in \Pi_C, \{\{C\}\} \neq \breve{\pi}_C \subset \breve{\pi}_N\}.$$

Denote $\breve{\Pi}^2_N(C, \pi_N(C)) = \underset{S \in \pi_N(C)/\{C\}}{\times} \breve{\Pi}_N(C,S)$, the collection of second-order beliefs.

Example (cont.) Take as notional (first-order belief) partition $\pi_N(C) = \{12|34|5\}$*. Let's denote the residual coalitions* $S_1 = \{3,4\}$*,* $S_2 = \{5\}$*. Then the possible second-order believed partitions are* $\breve{\Pi}_N(C,S_1) = \{\{1|2|34|5\}\}$ *corresponding to coalition* S_1*, and* $\breve{\Pi}_N(C,S_2) = \{\{1|2|34|5\}, \{1|2|3|4|5\}\}$ *corresponding to* S_2*. The set of admissible second-order beliefs for* C *is therefore given by the product* $\breve{\Pi}_N(C,S_1) \times \breve{\Pi}_N(C,S_2)$*.*

Finally, we fix the realization θ_C of all private information that is shared within coalition C. Now are ready to define the best-response correspon-

---

[5] Partition π_C is non-trivial iff $\pi_C \neq \{\{C\}\}$.

dance for coalition C playing game M. Denote $U_i^M(r;\theta_i)$ payoff to player i of type θ_i in game M for strategy profile $r \in \underset{i\in N}{\times} \mathcal{R}_i$, where $\mathcal{R}_i$ is player i's strategy set.[6]

$$BR_C(r_C, \theta_C, \pi_N, \check{\pi}_N) = \underset{i\in C}{\times} BR_i(r_{C/i}, \theta_C, \pi_N, \check{\pi}_N) \in \underset{i\in C}{\times} \mathcal{R}_i \ ,$$

$$BR_i(r_{C/i}, \theta_C, \pi_N, \check{\pi}_N) =$$

$$= \bigcup_{r_{N/C}(\cdot)\in NE_{\pi_N/\{C\}}(\cdot,\check{\pi}_N)} \underset{r_i\in\mathcal{R}_\mathrm{i}}{argmax} \left\{ E_{\tilde{\theta}_{N/C}} U_i^M \left(\left[r_C, r_{N/C}\left(\tilde{\theta}_{N/C}\right)\right]; \theta_i \right) \right\}$$

where $NE_{\pi_N/\{C\}}(\tilde{\theta}_{N/C}; \check{\pi}_N) = \underset{S\in\pi_N/\{C\}}{\times} NE_S(\tilde{\theta}_S; \check{\pi}_N(S))$

Strategies defined by this relation are individual best-responses within C given its information θ_C and the assumption that any outsider-coalition S believes in a partition $\check{\pi}_N(S)$ (the collection of their beliefs is referred to as $\check{\pi}_N$) and plays a strategy in the set $NE_S(\tilde{\theta}_S, \check{\pi}_N(S))$. The latter is what we call a set of strategies partition-rationalizable w.r.t. $\check{\pi}_N(S)$; the definition comes next.

3.2.2 Fixed Point and Partition Rationalizability

Our notion of partition-rationalizability is recursive.

Definition. Fix C, θ_C, $\pi_N \in \Pi_N(C)$, and $\check{\pi}_N \in \check{\Pi}_N^2(C, \pi_N(C))$. Call $NE_C(\theta_C; \pi_N) = \bigcup_{\check{\pi}_N\in\check{\Pi}_N^2(C,\pi_N(C))} \left\{ r_C \in \underset{i\in C}{\times} \mathcal{R}_i | r_C = BR(r_C, \theta_C, \pi_N, \check{\pi}_N) \right\}$ *partition-rationalizable w.r.t* π_N if $\forall S \in \pi_N/\{C\}$ $NE_S(\tilde{\theta}_S, ; \pi_N(S))$ partition-rationalizable.

Observe that requiring $\check{\pi}_N \in \check{\Pi}_N^2(C, \pi_N(C))$ guarantees that rationalizability is well-defined as soon as we pin down the initial fixed point, namely

---

[6] For notational brevity, we assume state-independence of the strategy set $\mathcal{R}_i$

$NE_i(\theta_i;\pi_N^f)$. The obvious candidate for it is Bayesian-Nash equilibrium profile. (Note that multiplicity at any stage does not prevent rationalizabitlity from being well-defined).

In words, a strategy of player i in coalition C is partition-rationalizable, if there exists a tree of higher-order beliefs such that each coalition (including C) to which the belief is applicable, plays best response, holding that belief. In the example of $C = \{1,2\}$, $\pi_N(C) = \{12|34|5\}$ that we looked at as illustration, there are two admissible belief trees.

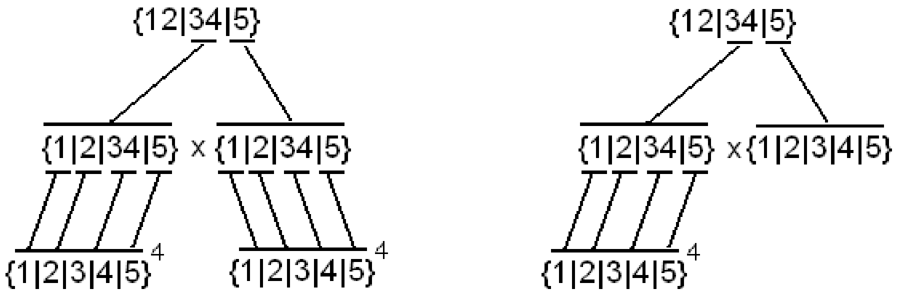

3.2.3 Defining Ex Ante Payoffs

Finally, to determine the payoff correspondence $v_i : \Pi_N \mapsto R$ for each player i, we take the expectations over all possible outcomes satisfying partition rationalizability. The image of a fixed partition π_N given by payoff correspondence v_i clearly does not have to be a singleton, or even a joint set. Formally, $v_i(\pi_N) = \left\{ y \in R |\, y = E_{\widetilde{\theta}_N} U_i^M\left(r\left(\cdot\right);\theta_i\right)\ where\, r\left(\cdot\right)\ rationalizable\, w.r.t.\, \pi_N \right\}$

3.3 Discussion and Application to Collusion in Mechanisms

The notion of partition-rationalizability, contributed in this section, can be applied to any games of imperfect information, when information is jointly held by separate groups of players. Note that this case would be virtually equivalent the standard private-information case, if the partition itself were common knowledge - however if groups form secretly, a different approach that incorporates a belief structure, is required. For example, our notion of partition-rationalizabity admits all strategies, undominated under at least one admissible belief structure.

Observe that, though we frame the strategy choices as best-responses, this framework admits, in particular, coalition surplus maximization, as soon as the eventual transfers are incorporated into the final payoffs. For example, the typical case we look at in this paper, is when surplus is shared unless at least one deviation from surplus maximization has been detected. [7]

Since this paper's primary interest is collusion-robustness, we want to know, whether for at least some partition-rationalizable strategy profiles coalitions generate extra surpluses with respect to the finest partition. This translates to the interim stage, a positive extra surplus is sustainable with transfers.

In modified Vickrey auction, an extra surplus (above truthful bidding) could arise only if the coalition's second highest valuation bidder placed a bid lower than his true valuation. But as we have seen, no transfer could support such an action as a best-response. Since there is no surplus for any realization, there is no ex ante surplus either. Hence, without computing any Nash equilibria and respective ex ante payoff, we conclude that the finest partition is not dominated. We apply similar logic to the expected externality mechanism in the following session.

---

[7] We suggest that players' actions in the mechanism - like type reports - should always be concealed by the mechanism designer, who seeks to prevent collusion. There is no reason to think this should invovle any non-negligible costs. However, as we will see, defections can be induced from outcomes that remain observable in case of standard AGV, as well as in the Vickrey auction.

4 The Expected Externality (AGV) Mechanism

Similarly to the Vickrey auction, the expected externality (AGV) mechanism in its standard formulation fails against ex ante collusion, if transfers within coalitions are feasible. To show this, we present an environment

In this section we first show an example of failure of the standard expected externality mechanism when collusion is possible, and second, introduce a modification to the budget-balance part of transfer, such that collusion is precluded.

4.1 Non-Robustness of the Standard AGV

Recall that payoff of player i in the expected externality mechanism are given by

$$v_i(x^M(r_i, r_{N/i}); \theta_i) + \sum_{j \neq i} E_{\widetilde{\theta}_{N/i}} v_j(x^M(r_i, \widetilde{\theta}_{N/i}); \widetilde{\theta}_j) + H_i(r_{N/i}),$$

where $x^M(r) = \underset{x}{argmax} \sum_{i \in N} v_j(x; r_i)$, and $H_i(r_{-i}) = -\frac{1}{n-1} \sum_{k \neq i} T_k(r_k)$, denoting $T_i(r_i) = \sum_{j \neq i} E_{\widetilde{\theta}_{N/i}} v_j(x^M(r_i, \widetilde{\theta}_{N/i}); \widetilde{\theta}_j)$. This mechanism implements the efficient allocations in Bayesian Nash equilibrium.

We look at a symmetric environment, where the players' utility function is quadratic in the amount of allocation and linear in money. This implies payoff function $v_i(x; \theta_i) = -k(x - \vartheta_i)^2 + m_i$ for an individual of type $\vartheta_i \in R$, when the allocation equals $x \in R$ and the transfer received is $m_i \in R$. Under this specification type ϑ_i corresponds to the most preferred alternative. Without loss of generality, assume $k = 1$.

Lemma Under the assumptions of symmetric quadratic preferences, risk neutrality, and zero correlation between payoffs, for any $\pi_N \neq \pi_N^f$ there exists a partition-rationalizable profile of strategies $r_i(\theta_i, \pi_N)$:

- such that $r_i(\theta_i, \pi_N) \neq \theta_i$

- $r_i(\theta_i, \pi_N)$ induces payoffs in the expected externality mechanism $v()$, such that π_N dominates π_N^f

This lemma states that the standard externality is non-robust to collusion.

4.2 Modification of the Expected Externality Mechanism

We depart from the standard externality mechanism, where, in order to balance the budget, the incentive payment $T_i(r_i)$ to player i is covered in equal parts by all other players (recall, $H_i(r_{-i}) = -\frac{1}{n-1}\sum_{k\neq i} T_k(r_k)$). We first change the compensation rule and assign to each player another one to reimburse the incentive payment, so that budget balance is preserved. Without loss of generality, we say that player $i+1$ pays the incentive transfer of player i (and 1 pays for n). Second, we change the incentive transfer to “almost dominant strategy” implementation one. Set $T_i(r_{N/(i+1)}) = E_{\widetilde{\theta}_{(i+1)}}\left[\sum_{j\neq i,(i+1)} v_j(x^M(r_{N/(i+1)}, \widetilde{\theta}_{(i+1)}); r_j) + v_{(i+1)}(x^M(r_i, \widetilde{\theta}_{-i}); \widetilde{\theta}_{(i+1)})\right]$. Then, player i’s budget-balancing transfer is $H_i(r_{N/i}) = T_{(i-1)}(r_{N/i})$.

Proposition. With transfers, defined as above, for all π_N the only partition-rationalizable strategy for any player is truth-telling: $r_i(\theta_i, \pi_N) = \theta_i$. Hence π_N^f is trivially undominated and lies in the optimistic core.

Most importantly, note that we did not assume anything about the preferences and type distributions (in particular, the type space can be finite or infinite). Even the common knowledge of type distributions is not required. It is only necessary that player i has in mind the same distribution of $\widetilde{\theta}_{(i+1)}$, as the mechanism designer. In that sense our mechanism is almost detail-free.

5 Conclusion

Collusion between bidders in a Vickrey auction destroys the truthful-revelation property and reduces the seller's expected revenue. Whenever transfers within coalitions are feasible, such collusion can be sustained. We presented an amendment to the Vickrey auction that precludes collusion within any subset of bidders, introducing a conflict of interest within symmetrically-informed coalitions. The modified auction is detail-free, achieves ε-efficiency and dominant-strategy incentive compatibility.

Similarly to the Vickrey auction, the expected externality (AGV) mechanism is shown to fail against ex ante collusion. We introduce a change in the assignment of budget-balancing transfers that precludes collusion and induces truth-telling, whereby requiring less of common knowledge.