UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Institute for New Economic Thinking

Cambridge-INET Institute

Cambridge-INET Working Paper Series No: 2016/03

Cambridge Working Paper Economics: 1619

A COASIAN APPROACH TO EFFICIENT MECHANISM DESIGN

Mikhail Safronov
(University of Cambridge)

Revised September 2016

Vickrey-Clarke-Groves and d'Apresmont Gerard-Varet mechanisms implement efficient social choice by paying each agent the externality her report imposes on all other agents. This paper provides an alternative mechanism, where each pair of agents directly compensates each other according to the externalities they impose. In the new mechanism any agent can guarantee to receive her ex ante efficient payoff by committing to report truthfully, regardless of others' strategies. This absence of ex ante externalities makes all equilibria efficient, whether agents act individually or in coalitions. Truthful report is a coalition-proof equilibrium, and under mild assumptions it is the uniquely rationalizible strategy. The mechanism is extended to repeated interactions.

A Coasian Approach to Efficient Mechanism Design

Mikhail Safronov*

University of Cambridge

August 31, 2016

Abstract

Vickrey-Clarke-Groves and d'Apresmont Gerard-Varet mechanisms implement efficient social choice by paying each agent the externality her report imposes on all other agents. This paper provides an alternative mechanism, where each pair of agents directly compensates each other according to the externalities they impose. In the new mechanism any agent can guarantee to receive her ex ante efficient payoff by committing to report truthfully, regardless of others' strategies. This absence of ex ante externalities makes all equilibria efficient, whether agents act individually or in coalitions. Truthful report is a coalition-proof equilibrium, and under mild assumptions it is the uniquely rationalizible strategy. The mechanism is extended to repeated interactions.

*E-mail: mikhailsafronov2014@u.northwestern.edu. This project was started during my PhD study at Northwestern University. I am grateful to Matt Elliott, Aytek Erdil, Robert Evans, Ben Golub, Qingmin Liu, Roger Myerson, Alessandro Pavan, Antonio Penta, Doron Ravid, Soenje Reiche, Ludovic Renou, Hamid Sabourian, Bruno Strulovici, Juuso Toikka, Asher Wolinsky and participants of the 2016 Annual Conference of the Royal Economic Society, 21st Coalition Theory Network workshop, the 2016 North American Summer Meeting of the Econometric Society, the Stony Brook 27th International Conference on Game Theory and the 5th World Congress of the Game Theory Society for fruitful discussion. I am thankful for financial support provided by the Center for Economic Theory and the Graduate School at Northwestern University; and by the Cambridge-INET Institute. All errors are mine.

1 Introduction

The Coase theorem states that the problem of externalities causing economic inefficiency can be solved if the externalities can be traded. When the externalities are traded, their effects are internalized, inducing any agent responsible for the externalities to make the efficient social choice. This paper shows that the similar idea can be used in a benevolent social planner's problem with agents having independent private types and quasilinear preferences. The social planner asks the agents to report their private types, and then she makes the social choice which maximizes the total payoff of the agents. Since any agent's report affects the social choice, the agents impose externalities on each other and may benefit from misreporting their types. In order to induce truthful reports, the agents should be made to compensate each other for these externalities.

The idea of internalizing of externalities has been already used in the classic Vickrey-Clark-Groves (VCG) and d'Aspremont-Gerard-Varet (AGV) mechanisms, though the way it is implemented in these mechanisms is not as in the Coase theorem. In VCG mechanism it is the social planner who compensates the agents for the externalities. In AGV mechanism the compensation is unfair: if agent i's report causes externalities on agent j and no externalities on agent k, agent k still has to pay agent i. As a result, both VCG and AGV mechanisms internalize the externalities partially and are not resistant to group deviation. In these mechanisms each agent individually prefers to report truthfully, but a group of agents can coordinate on a misreport and jointly benefit.

The current paper presents an alternative mechanism, where each pair of agents directly compensates each other for externalities. In the mechanism the agents are ordered in an *arbitrary* sequence and report their types sequentially according to that ordering. Each report is publicly observed, including the agents who has yet to report.

When any agent i reports own type, social planner updates her beliefs on the efficient choice she will make at the end, and she updates the expected payoffs of agents from that choice. The mechanism prescribes any other agent $j \neq i$ to pay agent i the change in j's expected payoff which occurred from i's report. These payments are made for the report of each agent.

In the new mechanism each pair of agents directly compensates each other for the externalities of their reports. As a result, all externalities are removed at ex ante level. If any agent i, before learning her type, commits to report truthfully, she guarantees to get her ex ante efficient payoff, regardless of others' strategies. This result follows from the way the payments are being made. First, agent i receives a payment from any other agent j, equal to the change of j's payoff caused by i's report. Since agent i reports truthfully, in expectation over i's report that change is zero, and so is j's payment to i. Second, agent i makes payment to j, equal to the change in i's expected payoff caused by j's report. Effectively the utility of agent i (payoff from social choice plus payments) does not change with j's report. Therefore, i's utility does not change with reports of other agents and equals to its ex ante value, to i's ex ante efficient payoff.

The idea of the mechanism is similar to the property rights of the Coase theorem. Before the mechanism the social planner expects each agent i to obtain ex ante efficient payoff, and provides agent i with the guarantee for that payoff. This guarantee is an attractive feature of the mechanism, since agent i does not need to hold any beliefs about others' strategies.

The property of no ex ante externalities guarantees the social choice to be efficient in any equilibrium, (i.e., full efficient implementation). Since reporting truthfully is always an option, in any equilibrium each agent gets at least her ex ante efficient payoff. The total utility of all the agents equals to at least total efficient payoff. Since the mechanism is ex post budget balanced, the total payoff of all the agents is efficient, and so is the social choice. The full efficient implementation holds regardless of whether agents act individually or being in coalitions. If the process of coalition formation is endogenous, any agent guarantees ex ante efficient payoff by refusing to join the coalition and reporting truthfully. If the process of coalition formation is exogenous, happening before the mechanism and making the agents in coalition to act as a single player, then truthful report guarantees efficient payoff to coalition as well as to any agent outside the coalition.

The set of equilibria in the mechanism always contains truthful equilibrium. This equilibrium is coalition-proof: it is not profitable to misreport for any coalition. Otherwise, if a coalition could strictly benefit by misreporting, it would decrease the utility of some agent outside the coalition below ex ante efficient payoff. This is impossible, however, since that other agent reports truthfully.

In truthful equilibrium the payment made to each agent i equals to the expected payoff of all other agents; that payoff is estimated given the reports before i and assuming agents after i report truthfully. The incentives to report truthfully thus lie between VCG and AGV mechanisms. The solution concept for the truthful equilibrium can also be made to lie between weak dominance of VCG and Bayesian Nash equilibrium of AGV. With a mild assumption of efficient social choice being unique, truthful report becomes a uniquely rationalizible strategy. The last agent strictly prefers to report truthfully regardless of others' reports. Knowing that, the second-to-last agent strictly prefers to report truthfully as well. By induction, all the agents have truthful report as the uniquely rationalizible strategy.

The mechanism has other features. The ordering in which the agents report their types, determines the monetary transfers to each agent from the mechanism. However, even if one fixes the ordering, it is possible to find alternative transfers achieving the main property of ex ante removal of externalities. One could also impose symmetry by making all the agents to report their types simultaneously and then choose the ordering of 'revealing' the reports uniformly. Since the mechanism works for any arbitrary deterministic ordering of agents, it works for random ordering as well. In the resulting symmetric mechanism each agent pays the externality other agents impose on her, and gets paid the Shapley value of externalities her report imposes on others. Switching to simultaneous reports, however, prevents unique rationalizibility of truthful report.

The mechanism built in this paper can be extended to other environments. First, the mechanism is shown to work in a dynamic setting with agents' types changing over time, as in paper by Athey and Segal (2013). In the dynamic setting the agents compensate each other for the changes in their expected continuation payoffs. As a result, on equilibrium path one achieves full efficient implementation and coalition-proofness. Second, under some conditions, the idea of trading externalities can be applied when taken into account interim participation constraint.

The paper is organized as follows. Section 2 discusses the relevant literature. Section 3 builds the mechanism and shows the properties of full efficient implementation and coalition-proofness. Section 4 describes other properties of the mechanism, including unique rationalizibility of truthful report. Section 5 adjusts the mechanism to a dynamic setting and to a setting with participation constraints. Section 6 concludes.

2 Literature review

The idea of internalizing the externalities in efficient mechanism has given rise to classic Vickrey-Clarke-Groves (VCG) and d'Aspremont Gerard-Varet (AGV) mechanisms. In VCG mechanism, introduced by Vickrey (1961), Clarke (1971), Groves (1973), each agent is paid the externality her report imposes on other agents. As a result, the truthful report is a weakly dominant strategy. AGV mechanism from the paper by d'Aspremont Gerard-Varet (1979), uses similar idea: each agent is paid the expected externality her report imposes on other agents. The payment is made budget-balanced by taking it with equal shares from other agents. As a result, AGV mechanism is budget-balanced, though the solution concept is weaker: truthful report is Bayesian incentive-compatible.

The mechanisms in Samuelson (1985) and Cramton, Gibbons, Klemperer (1987) perform similar to the Coase theorem. They consider the environments where the agents have property rights of an asset and trade them through efficient mechanisms. Each agent owning share of the asset imposes individual rationality constraints and makes it impossible to always have efficient allocation. The authors find the conditions on the initial shares under which one gets efficiency. In comparison, my paper builds the mechanism in which each agent is expected to get her ex ante efficient payoff from the social choice, similar to initial share of the asset. By submitting their private types, agents change their efficient payoffs, and compensate each other for those changes. Since there is no participation constraint, the mechanism always achieves efficiency.

The idea of internalizing the externalities has been considered in dynamic environments in the papers by Bergemann and Välimäki (2010) and Athey and Segal (2013). Athey and Segal extend AGV mechanism to a dynamic setting: agents' private types evolve stochastically over time. In order to incentivize the agents to report truthfully, each agent i is paid the change of efficient continuation payoff of other agents, caused by i's report. The mechanism is budget-balanced by taking that payment with equal shares from other agents. As a result, truthful report is Bayesian Nash equilibrium. The same idea of compensating the agents for the *change* in their payoff is used in my paper. The difference is that, unlike AGV mechanism, each pair of agents directly compensates each other.

Another series of papers studies the problem of collusion in mechanism design. Laffont, Martimort (1997, 1998, 2000) consider the environment with two agents and show the optimal outcome to be collusion-proof in case of independent types. The paper by Che and Kim (2006) extends the model to an arbitrary number of agents and more general environment with object allocation. Making the principal to 'sell' the object to the grand coalition, Che and Kim show that any incentive compatible, individually rational mechanism can be adjusted to be collusion-proof in case the grand coalition is formed. With an additional requirement of ex-post incentive compatibility, the same result holds if a subgroup of agents can form a coalition and the principal knows at least two agents in the subgroup. In another paper on auctions Che and Kim (2009) show that with passive beliefs, assuming impossibility of forming the grand coalition, the seller can achieve the same revenue as in case of no collusion.

The assumption of passive beliefs, used widely in the models of collusion, was motivated in Myerson (2007). The agents report their types to the social planner, though they are not yet committed to them. The third party proposes a collusion, and if successful, the involved agents resubmit their reports. Otherwise, if the collusion fails, the reports are unchanged.

The problem of different aspects of the mechanism with collusion has been studied more extensively in auctions. McAfee and McMillan (1992) show that the inability of the cartel members to pay each other cuts down their payoffs. Later, Che, Condorelli and Kim (2013) show that in this case the seller is not hurt by the collusion possiblity. Erdil and Klemperer (2011) propose a new class of payment rules to make the agents less willing to submit non-truthful bids if colluding. Biran and Forges (2011) consider the stability of a collusion in auctions with respect to externalities each bidder may impose on others if getting the object. Chen, Micali (2012) allow the agents to report not only their value but also the coalition they belong to. If several agents consistently report being in a same coalition, and one of them wins the good, the bids of other coalition members do not increase the payment, inducing the agents to reveal being in a coalition.

An independent branch of literature is devoted to *full implementation*: it considers mechanism design in which all equilibria achieve the desired social choice. In the environment with observable types one requires Maskin monotonicity condition (described in Maskin (1998)). This condition is extended to Bayesian monotonicity in the environments with incomplete information and interdependent types, as shown in Jackson (1991). The idea is that for any undesirable outcome, there is an agent who can credibly inform the designer if this outcome is being played and get rewarded. However, one needs a non-direct mechanism for this communication to be possible. Matsushima (1993) shows that with quasilinear utilities and side payments one can replace Bayesian monotonicity with much weaker condition, which is satisfied for a generic class of social choices. This result is further developed by Chen, Kunimoto, Sun (2015) where one needs only small transfers for full implementation. A recent paper by Ollar, Penta (2015) shows the full implementation using a direct mechanism. In their paper mechanism designer uses moment conditions, commonly known to both the designer and the agents, and make truthful report the uniquely rationalizible strategy.

3 Mechanism

I consider a setup with n agents. Each agent $i \in \{1, ..., n\}$ privately observes own private type θ_i. The total type profile over all agents is denoted as θ. Types are independently distributed across the agents, the ex ante distribution of types is publicly known. There is a set of social choices S, each choice $s \in S$ gives agent i a payoff $u_i(\theta_i, s)$. I allow for monetary transfers and assume the agents to have quasilinear utilities: if agent i receives amount x_i of money, her total utility equals $u_i(\theta_i, s) + x_i$. Later in the paper I will refer to the 'payoff' as the payoff $u_i(\theta_i, s)$ from social choice, and the 'utility' as the payoff plus monetary transfers.

I assume that for any type profile θ there exists an *efficient* social choice $s^*(\theta) \equiv argmax_s \sum_i u_i(\theta_i, s)$, which maximizes the sum of agents' payoffs[1], given θ. The payoff of agent i at efficient choice $s^*(\theta)$ is denoted as $u_i(\theta)$. I assume there exists a social planner, who implements an efficient mechanism. In the mechanism each agent i reports her private type $\hat{\theta}_i$, and then the social planner chooses $s^*(\hat{\theta})$ as a function of total report profile $\hat{\theta}$. Since each report affects the social choice, agents impose externalities on each other. I will now introduce transfers $x_i(\hat{\theta})$ into the efficient mechanism, which will remove the externalities at ex ante level; achieving full efficient implementation, coalition-proofness and ex post budget balance.

The idea of the mechanism is to make the agents to compensate each other for the externalities they impose when reporting. The agents are *arbitrarily* ordered into a sequence $1, 2, 3, ..., n$, which is publicly known before the mechanism. The agents report their types according to that sequence, each report being publicly observed [2]. Since the social planner knows the type distribution for each agent, she knows the distribution over efficient social choices, assuming that all agents report truthfully. The social planner can calculate ex ante expected payoff $E_\theta u_i(\theta)$ for each agent i, before the mechanism starts. When agent 1 submits report $\hat{\theta}_1$, this new information allows the social planner to update the expected payoffs of all agents. When agent 2 submits report $\hat{\theta}_2$, social planner makes another update, *given* the already known report $\hat{\theta}_1$ of agent 1. In general, when any agent reports own type, the expected payoffs of all the agents get updated.

After agent i has submitted report $\hat{\theta}_i$, the mechanism prescribes other agents to make the following monetary transfers to i [3]. Any other agent $j \neq i$ pays i the change in

[1] In case of several choices maximizing the total payoff, one is arbitrarily chosen as $s^*(\theta)$.

[2] The results will hold if some of the agents do not observe some of the previous reports. In particular, the agents may all report simultaneously, and an ordering being put afterwards.

[3] These transfers can be made immediately after agent i's report, or at the end, when all the agents have submitted their reports.

j's expected payoff caused by i's report:

DEFINITION 1 *Given the total submitted report to be $\hat{\theta}$, agent j pays agent i the change in expectation of j's payoff:*

$$E_{\theta_{i+1},...,\theta_n} u_j(\hat{\theta}_1, \hat{\theta}_2, ..., \hat{\theta_{i-1}}, \hat{\theta}_i, \theta_{i+1}, ..., \theta_n) - E_{\theta_i,\theta_{i+1},...,\theta_n} u_j(\hat{\theta}_1, \hat{\theta}_2, ..., \hat{\theta_{i-1}}, \theta_i, \theta_{i+1}, ..., \theta_n)$$

caused by report $\hat{\theta}_i$.

This payment from j to i may be negative (i.e., agent j receives a positive transfer from agent i), if the change in expectation of payoff u_j is negative.

Such payments are made for the report of each agent, making each pair of agents to exchange monetary transfers according to the externalities they impose on each other. Let's show that such monetary transfers lead to:

THEOREM 1 *If any agent i commits to report truthfully before observing her type, her expected utility equals to her ex ante efficient payoff $E_\theta u_i(\theta)$, regardless of others' strategies.*

Proof.

The proof will be shown for the concrete example of $n = 5$ agents to illustrate the general idea. The agents are numbered in order of submitting their reports, as $1, 2, 3, 4, 5$. Let agent 3 commit to report truthfully. Agent 3's transfers in the mechanism consist of two parts: a) others pay agent 3 since 3's report affect their payoffs, and b) agent 3 pays others since their reports affect 3's payoff.

First observation is that since agent 3 reports truthfully, in expectation over 3's report, any other agent pays zero to 3. For example, agent 1 pays agent 3 the value

$$E_{\theta_4,\theta_5} u_1(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \theta_4, \theta_5) - E_{\theta_3,\theta_4,\theta_5} u_1(\hat{\theta}_1, \hat{\theta}_2, \theta_3, \theta_4, \theta_5)$$

This value does not depend on the reports $\hat{\theta}_4, \hat{\theta}_5$ made after 3's report. Moreover, the reports $\hat{\theta}_1, \hat{\theta}_2$ cannot be conditioned on 3's report, $\hat{\theta}_3$. If one substitutes 3's report $\hat{\theta}_3$ for a true value θ_3, and takes expectation over 3's type, the value of 1's payment becomes zero:

$$E_{\theta_3}[E_{\theta_4,\theta_5} u_1(\hat{\theta}_1, \hat{\theta}_2, \theta_3, \theta_4, \theta_5) - E_{\theta_3,\theta_4,\theta_5} u_1(\hat{\theta}_1, \hat{\theta}_2, \theta_3, \theta_4, \theta_5)] = 0$$

Indeed, agent 1 pays agent 3 the change in 1's expected payoff, where the expectation is calculated given 3's true type distribution. Since agent 3 reports truthfully, the average of 1's expected payoff does not change, giving zero expected transfers. Similarly, agents $2, 4, 5$ expect each to pay zero to agent 3, too.

Let's now find transfers which agent 3 pays others since their reports affect 3's payoff. Agent 3 pays agent 1 the value

$$E_{\theta_2,\theta_3,\theta_4,\theta_5} u_3(\hat{\theta}_1, \theta_2, \theta_3, \theta_4, \theta_5) - E_\theta u_3(\theta)$$

and pays agent 2 the value

$$E_{\theta_3,\theta_4,\theta_5} u_3(\hat{\theta}_1, \hat{\theta}_2, \theta_3, \theta_4, \theta_5) - E_{\theta_2,\theta_3,\theta_4,\theta_5} u_3(\hat{\theta}_1, \theta_2, \theta_3, \theta_4, \theta_5)$$

In total, agent 3 pays both agents $1, 2$ the externality of their reports on 3's payoff:

$$E_{\theta_3,\theta_4,\theta_5} u_3(\hat{\theta}_1, \hat{\theta}_2, \theta_3, \theta_4, \theta_5) - E_\theta u_3(\theta) \tag{1}$$

Similarly, agent 3 pays both agents $4, 5$ the value

$$u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5) - E_{\theta_4,\theta_5} u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \theta_4, \theta_5) \tag{2}$$

Thus, agent 3 pays other agents the value

$$E_{\theta_3,\theta_4,\theta_5} u_3(\hat{\theta}_1, \hat{\theta}_2, \theta_3, \theta_4, \theta_5) - E_\theta u_3(\theta) + u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5) - E_{\theta_4,\theta_5} u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \theta_4, \theta_5) \tag{3}$$

When taking expectation over 3's truthful report $\hat{\theta}_3 = \theta_3$, the first and the last terms in expression 3 disappear, leaving the value $u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5) - E_\theta u_3(\theta)$ as 3's transfer to others. Since agent 3 reports truthfully, the payoff of agent 3 from a social choice equals to $u_3(\hat{\theta}_1, \hat{\theta}_2, \theta_3, \hat{\theta}_4, \hat{\theta}_5) = u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5)$. Therefore, 3's utility equals 3's payoff $u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5)$ minus 3's transfer $u_3(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5) - E_\theta u_3(\theta)$, and equals to 3's ex ante efficient payoff, $E_\theta u_3(\theta)$.

In general case of n agents the proof is the same. If agent i commits to report truthfully, every other agent $j \neq i$ expects to pay zero to i. Moreover, when any other agent $j \neq i$ submits report $\hat{\theta}_j$, the expected payoff of agent i changes, and agent i pays that change to agent j. Therefore, the expected utility of agent i does not change. After all reports are submitted, the utility of agent i remains the same as it was before the mechanism, and equals to ex ante value $E_\theta u_i(\theta)$. *Q.E.D.*

3.1 Full efficient implementation and coalition-proofness

Theorem 1 shows that all the externalities among the agents are removed at ex ante level. This property leads to all equilibria achieving efficient social choice with probability 1, regardless of the agents forming coalitions:

Theorem 2 *In any equilibrium of the mechanism, the efficient social choice is achieved with probability* 1*. This result does not depend on whether agents act individually or being in coalitions.*

Proof.

Let's first consider the agents acting individually. Each agent i has always an option to report truthfully. Therefore, in any equilibrium, from ex ante point of view, any agent i gets at least the value $E_\theta u_i(\theta)$ as her utility. The total utility of all the agents equals at least $\sum_i E_\theta u_i(\theta)$. Since the mechanism is budget-balanced, from ex ante point of view the total *payoff* of the agents equals to at least ex ante efficient value of $\sum_i E_\theta u_i(\theta)$. This can happen only if the efficient social choice is achieved with probability 1.

Now let the agents form coalitions, which can happen in two ways. With an *endogenous* process of coalition formation, the coalition is being formed after the mechanism was announced. There is a third party which defines rules of the coalition and proposes a group of agents to join the coalition. If all agents in the group agree, they act in the original mechanism according to the rules of coalition. However, any agent can refuse joining the coalition and report independently in the original mechanism. By refusing to join the coalition and reporting truthfully, any agent can guarantee her ex ante efficient payoff. Full efficient implementation result follows.

With an *exogenous* process of coalition formation the coalition is formed before the mechanism. Agents in the coalition act as a single player: their private types are a common knowledge within the coalition, and they can make monetary transfers to each other. Coalition members jointly report their types in the mechanism to maximize their total utility. With this concept of coalition formation, any agent outside of the coalition guarantees her ex ante efficient payoff, and so does the coalition by making all its members to report truthfully. The total payoff is thus ex ante efficient, showing full implementation.

Q.E.D.

The property of full implementation does not guarantee by itself the existence of equilibrium. The next theorem shows that truthful report is always an equilibrium:

Theorem 3 *Truthful report is an equilibrium, whether agents act individually or being in coalitions.*

Proof.

The theorem will be shown to hold under the extreme case of several agents exogenously forming a coalition C and behaving as a single player. First let's notice that if everyone reports truthfully, everyone receives ex ante efficient payoff as utility. Let the agents outside C report truthfully, and let's assume that there exists a type θ'_C of the coalition, at which the coalition could misreport and strictly benefit. Since the mechanism is efficient and budget-balanced, coalition C misreporting and benefitting would lead to some agent $i \notin C$ outside the coalition to strictly lose her utility. If probability of type θ'_C was positive, by reporting truthfully at all types except θ'_C, and misreporting at θ'_C, coalition C can make i's utility to fall below ex ante efficient payoff $E_\theta u_i(\theta)$. By Theorem 1, that is impossible, since i always reports truth. Thus, truthful reporting is an equilibrium, any agent (coalition) reports truthfully with probability 1.

Q.E.D.

4 Properties of the mechanism

4.1 Truthful report as a uniquely rationalizible strategy

In truthful equilibrium the incentives to report truthfully lie between VCG and AGV mechanisms. Each agent i prefers to report truthfully no matter of previous reports (similar to VCG) and anticipating future reports to be truthful (similar to AGV). The solution concept for truthful report to be an equilibrium is weaker than in weakly dominant strategies of VCG; although truthful report is a uniquely rationalizible strategy under the following assumption:

ASSUMPTION 1 *For any two type profiles $\theta \neq \theta'$ one has*

$$\sum_i u_i(\theta) = \sum_i u_i(\theta_i, s^*(\theta)) > \sum_i u_i(\theta_i, s^*(\theta'))$$

Assumption 1 imposes the requirement that all reports have to be truthful for the total payoff to be efficient.

PROPOSITION 1 *Assumption 1 makes the strategy to report truthfully to be uniquely rationalizible choice for each type of each agent.*

Proof.

The proof will be provided for the case of five agents, ordered as they submit their reports as $1, 2, 3, 4, 5$. Agent 5's report, $\hat{\theta}_5$, makes the transfers paid to 5 to be equal to:

$$\sum_{i \neq 5} [u_i(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5) - E_{\theta_5} u_i(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \theta_5)]$$

The second term does not depend on $\hat{\theta}_5$. Essentially agent 5 gets paid the value $\sum_{i \neq 5} u_i(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4, \hat{\theta}_5)$, which is the payoff of other agents from efficient social choice. This makes 5's utility to be equal to the total payoff, which is uniquely maximized at truthful report $\hat{\theta}_5 = \theta_5$.

Now let's look at incentives of agent 4. Since agent 5 reports truthfully, agent 4 expects to pay zero to agent 5, regardless of 4's report. Thus, 4's report only affects transfers made to 4. These transfers equal to:

$$\sum_{i\neq 4}[E_{\theta_5}u_i(\hat{\theta}_1,\hat{\theta}_2,\hat{\theta}_3,\hat{\theta}_4,\theta_5)-E_{\theta_4,\theta_5}u_i(\hat{\theta}_1,\hat{\theta}_2,\hat{\theta}_3,\theta_4,\theta_5)]$$

Agent 4's report affects only the first term, which is the expected payoff of other agents; the expectation is taken over 5's truthful report. Agent 4's utility equals to the expected payoff of all agents. Since any misreport of agent 4 leads to a positive loss in the expected total payoff, agent 4 has a unique best strategy to report truthfully.

By induction, agents $1, 2, 3$ can be shown each to have truthful report as the only rationalizible strategy as well. *Q.E.D.*

4.2 Symmetric mechanism and Shapley value

The mechanism achieves coalition-proofness and full efficient implementation for any ordering, in which agents submit their types. The sequential reporting makes the mechanism non-symmetric for the agents. One could impose symmetry by making the agents to report their types simultaneously, and then to uniformly randomize over the orderings, in which agents' reports are 'revealed'. In the resulting mechanism, the transfer to agent i equals:

$$x_i(\hat{\theta})=E_\theta u_i(\theta)-u_i(\hat{\theta})+$$

$$+\sum_{m=0}^{m=n-1}[\frac{1}{nC_{n-1}^m}\sum_{j_1,j_2,\dots,j_m\neq i}(E_{\theta_{j_1},\dots,\theta_{j_m}}a(\{\hat{\theta}_k\}_{k\neq i,j_1,j_2,\dots,j_m},\hat{\theta}_i,\theta_{j_1},\dots,\theta_{j_m})- \qquad (4)$$
$$-E_{\theta_i,\theta_{j_1},\dots,\theta_{j_m}}a(\{\hat{\theta}_k\}_{k\neq i,j_1,j_2,\dots,j_m},\theta_i,\theta_{j_1},\dots,\theta_{j_m}))]$$

where $a(\theta)=\sum_i u_i(\theta)$ is the efficient total payoff. One has then:

COROLLARY 1 *Transfers given by* (4) *achieve full efficient implementation, coalition-proofness and budget balance, with all agents treated in a symmetric way.*

One way to think of those transfers is as follows. For any agent i the first line in expression (4) is how much she should pay other agents as their reports affect her payoff. Two bottom lines in expression (4) indicate how much agent i should be paid as her report affects the payoffs of others. One takes any subset $J = \{i, j_1, ..., j_m\}$ of agents including i and calculates the externality of agent i's report given the reports of agents outside J. In other words, one calculates externality of i's report, as if agents outside J reported before i, and agents $j_1, ..., j_m$ reported after i. One then takes the average across all possible J-s, which is a Shapley value of externalities of i's report, and pays it to agent i.

It should be noted that with simultaneous reports the truthful report is no longer a uniquely rationalizible strategy. This leads to the choice of either agents reporting sequentially and having the truthful report as a uniquely rationalizible strategy, or agents reporting simultaneously and the mechanism being symmetric.

4.3 Uniqueness

The transfers in the mechanism depend on the ordering in which the agents report their types. However, even if one fixes the ordering, the transfers achieving ex ante removal of externalities (Theorem 1) are not unique. When agent i submits her report, the current mechanism prescribes any other agent j to pay i the change in expectation of j's payoff, u_j. Instead, one could prescribe agent j to pay the change in the expectation of the expression $E_{\theta_j} u_j$. That is, one takes the original transfers given by Definition 1, and applies the expectation operator $E_{\theta_j}(.)$ If agent j reports after i, the new transfers coincide with the original ones given by Definition 1. If agent j reports before i, the new transfers are calculated as if the social planner has forgotten j's report.

With the new transfers one still achieves ex ante removal of externalities. The proof is the same as in Theorem 1:

Proposition 2 *Theorem 1 holds, if agent j pays to i the change in expectation of $E_{\theta_j} u_j$.*

Proposition 2 implies that if one imposes symmetry on the mechanism by making the agents to submit their types simultaneously, the transfers required for ex ante removal of externalities are not unique.

4.4 Impossibility of coalitional weak dominance

The mechanism constructed in the paper can be thought of as a mixture between AGV and VCG mechanisms, for both agents and coalitions. Since for agents one has VCG achieving weak dominance, one may wonder whether the similar result holds for coalitions (dropping budget balance property). That is, after all the agents have submitted their reports, even if several agents are in a coalition and behave as a single player, they do not with to change their reports. The following statement shows that the required condition demands full separability for total efficient payoff[4]:

Proposition 3 *In the efficient mechanism a truthful report is a weakly dominant strategy for all exogenously formed coalitions, if and only if there exists a set of functions $f_i(\theta_i)$ such that*

$$\sum_i u_i(\theta) = \sum_i f_i(\theta_i) \tag{5}$$

Proof.

[4]Crémer (1996) shows the general impossibility of weak dominance for coalitions, even under weaker version of coalition formation.

If one wants the coalition C_{-i}, which contains everyone but agent i, to have a weakly dominant strategy to report truthfully, the total transfer to C_{-i} has to satisfy[5]:

$$\sum_{j \neq i} x_j(\hat{\theta}) = u_i(\hat{\theta}) - f_i(\hat{\theta}_i) \tag{6}$$

That is, the coalition C_{-i} gets as a total transfer the payoff of agent i minus some term $f_i(\hat{\theta}_i)$, which does not depend on the report $\hat{\theta_{-i}}$ of coalition C_{-i}. The equation (6) has to hold for all i.

In order to incentivize the grand coalition to report truthfully, the total transfer $\sum_i x_i(\hat{\theta})$ has to be constant, independent of $\hat{\theta}$. This means, that if one sums up the equations (6) for all i, then on the left hand side there will be a constant. On the right hand side one will have the total efficient payoff $\sum_i u_i(\hat{\theta})$ and sum of terms $-f_i(\hat{\theta}_i)$, each depending only on the type of agent i, yielding the result. *Q.E.D.*

Remark 1 *In the proof the incentives for the grand coalition were checked, and one may find the assumption for the grand coalition (and any big coalition) to be formed as too strong. However, even if one wants truthful report to be a weakly dominant strategy for two agents i, j and a coalition $\{i, j\}$, it has to be that the total efficient payoff has to be separable across types θ_i, θ_j, which is a restrictive condition.*

Proposition 3 shows that type independence is essential for the mechanism and for Theorem 1. The property of ex ante removal of externalities is stronger than the coalition-proofness. The coalition-proofness cannot be obtained in weakly dominant strategies, and respectively it cannot be obtained either in case of very strong type correlation among the agents. Indeed, in truthful equilibrium any coalition could predict others' reports with a very high probability, and the Bayesian equilibrium would be almost equivalent to the equilibrium in weakly dominant strategies. Therefore,

[5]These conditions are the same as in VCG mechanism, written for a coalition.

type correlation may prevent coalition-proofness and ex ante removal of externalities. The precise bounds in which the current mechanism works are yet to be found.

5 Extensions

5.1 Individual rationality

The current mechanism, built similar to VCG and AGV mechanisms, does not impose individual rationality constraint: no agent can quit the mechanism. Since the mechanism satisfies efficiency, incentive compatibility and budget balance, due to Myerson-Satterwaite impossibility theorem individual rationality cannot be added for general case. However, a different question may be asked. Let there be an efficient mechanism which already satisfies interim incentive compatibility and individual rationality, and ex post budget balance. Is it possible then to add full efficient implementation and coalition-proofness?

The answer is positive under the some conditions, as shown below.

Proposition 4 *There exists a budget-balanced mechanism, which satisfies Theorem 1 and individual rationality, if either*

a) there exists an efficient mechanism M, which satisfies interim incentive compatibility and individual rationality, and ex post budget balance. The type profile θ and social choice s are elements of a compact set in Euclidean space. Efficient social choice $s^*(\theta)$ is differentiable over type profile θ, and payoff of each agent i $u_i(\theta_i, s)$ is differentiable over social choice s and type θ_i;

or

b) standard AGV mechanism satisfies individual rationality.

The proof of this result is straightforward. Consider the (current) mechanism built in this paper. Part a) guarantees the payoff equivalence between the current mechanism and mechanism M: for any type of any agent, the expected interim utility differs in both mechanisms at a type-independent constant. One can adjust the transfers to the agents in a budget-balanced type-independent manner and make the current mechanism to satisfy individual rationality. Moreover, for each agent the interim utility in the current mechanism differs at a type-independent constant from AGV mechanism, allowing for a similar construction in part b).

5.2 Dynamics

It is possible to have a dynamic version of the current mechanism, in a multiple-period environment from Athey and Segal (2013). In that environment, in each period the social planner makes a social choice, and each agent i gets a payoff as a function of social choice and i's current type. The type of each agent evolves over time according to a Markov chain, depending on the social choice and independently of other agents' types. The social planner maximizes the total expected discounted sum of payoffs of all the agents. Athey and Segal extend AGV mechanism to this dynamic setting. In their mechanism in each period agent i is paid the change in expected efficient continuation payoff of others, caused by i's report. Their efficient mechanism makes truthful report a Bayesian equilibrium, and is ex post budget-balanced in each period.

The idea of the current mechanism makes it possible to add full implementation and coalition-proofness. To show this, let's first introduce some notations. Assuming all the agents always report truthfully, let the type profile to be θ in the current period. Respectively, the efficient social choice $s^*(\theta)$ is chosen, and each agent i obtains payoff $u_i(\theta)$. Agent i's expected efficient continuation payoff for the future period is denoted as $V_i(\theta)$; the expectation is taken at the beginning of the future period over the future type profile. Let's also denote the discount factor as δ, which is the same for all the agents. At the moment between the agents having announced type θ and the social planner making choice $s^*(\theta)$, the continuation payoff for agent i is denoted as $U_i(\theta) \equiv u_i(\theta) + \delta V_i(\theta)$. Since the social planner knows current type θ, she has correct beliefs about the future type distribution, θ^{+1}. In the mechanism the agents report their types sequentially; as the reports are getting submitted, i's continuation payoff changes in the future period from $V_i(\theta)$ to $U_i(\theta^{+1})$. When agent i reports own type, any other agent j pays i the related change in j's expected continuation payoff.

Proposition 5 *If agent i reported truthfully her type θ_i^{-1} in the past period and she commits to always report truthfully before learning current type θ_i, then she guarantees herself the continuation payoff of $V_i(\hat{\theta}_{-i}^{-1}, \theta_i^{-1})$, where $\hat{\theta}_{-i}^{-1}$ is others' report in the past period.*

Proof.

The proof will use induction. Let's suppose at period $t+1$ agent i can guarantee herself the continuation payoff $V_i(\hat{\theta}_{-i}, \theta_i)$[6] if she commits to report truthfully from period $t+1$ onwards, and she reports truthfully type θ_i at period t. Let's show that at period t agent i can guarantee herself the continuation payoff $V_i(\hat{\theta}_{-i}^{-1}, \theta_i^{-1})$, if she reported truthfully type θ_i^{-1} at period $t-1$ and commits to report truthfully at period t. Agent i's continuation payoff, estimated just before the social choice is made at period t, equals $U_i(\hat{\theta}_{-i}, \theta_i) = u_i(\hat{\theta}_{-i}, \theta_i) + \delta V_i(\hat{\theta}_{-i}, \theta_i)$. Moreover, since agent i reported truthfully at period $t-1$, the social planner has correct beliefs about the distribution of type θ_i in period t. In period t any other agent $j \neq i$ is expected to pay zero to i. On the other hand, if other agents report their types in period t as $\hat{\theta}_{-i}$, agent i will pay other agents the value $U_i(\hat{\theta}_{-i}, \theta_i) - V_i(\hat{\theta}_{-i}^{-1}, \theta_i^{-1})$. This means that the continuation payoff of agent i at the beginning of period t equals to

$$U_i(\hat{\theta}_{-i}, \theta_i) - [U_i(\hat{\theta}_{-i}, \theta_i) - V_i(\hat{\theta}_{-i}^{-1}, \theta_i^{-1})] = V_i(\hat{\theta}_{-i}^{-1}, \theta_i^{-1})$$

[6] The continuation payoff $V_i(\hat{\theta}_{-i}, \theta_i)$ is estimated before agent i learns her type in period $t+1$.

Q.E.D.

Since at the beginning of the game social planner has the correct beliefs about future types' distribution, every agent guarantees herself her ex ante efficient continuation payoff, making any equilibrium efficient and coalition-proof. However, off-equilibrium path one may have inefficient equilibria. That is, if it so happens that agent i has misreported at some period and induced incorrect beliefs for social planner, agent i no longer guarantees to get the efficient continuation payoff and may prefer to continue misreporting in future periods.

6 Discussion

This paper provides an efficient mechanism with agents having independent private types and quasilinear preferences. In the mechanism each pair of agents directly exchanges monetary transfers according to the externalities they impose on each other. The transfers internalize externalities, achieving coalition-proofness and full efficient implementation. This method of removal of externalities can also be applied in a dynamic setting.

The assumption of independent private types is crucial for the mechanism. The removal of externalities is based on two effects: if any agent i commits to report truthfully, then a) other agents expect each to pay zero to i, and b) agent i makes the payments to other agents which keep i's utility constant. If types of the agents are interdependent, then part a) may no longer hold. Any agent $j \neq i$ would hold beliefs about i's type dependent on own private type θ_j and therefore might expect to make a non-zero payment to i. This may induce agent j to misreport and make the social planner to underestimate the externality of i's report on j's payoff. The current method of trading externalities cannot therefore be applied to the general setting, and it remains an open question how to extend it.

References

Arrow, K. (1979) "The Property Rights Doctrine and Demand Revelation Under Incomplete Information," in , *Economics and Human Welfare*, ed. by M. Boskin. New York: Academic Press.

Athey, S., and I. Segal (2013) "An Efficient Dynamic Mechanism," *Econometrica*, Vol. 81, pp. 2463-2485.

Ausubel, L., and O. Baranov. (2013) "Core-Selecting Auctions with Incomplete Information," *working paper*.

Biran, O., and F. Forges (2011) "Core-Stable Rings in Auctions with Independent Private Values," *Games and Economic Behavior*, Vol. 73, pp. 52-64.

Bergemann, D., and J. Välimäki (2010) "The Dynamic Pivot Mechanism," *Econometrica*, Vol. 78, pp. 771-789.

Che, Y., and J. Kim (2006) "Robustly Collusion-Proof Implementation," *Econometrica*, Vol. 74, pp. 1063-1107.

——— (2009) "Optimal Collusion-Proof Auctions," *Journal of Economic Theory*, Vol. 144, pp. 565-603.

Che, Y., Condorelli, D., and J. Kim (2013) "Weak Cartels and Collusion-Proof Auctions," *working paper*.

Chen, J., and S. Micali (2012) "Collusive Dominant-Strategy Truthfullness," *Journal of Economic Theory*, Vol. 147, pp. 1300-1312.

Chen, Y.-C., Kunimoto, T., and Y. Sun (2015) "Implementation with Transfers," *Discussion Paper No, 2015-04*.

Clarke, E. (1971) "Multipart Pricing of Public Goods," *Public Choice*, Vol. 11, pp. 19-33.

Cramton, P., Gibbons, R., and P. Klemperer (1987) "Dissolving a Partnership Efficiently," *Econometrica*, Vol. 55, pp. 615-632.

Crémer, J. (1996) "Manipulations by Coalitions Under Asymmetric Information: The Case of Groves Mechanism," *Games and Economic Behavior*, Vol. 13, pp. 39-73.

D'Aspresmont, C., and L. Gerard-Varet. (1979) "Incentives and Incomplete Information," *Journal of Public Economics*, Vol. 11, pp. 25-45.

Erdil, A., and P. Klemperer (2011) "A New Payment Rule For Core-Selecting Package Auctions," *Journal of the European Economic Association*, Vol. 8, pp. 537-547.

Groves, T. (1973) "Incentives in Teams," *Econometrica*, Vol. 41, pp. 617-631.

Jackson, M. (1991) "Bayesian Implementation," *Econometrica*, Vol. 59, pp. 461-477.

Krishna, V., and M. Perry. (2000) "Efficient Mechanism Design," *unpublished.*

Laffont, J., and D. Martimort (1997) "Collusion Under Asymmetric Information," *Econometrica*, Vol. 65, pp. 875-911.

——— (1998) "Collusion and Delegation," *The Rand Journal of Economics*, Vol. 29, pp. 280-305.

——— (2000) "Mechanism Design With Collusion and Correlation," *Econometrica*, Vol. 68, pp. 309-342.

Maskin, E. (1998) "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, Vol. 66, pp. 23-38.

Matsushima, H. (1993) "Bayesian Monotonicity with Side Payments," *Journal of Economic Theory*, Vol. 59, pp. 107-121.

McAfee, R. and J. McMillan (1992) "Bidding Rings," *American Economic Review*, Vol. 82, pp. 579-599.

Myerson, R. (2007) "Virtual utility and the core for games with incomplete information," *Journal of Economics Theory*, Vol. 136, pp. 260-285.

Ollar, M. and A. Penta (2015) "Full Implementation and Belief Restrictions," under revision in *American Economic Review*.

Palfrey, T. and S. Srivastava (1989) "Implementation with Incomplete Information in Exchange Economies," *Econometrica*, Vol. 57, pp. 115-134.

Palfrey, T. and S. Srivastava (1989) "Mechanism Design with Incomplete Information: A Solution to the Implementation Problem," *Journal of Political Economy*, Vol. 97, pp. 668-691.

Samuelson, W. (1985) "A Comment on the Coase Theorem," in *Game-Theoretic Models of Bargaining*, ed. by A. Roth. Cambridge University Press.

Vickrey, W. (1961) "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, Vol. 16, pp. 8-37.

Williams, S. (1999) "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms" *Economic Theory*, Vol. 14, pp. 155-180.