Balanced Bayesian Mechanisms*

Claude d'Aspremont
CORE, Université catholique de Louvain, Louvain-la-Neuve, Belgium
Jacques Crémer
CNRS, IDEI and GREMAQ, University of Toulouse 1, France
Louis-André Gérard-Varet
EHESS, GREQAM, Marseille, France

Abstract

In the transferable utility case, a number of authors have identified conditions on beliefs that guarantee the existence of Bayesian incentive compatible mechanisms with balanced transfers. We present a new, easy to interpret, condition and we show that it is (strictly) more general than all the other conditions found in the literature. We also study conditions guaranteeing the Bayesian implementability of all social decision rules with balanced budget mechanisms.

JEL Classification: C7; D7; D82; H4
Keywords: Mechanism design; Bayesian implementation; Bayesian incentive compatibility

1. Introduction

In many resource allocation problems, the general properties of decentralized information structures may be exploited in the construction of decision procedures or contracts to achieve an optimal allocation, despite market failures due to externalities or to public goods, and despite strategic behavior, free-riding and misrepresentation of preferences.[1] In the transferable utility case, and for dominant strategy mechanisms, Green and Laffont (1979) and Walker (1980) have shown that it is in general impossible to balance the budget. On the other hand, as is by now well known, in Bayesian frameworks where the structure of agents beliefs is explicitly taken into account, one can find Bayesian incentive compatible mechanisms (BIC-mechanisms) that balance the budget.[2]

The main purpose of this note is to clarify, in the transferable utility case, the relationship between the conditions on beliefs that have been presented in the literature and shown to guarantee the existence of BIC-mechanisms that implement efficient decision rules with balanced transfers.[3] This clarification is based on a new condition, condition C, that is less restrictive than previous ones. This condition is easier to interpret than the already existing equivalent conditions, namely condition

$*$. Reprinted from *Journal of Economic Theory*, 115, 385–396, 2004.
The research presented in this paper was entirely completed before Louis-André Gérard's death in January 2001. The other two authors are responsible for a substantial amount of editing. Support by the Belgian Federal Public Planning Service for Science Policy is gratefully acknowledged.

1. The first applications belonged to public economics and to the study of collective decision making and auctions (see Clarke, 1971; Groves, 1973), but there are many other applications.
2. See among others, d'Aspremont and Gérard-Varet (1975, 1979, 1982), Arrow (1979), Laffont and Maskin (1979), d'Aspremont et al. (1990), Johnson et al. (1990).
3. Here we mean balancing the budget ex post (for all states of the world). That no condition is required to balance the budget ex ante is a known fact (d'Aspremont and Gérard-Varet, 1982, Theorem 9).

C^* introduced[4] by d'Aspremont and Gérard-Varet (1982) and LINK introduced by Johnson et al. (1990).

Condition C has an important consequence: it "guarantees budget balance", in the sense that it ensures that any BIC-mechanism can be transformed into a BIC-mechanism that balances the budget. Since for efficient decision rules, Vickrey-Clarke-Groves mechanisms are BIC-mechanisms and always exist, it is clear that when condition C holds it is possible to build a balanced budget BIC-mechanism. Other authors have presented conditions that guarantee budget balance (*pairwise identifiability* introduced by Fudenberg et al. (1994, 1996) as well as Assumption I (i) in Aoyagi (1998), rebaptized *weak regularity* by Chung (1999), following Matsushima (1991)). Because, as we show, condition C is both necessary and sufficient for an information structure to guarantee budget balance, it is less restrictive than these other conditions, and, through examples, we show that it is strictly less restrictive.

In the final section of the note, we study stronger conditions that guarantee that all decision rules, even those that are not efficient, can be implemented while balancing the budget, a property that might be useful in many specific problems (typically those involving only a subset of all agents). We show that a necessary and sufficient condition for this property is the already known (see d'Aspremont and Gérard-Varet, 1982) condition B and that it is a weaker condition than the *strict regularity* condition of Aoyagi (1998). Finally, we show that the mechanisms can easily be constructed through a 'scoring rules' method. We also show that, loosely speaking, the class of beliefs satisfying condition C can be partitioned into those that satisfy condition B and those that satisfy a (very weak) independence property.

The results presented here leave open a number of important questions, some of which we answer in d'Aspremont et al. (2003). In particular, among other results, we show (a) that, even though it is a very general condition, condition C is not necessary for implementation of efficient Bayesian mechanisms and (b) that it is not true that efficient Bayesian mechanisms always exist.

2. Bayesian incentive compatible mechanisms

We consider a set $\mathcal{N}$ of $n \geq 3$ agents.[5] All the private information of agent $i \in \mathcal{N}$ is represented by his type α_i which belongs to a finite set (with at least two elements) $\mathcal{A}_i$. An n-vector of possible types is denoted α and is an element of $\mathcal{A} = \prod_{i \in \mathcal{N}} \mathcal{A}_i$.

The utility function of agent i of type α_i is defined over a set $\mathcal{X}$ of public decisions, and utility is 'transferable': for $x \in \mathcal{X}$ and a monetary transfer $t_i \in \mathbb{R}$, his utility is $u_i(x; \alpha_i) + t_i$. Some of our results hold when utility functions have the more general form $u_i(x; \alpha) + t_i$, i.e. are *mutually payoff-relevant* (Johnson et al., 1990).

The type of agent i also determines his beliefs about the types of the other agents. When he is of type α_i they are represented by a probability distribution over $\mathcal{A}_{-i} = \prod_{j \in \mathcal{N}-i} \mathcal{A}_j$, the set of the possible types of the other agents. A generic element of $\mathcal{A}_{-i}$ will be denoted α_{-i}; we will sometimes use the notation $\alpha_{-i-j} \in \mathcal{A}_{-i-j}$ to denote a vector of possible types of all agents but i and j. We

Also, in this paper, we shall neglect the issue of unicity of equilibrium, so that our use of terms is different from the language in the literature on implementation (see Palfrey, 2002). In d'Aspremont et al. (1999, 2003), we discuss conditions that guarantee 'unique implementation' both for auctions and balanced Bayesian mechanisms.

4. It was introduced as condition C (without a star) but can be shown equivalent to our new condition (that we will also call condition C) by a simple duality argument.

5. When $n = 2$, condition C is equivalent to independence of types, and condition B never holds (see d'Aspremont and Gérard-Varet, 1982).

assume that there exists a probability distribution p over $\mathcal{A}$ such that the beliefs $p(\alpha_{-i} \mid \alpha_i)$ of agent i of type α_i are obtained by conditioning p with respect to α_i (the beliefs are 'consistent'), and that $p(\alpha_i) > 0$ for all i and all α_i. Similarly, for $i \neq j$, the beliefs of agent i on the types of agents other than himself or j are $p(\alpha_{-i-j} \mid \alpha_i) = \sum_{\alpha_j \in \mathcal{A}_j} p(\alpha_{-i} \mid \alpha_i)$. The triplet $(\mathcal{N}, \mathcal{A}, p)$ is called an *information structure*. An *environment* is composed of an information structure, together with a set of outcomes and utility functions for the agents and is denoted $(\{\mathcal{N}, \mathcal{A}, p\}, \mathcal{X}, \{u_i\}_{i\in\mathcal{N}})$.

A *public decision rule* s is a function from $\mathcal{A}$ into $\mathcal{X}$: for a vector α of types the public decision $s(\alpha)$ is taken. It is *efficient* if $\sum_{i\in\mathcal{N}} u_i(s(\alpha); \alpha_i) \geq \sum_{i\in\mathcal{N}} u_i(x; \alpha_i)$ for all $\alpha \in \mathcal{A}$ and all $x \in \mathcal{X}$.

The problem is to implement a decision rule s when the decision mechanism must be based on private information revealed by the agents. Invoking the revelation principle[6], we restrict ourselves to *direct mechanisms* in which agents are induced to truthfully reveal their type to the planner; such a direct mechanism is defined by a decision rule $s : \mathcal{A} \to \mathcal{X}$ and a transfer rule $t : \mathcal{A} \to \mathbb{R}^n$.

We will say that an information structure $(\mathcal{N}, \mathcal{A}, p)$ *guarantees implementation of efficient public decision rules* if for every outcome set[7] $\mathcal{X}$, every utility functions $u_i : \mathcal{X} \times \mathcal{A}_i \to \mathbb{R}$, $i = 1, \ldots, n$, and any efficient public decision rule s, we can find a transfer rule t which *balances the budget*, i.e., that satisfies $\sum_{i\in\mathcal{N}} t_i(\alpha) = 0$, for all $\alpha \in \mathcal{A}$, and such that the associated direct mechanism (s, t) satisfies the *Bayesian incentive compatibility (BIC)* constraints

$$\begin{aligned} &\sum_{\alpha_{-i}\in\mathcal{A}_{-i}} p(\alpha_{-i} \mid \alpha_i)[u_i(s(\alpha_i, \alpha_{-i}); \alpha_i) + t_i(\alpha_i, \alpha_{-i})] \\ \geq &\sum_{\alpha_{-i}\in\mathcal{A}_{-i}} p(\alpha_{-i} \mid \alpha_i)[u_i(s(\tilde{\alpha}_i, \alpha_{-i}); \alpha_i) + t_i(\tilde{\alpha}_i, \alpha_{-i})] \end{aligned} \tag{1}$$

for all $i \in \mathcal{N}$, and all α_i and $\tilde{\alpha}_i \in \mathcal{A}_i$. *Strict implementation* is obtained when all (*BIC*)-inequalities hold strictly.[8] If, in the preceding definition, 'any efficient decision rule' is replaced by 'any decision rule', the information structure *guarantees implementation of all decision rules.*

In this framework, most effort in the literature has been devoted to finding information structures that guarantee implementation of efficient public decision rules with no additional restriction on the utility functions. We keep this approach. We introduce no individual rationality constraint. However, all our results hold true if we add an *ex ante* individual rationality constraint of the form $\sum_{\alpha\in\mathcal{A}} p(\alpha)[u_i(s(\alpha); \alpha_i) + t_i(\alpha)] \geq 0$, as long as there is a status quo decision that guarantees each agent a utility of 0. This is appropriate for many applications in which the contract is signed before the agents acquire information about their types. For instance, this model has been used to study agreements to reduce pollution (Duggan and Roberts, 1999), joint research ventures (d'Aspremont et al., 1998; Bhattacharya et al., 1992), and the contracts between a firm and suppliers (Cremer and Riordan, 1987; Riordan, 1983, 1984). In all these cases, it is assumed that the parties have symmetric but imperfect information before contracting, and acquire private information afterwards. The same dynamic of information acquisition is supposed in the recent theory of the core solution concept in cooperative games of incomplete information (e.g. Forges et al., 2002) using balanced Bayesian

6. For a statement and references see Fudenberg and Tirole (1992).
7. Because we have only a finite set of types, only a finite subset of decisions are really relevant. The fact that the set $\mathcal{X}$ varies does not create any difficulty, and we could keep it fixed without changing the results it its cardinality was at least equal to that of $\mathcal{A}$.
8. Notice that we have imposed no uniqueness of equilibrium requirement. In d'Aspremont et al. (1999), we show how equivalent mechanisms with a single equilibrium can be constructed in nearly all environments.

mechanisms (and some of the conditions presented below) to represent the bargaining that takes place within coalitions.

3. Conditions that guarantee implementation of efficient decision rules

Many conditions have been introduced to guarantee implementation of efficient public decision rules. We will review them, but we start by introducing a condition that will turn out to be weaker than all others, and that we call condition C because a simple duality argument shows it to be equivalent to a condition introduced by d'Aspremont and Gérard-Varet (1982) to which we shall refer here as condition[9] C^*. It simply states that beliefs are such that we can collect from the agents any aggregate transfer, dependent on the state of nature, without inciting them to lie.

3.1 Condition C

An information structure satisfies *condition C* if and only if for every function $R : \mathcal{A} \to \mathbb{R}$, there exists a transfer rule t^C such that for all $\alpha \in \mathcal{A}$

$$\sum_{i\in\mathcal{N}} t_i^C(\alpha) = R(\alpha) \tag{2}$$

and such that for all $i \in \mathcal{N}$ and all α_i and $\tilde{\alpha}_i$ in $\mathcal{A}_i$, $\alpha_i \neq \tilde{\alpha}_i$, we have

$$\sum_{\alpha_{-i}\in\mathcal{A}_{-i}} t_i^C(\alpha_{-i}, \alpha_i) p(\alpha_{-i} \mid \alpha_i) \geq \sum_{\alpha_{-i}\in\mathcal{A}_{-i}} t_i^C(\alpha_{-i}, \tilde{\alpha}_i) p(\alpha_{-i} \mid \alpha_i). \tag{3}$$

We show next that condition C is both necessary and sufficient to ensure that any BIC-mechanism can be modified into a BIC-mechanism that balance the budget.[10] To state this formally, we will say that an information structure *guarantees budget balance* if, given any set $\mathcal{X}$ of public decisions and any utility functions $\{u_i\}_{i\in\mathcal{N}}$, there exists, for any BIC-mechanism (s, t), another transfer rule t' that balances the budget and such that (s, t') is also a BIC-mechanism.

Lemma 1 *An information structure guarantees budget balance if and only if it satisfies condition C.*

Proof Given $\mathcal{X}$ and $\{u_i\}_{i\in\mathcal{N}}$, consider an information structure $(\mathcal{N}, \mathcal{A}, p)$ that satisfies condition C and a BIC-mechanism (s, t). Let $R = -\sum_{i\in\mathcal{N}} t_i$. By condition C, there exists a transfer rule t^C that satisfies (2) and (3). The transfer rule $t' = t + t^C$ balances the budget and provides the correct incentives.

To show the reverse implication, consider an information structure $(\mathcal{N}, \mathcal{A}, p)$ that guarantees budget balance. Choose any function $R : \mathcal{A} \to \mathbb{R}$. Pick a payoff structure and a decision rule s such

9. In the consistent case, condition C^* is: If, $\forall\, \alpha \in \mathcal{A}, \forall\, i \in \mathcal{N}$,

$$[p(\alpha_{-i} \mid \alpha_i) \sum_{\tilde{\alpha}_i\neq\alpha_i} \lambda_i(\tilde{\alpha}_i, \alpha_i) - \sum_{\tilde{\alpha}_i\neq\alpha_i} \lambda_i(\alpha_i, \tilde{\alpha}_i) p(\alpha_{-i} \mid \tilde{\alpha}_i)] = \kappa(\alpha),$$

for some $\lambda_i : \mathcal{A}_i \times \mathcal{A}_i \to \mathbb{R}_+$ $(i = 1, \ldots, n)$ and $\kappa : \mathcal{A} \to \mathbb{R}$, then κ must be identically zero. To obtain the *compatibility condition* of d'Aspremont and Gérard-Varet (1979), just replace each $\lambda_i(\tilde{\alpha}_i, \alpha_i)$ by $\lambda_i(\alpha_i, \tilde{\alpha}_i)$ on the left-hand side of the equalities. It was shown to be strictly weaker than C^* by Johnson et al. (1990). Through a counterexample, Chung (1999) shows that the regularity condition in Matsushima (1991) does not imply the compatibility condition. However, as we shall prove, it does imply condition C^*.

10. Johnson et al. (1990), Proposition 5.4, proves this result for condition C^*, but the argument is more intricate.

that $u_i(s(\alpha); \tilde{\alpha}_i) = R(\alpha)/n$ for all α and all $\tilde{\alpha}_i$. It is easy to verify that if we set $t_i(\alpha) = -R(\alpha)/n$ for all α, the inequalities (BIC) hold (with both sides being equal to each other).

Because budget balanced is guaranteed, there exists a transfer function $\tilde{t}$ that satisfies (1) (with t replaced by $\tilde{t}$) for all α, all $\tilde{\alpha}_i$ and all i, as well as $\sum_i \tilde{t}_i(\alpha) = 0$ for all α. The transfers t^C, with $t_i^C(\alpha) = u_i(s(\alpha); \alpha_i) + \tilde{t}_i(\alpha) = (R(\alpha)/n) + \tilde{t}_i(\alpha)$, satisfy (2) and (3), which proves the result.[11] ■

Lemma 1 yields an immediate proof of the following theorem.[12]

Theorem 1 *Any information structure that satisfies condition C guarantees implementation of efficient decision rules.*

Proof Consider any $\mathcal{X}$, any $\{u_i\}_{i \in \mathcal{N}}$ and any efficient decision rule s. The (BIC) constraints can be satisfied using transfers of the Vickrey-Clarke-Groves type, i.e., $t_i^G(\alpha) = \sum_{j \neq i} u_j(s(\alpha); \alpha_j)$, since, as well known, these transfers implement any efficient decision rule in dominant strategies. The result follows from Lemma 1. ■

Notice that the definition of condition C, as well as the argument of the theorem (and Lemma 1) could be generalized to environments where the sets of types are not assumed to be finite.[13] However, in the finite case, to verify that condition C holds, it is sufficient to show that it can be solved for a finite number of functions $\tilde{R}$, those which satisfy $|\tilde{R}(\alpha')| = 1$ for some $\alpha' \in \mathcal{A}$ and $\tilde{R}(\alpha) = 0$ for all $\alpha \neq \alpha'$. Indeed any function $R : \mathcal{A} \to \mathbb{R}$ is a positive linear combination of these functions $\tilde{R}$. Of course, if we are interested by a specific environment, by the argument of Lemma 1, we can simply start by constructing a BIC-mechanism (s, t), define $R = -\sum_{i \in \mathcal{N}} t_i$, and check whether there exists a transfer rule $\tilde{t}$ that satisfies equations (2) and (3). If it exists, then the transfer rule $t' = t + \tilde{t}$ balances the budget and provides the correct incentives.

Finally, in d'Aspremont et al. (1990), we showed that, with finitely many types, condition C holds for nearly all[14] information structures (see also d'Aspremont et al., 2003, for a simpler proof).

3.2 Other conditions

As known, since d'Aspremont and Gérard-Varet (1975) and Arrow (1979), guaranteeing implementation of any efficient decision rule can be obtained by assuming some form of independence of types. Formally, an agent i is said to have *free beliefs on a pair* $\{\alpha_i, \alpha_i'\}$ if for any α_{-i} we have $p(\alpha_{-i} \mid \alpha_i) = p(\alpha_{-i} \mid \alpha_i')$. Agent i *has free beliefs* if he has free beliefs on all pairs of types. *Independence of types* holds when all agents have free beliefs. Independence of types implies that condition C holds. As shown below, it holds even if only one agent has free beliefs (Cremer and Riordan, 1985, had shown directly that, in this case, it is possible to implement efficient decision rules).

Interestingly, the other conditions introduced in the literature and shown to guarantee implementation of efficient decision rules limit in some ways the degree of independence between the types.

11. It is clear that, with the definitions modified accordingly, this result holds for mutually payoff-relevant utility functions. This is not true for the following theorem, which relies on Vickrey-Clarke-Groves mechanisms.
12. This result was first proved for the 'compatibility condition' in d'Aspremont and Gérard-Varet (1979) and for condition C^* in d'Aspremont and Gérard-Varet (1982). The present proof is much more immediate.
13. It is much easier for C than for the dual condition C^* (see d'Aspremont and Gérard-Varet, 1982).
14. By 'nearly all' we mean on an open and dense subset of the set of probability distributions (which is itself a subset of $\mathbb{R}^K$, where $K = \sum_{i \in \mathcal{N}} \#\mathcal{A}_i$ and $\#\mathcal{A}_i$ is the cardinality of $\mathcal{A}_i$).

The first of these conditions is Matsushima (1991) *regularity condition*, which holds if, for some pair of agents (i, j), the vectors $\{p(\alpha_{-i-j} \mid \alpha_i)\}_{\alpha_{-i-j} \in \mathcal{A}_{-i-j}}$ are linearly independent (note that, for a given pair (i, j), the dimension of these vectors is equal to the cardinality of the set $\mathcal{A}_{-i-j}$, and their number is equal to the cardinality of $\mathcal{A}_i$).

Later, Chung (1999) analyzed a *weak regularity* condition, which is equivalent to Assumption I (i) in Aoyagi (1998). It holds if there exists a pair of agents (i, j) such that, for all pairs (α_i, α'_i) in $\mathcal{A}_i \times \mathcal{A}_i$,

$$\{p(\alpha_{-i-j} \mid \alpha_i)\}_{\alpha_{-i-j} \in \mathcal{A}_{-i-j}} \neq \{p(\alpha_{-i-j} \mid \alpha'_i)\}_{\alpha_{-i} \in \mathcal{A}_{-i-j}}.$$

Finally, Fudenberg et al. (1994, 1996) introduced the 'pairwise identifiability condition'. For any $i \in \mathcal{N}$, any $\alpha \in \mathcal{A}$, and any 'deviation' $\tilde{a}_i(a_i)$ (a function from $\mathcal{A}_i$ into itself), define $p(\tilde{a}_i = \alpha_i, \alpha_{-i}) = \sum_{\{\alpha'_i \mid \tilde{a}_i(\alpha'_i) = \alpha_i\}} p(\alpha'_i, \alpha_{-i})$, and the $|\mathcal{A}_i|^{|\mathcal{A}_i|} \times |\mathcal{A}|$ matrix $\prod_i = [p(\tilde{a}_i = \alpha_i, \alpha_{-i})]_{(\tilde{a}_i, \alpha)}$ ($\prod_i$ is indexed both by the set of deviations and $\mathcal{A}$). Let

$$\prod_{ij} = \begin{bmatrix} \prod_i \\ \prod_j \end{bmatrix}.$$

An information structure satisfies *pairwise identifiability* if rank $\prod_{ij} =$ rank $\prod_i +$ rank $\prod_j -1$ for every pair of agents (i, j): deviations from truth telling by two agents generate sufficiently different probability distributions on the n-tuple of reports that a deviating agent can be identified.

As the following theorem shows, condition C is more general than all these conditions.

Theorem 2 *An information structure $(N, \mathcal{A}, p)$ satisfies condition C if any of the following condition holds: (i) at least one agent has free beliefs, (ii) weak regularity (or regularity) is satisfied, (iii) pairwise identifiability is satisfied. Furthermore, condition C is strictly less restrictive than these three conditions.*

Proof (i) If an agent i has free beliefs, say $p(\alpha_{-i})$, we can easily construct transfers to satisfy Eqs. (2) and (3). For any R, some i and all $j \neq i$, let:

$$t_i(\alpha) = R(\alpha) - \sum_{\alpha'_{-i} \in \mathcal{A}_{-i}} R(\alpha'_{-i}, \alpha_i) p(\alpha'_{-i})$$

and

$$t_j(\alpha) = \left[\sum_{\alpha'_{-i} \in \mathcal{A}_{-i}} R(\alpha'_{-i}, \alpha_i) p(\alpha'_{-i}) \right] / (n-1).$$

(ii) As proved by Aoyagi (1998), weak regularity (which is obviously implied by regularity) guarantees budget balance, and therefore implies condition C (by our Lemma 1). Furthermore, weak regularity is clearly more restrictive, since it is incompatible with independence of types, which implies C.[15]

(iii) Furthermore, Fudenberg et al. (1996) Lemma 1 shows that pairwise identifiability guarantees budget balance. Hence it implies condition C. To show that C is strictly less restrictive, consider the

15. Section 4 shows further that weak regularity is stronger than C and no independence of types.

following example. Let $\{\mathcal{N}, \mathcal{A}, q\}$ be an information structure, and for $i = 1, 2$, let $Q_i = [q(\tilde{a}_i = \alpha_i, \alpha_{-i})]_{(\tilde{a}_i, \alpha_{-i})}$. Assume that the rank of the matrix

$$Q_{12} = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix}$$

is strictly smaller than rank Q_1 + rank $Q_2 - 1$, so that q does not satisfy pairwise identifiability.[16] Add now an agent 0, so that we have a new information structure $\{\mathcal{N} \cup \{0\}, \mathcal{A} \times \mathcal{A}_0, p\}$, with $p(\alpha_{-0}, \alpha_0) \equiv q(\alpha_{-0})r(\alpha_0)$, where r is a probability distribution over $\mathcal{A}_0$. Defining the matrices Π_1 and Π_2 as above, it is straightforward that rank $\Pi_1 =$ rank Q_1, rank $\Pi_2 =$ rank Q_2 and rank $\Pi_{12} =$ rank Q_{12}. This implies that p does not satisfy pairwise identifiability, but it does satisfy condition C, since agent 0 has free beliefs. ■

It should be stressed, finally, that condition C can hold even when no agent has free beliefs (with $n \geq 3$); indeed, condition B introduced in the next section requires that there is no free beliefs and it implies condition C.

4. Conditions guaranteeing implementation of all decision rules

Sometimes a mechanism designer does not only try to maximize the interests of the participants in a mechanism; this will happen, for instance, if the participants in the mechanisms are representatives of the agents whose welfare the mechanism designer cares about, and if the incentives of the representatives are not perfectly aligned with the welfare of the agents. In order to study this problem, we will use another condition, called *condition B*, which was introduced by d'Aspremont and Gérard-Varet (1982). It assumes that there exists a balanced transfer rule t^B such that for all $i \in \mathcal{N}$ and all α_i and $\tilde{\alpha}_i$ in $\mathcal{A}_i$, $\alpha_i \neq \tilde{\alpha}_i$ we have

$$\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} t_i^B(\alpha_{-i}, \alpha_i) p(\alpha_{-i} \mid \alpha_i) > \sum_{\alpha_{-i} \in \mathcal{A}_{-i}} t_i^B(\alpha_{-i}, \tilde{\alpha}_i) p(\alpha_{-i} \mid \alpha_i). \tag{4}$$

Condition B is important because of the following theorem (which remains valid in the mutually payoff-relevant case).

Theorem 3 *Condition B is necessary and sufficient for an information structure to guarantee the implementation*[17] *of all decision rules.*

Proof It is straightforward to show that condition B is sufficient: for any environment one can multiply the transfers t_i^B by a sufficiently large positive number to ensure that the incentives for truthtelling derived from (4) dominate any incentives from misrepresentation stemming from the desire to change the public decision.[18]

16. Such information structures exist.
17. Note that Theorem 3 also holds if 'implementation' is replaced by 'strict implementation'.
18. This technique used here is similar to the techniques used by Crémer and McLean (1985); Cremer and McLean (1988), in the case of auctions. The mechanism designer convinces the agents to announce their true types by making them 'bet' on the announcements of the others. This can only yield truthful revelation when the agents' beliefs about the others depend on their own types.

To prove necessity, choose any $(i, \alpha_i^0) \in \mathcal{N} \times \mathcal{A}_i$, and a decision rule s that satisfy[19] $u_i(s(\alpha_{-i}, \alpha_i^0); \alpha_i^0) = -1$ for all α_{-i} and $u_j(s(\alpha); \tilde{\alpha}_j) = 0$ if $(j, \alpha_j, \tilde{\alpha}_j) \neq (i, \alpha_i^0, \alpha_i^0)$. Because the decision rule s can be implemented, there exists a balanced transfer function $t^{(i,\alpha_i^0)}$ that satisfies

$$\begin{aligned} &\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} p(\alpha_{-i} \mid \alpha_i^0)[t_i^{(i,\alpha_i^0)}(\alpha_{-i}, \alpha_i^0) - t_i^{(i,\alpha_i^0)}(\alpha_{-i}, \tilde{\alpha}_i)] \\ \geq &\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} p(\alpha_{-i} \mid \alpha_i^0)[u_i(s(\alpha_{-i}, \tilde{\alpha}_i); \alpha_i^0) - u_i(s(\alpha_{-i}, \alpha_i^0); \alpha_i^0)] = 1 \end{aligned} \tag{5}$$

for all $\tilde{\alpha}_i \neq \alpha_i^0$, and

$$\begin{aligned} &\sum_{\alpha_{-j} \in \mathcal{A}_{-j}} p(\alpha_{-j} \mid \alpha_j)[t_j^{(i,\alpha_i^0)}(\alpha_{-j}, \alpha_j) - t_j^{(i,\alpha_i^0)}(\alpha_{-j}, \tilde{\alpha}_j)] \\ \geq &\sum_{\alpha_{-j} \in \mathcal{A}_{-j}} p(\alpha_{-j} \mid \alpha_j)[u_j(s(\alpha_{-j}, \tilde{\alpha}_j); \alpha_j) - u_j(s(\alpha_{-j}, \alpha_j); \alpha_j)] = 0 \end{aligned} \tag{6}$$

for all $(j, \alpha_j) \neq (i, \alpha_i^0)$ and $\tilde{\alpha}_j \neq \alpha_j$.

The result is proved by repeating this construction for every $i \in \mathcal{N}$ and every $\alpha_i^0 \in \mathcal{A}_i$, and, using Eqs. (5) and (6), showing that the balanced budget transfer rule obtained by summing the $t^{(i,\alpha_i^0)}$'s over all $i \in \mathcal{N}$ and all $\alpha_i^0 \in \mathcal{A}_i$ satisfies (4). ■

Theorem 3 provides a characterization of the set of information structures in which asymmetry of information does not restrict the implementation of public decisions. When an information structure does not satisfy condition B, a mechanism designer will know that he must rely on properties of the information structure and of the utility functions (such as knowing that the decision rule is efficient). Theorem 3 also has a constructive side: its sufficiency part provides a method for building BIC-mechanisms.[20]

It is easy to see that condition B is incompatible with free beliefs, which would imply that the two sides of Eq. (4) are equal. But condition B is not simply a condition on the absence of free beliefs as there exist information structures where no agent has free beliefs, and which do not satisfy condition B (d'Aspremont et al., 2003). On the other hand, Aoyagi (1998) proposes a *strict regularity condition* adding to weak regularity the requirement that no agent has free beliefs on any two types. He shows that this condition guarantees implementation of all decision rules.

More precisely, the relationship between free beliefs, condition B, condition C and the strict regularity condition is summarized in the following theorem. In the proof, to construct the transfers t_i^B, we will use the 'scoring rules' method introduced by Good (1952), discussed by Savage (1974), and applied to Bayesian implementation in Johnson et al. (1990).

Theorem 4 *An information structure* $(\mathcal{N}, \mathcal{A}, p)$ *satisfies condition* B *if and only if it satisfies condition* C *and there exists no agent with free beliefs on any two of his types.*

The strict regularity condition is strictly more restrictive than condition B.

19. It is possible to find an environment where such a decision rule exists. Choose $\mathcal{X} = \{x_0, x_1\}$, $u_i(x_0; \alpha_i^0) = -1$, $u_i(x_1; \alpha_i^0) = 0$, $u_i(x; \alpha_i) = 0$ for all $x \in \mathcal{X}$ and all $\alpha_i \neq \alpha_i^0$, $u_j(x; \alpha_j) = 0$ for all $j \neq i$, all $x \in \mathcal{X}$ and all α_j. We are implementing a decision rule that minimizes the sum of the utilities of the agents!

20. This constructive argument can be transposed to environments where the sets of types are not finite, at least assuming that $\sup_{i,\alpha,\alpha_i'} |u_i(s(\alpha_{-i}, \alpha_i'); \alpha_i) - u_i(s(\alpha_{-i}, \alpha_i); \alpha_i)|$ is bounded.

In some cases, these mechanisms might have unpleasant properties as they could require large side payments; then, other techniques for finding mechanisms could be used.

Proof In the first statement, only the 'if' part remains to be proved. Assume that $(\mathcal{N}, \mathcal{A}, p)$ satisfies C and that no agent has free beliefs on two types. Take ε small enough and define the transfer rule θ by $\theta_i(\alpha) = \log p(\alpha_{-i} \mid \alpha_i)$ if $p_i(\alpha_{-i} \mid \alpha_i) > 0$ and $\theta_i(\alpha) = \varepsilon$ if $p_i(\alpha_{-i} \mid \alpha_i) = 0$. Then the following strict inequalities are easily verified due to the strict concavity of the function log

$$\sum_{\alpha_{-i} \in \mathcal{A}_{-i}} p_i(\alpha_{-i} \mid \alpha_i)\theta(\alpha_{-i}, \alpha_i) > \sum_{\alpha_{-i} \in \mathcal{A}_{-i}} p_i(\alpha_{-i} \mid \alpha_i)\theta(\alpha_{-i}, \tilde{\alpha}_i)$$

for all $(\alpha_i, \tilde{\alpha}_i) \in \mathcal{A}_i^2$.

Define R by $R(\alpha) = -\sum_{i \in \mathcal{N}} \theta_i(\alpha)$ for all α. Because the information structure satisfies C, there exists a transfer rule t^C satisfying Eqs. (2) and (3). For all α, let $t_i^B(\alpha) = \theta_i(\alpha) + t_i^C(\alpha)$. The transfer rule t^B is balanced and satisfies Eq. (4), and therefore $(\mathcal{N}, \mathcal{A}, p)$ satisfies condition B.

Aoyagi (1998), Theorem 1, shows that strict regularity of $(\mathcal{N}, \mathcal{A}, p)$ guarantees implementation of any public decision rule. Hence, by Theorem 3, it satisfies condition B. To show that it is strictly more restrictive than B, consider the following information structure $(\mathcal{N}, \mathcal{A}, p)$, where $\mathcal{N} = \{1, 2, 3\}$ and $\mathcal{A}_i = \{1, 2, 3\}$ for all i. Let p be equal to 0 if and only if exactly two of the agents have a type equal to either 1 or 3; for all other states of nature let p equal 1/15 (so that $p(1, 1, 2) = p(3, 1, 3) = 0$ but $p(1, 1, 1) = p(1, 3, 2) = 1/15$). This information structure is symmetric in the agents and in the types 1 and 3 for each agent. For all i we have

$$\{p(\alpha_j \mid \alpha_i = 1)\}_{\alpha_j \in \mathcal{A}_j} = \{p(\alpha_j \mid \alpha_i = 3)\}_{\alpha_j \in \mathcal{A}_j} = (1/4, 1/2, 1/4),$$

which contradicts weak regularity. To show that condition B holds, define the transfer rule t^B as follows:

$$\begin{array}{rclclclcl}
t^B(1,1,1) & = & t^B(2,2,2) & = & t^B(3,3,3) & = & (0,0,0), & & \\
t^B(1,2,3) & = & (1,-2,1); & & t^B(1,2,2) & = & t^B(3,2,2) & = & (8,-4,-4) \\
t^B(1,1,2) & = & (3,3,2) & = & t^B(1,1,3) & = & t^B(3,3,1) & = & (-8,-8,16).
\end{array}$$

All other transfers are constructed by permutations on the agents (for example: $t^B(2, 1, 3) = (-2, 1, 1)$). The transfer rule t^B is clearly balanced and may be checked to satisfy (4). The result follows.[21] ∎

As a final remark, notice that the 'scoring rule' used in the theorem suggests an easy technique to construct more generally the transfer rule t^B ensuring condition B. Consider an information structure $(\mathcal{N}, \mathcal{A}, p)$ and let addition and subtraction on the indices of agents be defined modulo n so that $n + 1 \equiv 1$ and $1 - 1 \equiv n$. For all i, all α_i, and all $\tilde{\alpha}_i$, we assume (and this holds generically[22]) either that $p_i(\alpha_{-i-(i-1)} \mid \alpha_i) \neq p_i(\alpha_{-i-(i-1)} \mid \tilde{\alpha}_i)$ for some $\alpha_{-i-(i-1)}$, or that $p_i(\alpha_{-i-(i+1)} \mid \alpha_i) \neq p_i(\alpha_{-i-(i+1)} \mid \tilde{\alpha}_i)$ for some $\alpha_{-i-(i+1)}$. Define $t^B(\alpha)$ by

$$\begin{aligned}
t_i^B(\alpha) = {} & [\log p_i(\alpha_{-i-(i-1)} \mid \alpha_i) - \log p_i(\alpha_{-(i-1)-i} \mid \alpha_{i-1})] \\
& + [\log p_i(\alpha_{-i-(i+1)} \mid \alpha_i) - \log p_i(\alpha_{-(i+1)-i} \mid \alpha_{i+1})].
\end{aligned}$$

21. It is important to remark that the results would still hold true if the 0 probabilities in the information structure were replaced by a small enough ε, adapting the other probabilities accordingly.

22. Genericity has been demonstrated before, not only for condition C^*, as already mentioned, but also for pairwise identifiability (see Fudenberg et al., 1996) and for strict regularity (see Aoyagi, 1998). Of course, it cannot hold for free beliefs.

The negative terms are constant in α_i and do not influence the incentives of agent i, but they ensure that the rule is balanced. The strict concavity of the function $\log$ implies that for all i and all α_i, inequality (4) holds, and therefore that condition B is satisfied.

References

Masaki Aoyagi. Correlated types and bayesian incentive compatible mechanisms with budget balance. *Journal of Economic Theory*, 79(1):142–151, 1998.

Kenneth J Arrow. The property rights doctrine and demand revelation under incomplete information. In M.J. Boskin, editor, *Economics and Human Welfare*. Academic Press, New York, 1979.

Sudipto Bhattacharya, Jacob Glazer, and David EM Sappington. Licensing and the sharing of knowledge in research joint ventures. *Journal of Economic Theory*, 56(1):43–69, 1992.

Kim-Sau Chung. A note on matsushima's regularity condition. *Journal of Economic Theory*, 87(2): 429–433, 1999.

Edward H Clarke. Multipart pricing of public goods. *Public choice*, II:19–33, 1971.

Jacques Crémer and Richard P McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53:345–361, 1985.

Jacques Cremer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56:1247–1257, 1988.

Jacques Cremer and Michael H Riordan. A sequential solution to the public goods problem. *Econometrica*, 53(1):77–84, 1985.

Jacques Cremer and Michael H Riordan. On governing multilateral transactions with bilateral contracts. *RAND Journal of Economics*, 18:436–451, 1987.

Claude d'Aspremont and Louis-André Gérard-Varet. Incentive games with incomplete information: An application to a public input model. In *Colloque International de Théorie des Jeux*. Institute des Hautes Etudes de Belgique, 1975.

Claude d'Aspremont and Louis-André Gérard-Varet. Incentives and incomplete information. *Journal of Public economics*, 11(1):25–45, 1979.

Claude d'Aspremont and Louis-André Gérard-Varet. Bayesian incentive compatible beliefs. *Journal of Mathematical Economics*, 10(1):25–45, 1982.

Claude d'Aspremont, Jacques Crémer, and Louis-André Gérard-Varet. Incentives and the existence of pareto-optimal revelation mechanisms. *Journal of Economic Theory*, 51(2):233–254, 1990.

Claude d'Aspremont, Sudipto Bhattacharya, and Louis-André Gérard-Varet. Knowledge as a public good: efficient sharing and incentives for development effort. *Journal of Mathematical Economics*, 30(4):389–404, 1998.

Claude d'Aspremont, Jacques Crémer, and Louis-André Gérard-Varet. Unique implementation in auctions and in public goods problems. In C. d'Aspremont, editor, *Social Organization and Mechanism Design*. De Boeck Université, Brussels, 1999.

Claude d'Aspremont, Jacques Crémer, and Louis-André Gérard-Varet. Correlation, independence, and bayesian incentives. *Social Choice and Welfare*, 21(2):281–310, 2003.

John Duggan and Joanne Roberts. Implementing the efficient allocation of pollution. *American Economic Review*, 92(4):1070–1078, 1999.

Françoise Forges, Jean-François Mertens, and Rajiv Vohra. The ex ante incentive compatible core in the absence of wealth effects. *Econometrica*, 70(5):1865–1892, 2002.

Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1992.

Drew Fudenberg, David K Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica*, pages 997–1039, 1994.

Drew Fudenberg, David K Levine, and Eric Maskin. Balanced-budget mechanisms with incomplete information. Dk levine's working papers archive, UCLA Department of Economics, 1996.

Irving John Good. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 107–114, 1952.

Jerry R Green and Jean-Jacques Laffont. *Incentives in public decision making*. North-Holland, 1979.

Theodore Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, 14: 617–631, 1973.

Scott Johnson, John W Pratt, and Richard J Zeckhauser. Efficiency despite mutually payoff-relevant private information: The finite case. *Econometrica*, 58:873–900, 1990.

Jean-Jacques Laffont and Eric Maskin. A differential approach to expected utility maximizing mechanisms. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North Holland, Amsterdam, 1979.

Hitoshi Matsushima. Incentive compatible mechanisms with full transferability. *Journal of Economic Theory*, 54(1):198–203, 1991.

Thomas R Palfrey. Implementation theory. In S. Hart R. Aumann, editor, *Handbook of game theory with economic applications*, volume III. North Holland, Amsterdam, 2002.

Michael H Riordan. Contracting in an idiosyncratic market. *Bell Journal of Economics*, 14:338–350, 1983.

Michael H Riordan. Uncertainty, asymmetric information and bilateral contracts. *Review of Economic Studies*, 51(1):83–93, 1984.

Leonard J Savage. Elicitation of personal probabilities and expectations. In A. Zellner E.S. Fienberg, editor, *Studies in Bayesian Econometrics and Statistics*. North Holland, Amsterdam, 1974.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of finance*, 16(1):8–37, 1961.

Mark Walker. On the impossibility of a dominant-strategy mechanism to optimally decide public question. *Econometrica*, pages 1521–1540, 1980.