No. 375

Incentive Compatible Mechanisms with Full Transferability

by
Hitoshi Matsushima

July 1988

Incentive Compatible Mechanisms
With Full Transferability

Hitoshi Matsushima

Institute of Socio-Economic Planning, University of Tsukuba
1-1-1, Tennodai, Tsukuba City, Ibaraki 305, Japan.

November 1987
Revised July, 1988.

Please address correspondence to: Hitoshi Matsushima
2-30-12, Sanno, Ohta-ku, Tokyo 143, Japan.

Abstract

In the n-person Bayesian collective choice problem with full transferability, it is shown that, if there are at least three players in the society and a certain regularity condition holds, there exists a Bayesian incentive compatible mechanism that is both efficient and budget-balancing.

1. Introduction

Many authors have considered the problem of implementing an optimal provision of public goods under incomplete information about players' preferences, where the problem has been viewed as constructing mechanisms with full transferability that induce players to reveal their respective true types in equilibrium. Groves intoduced in [6], so-called, Groves mechanisms, in which the truth-telling behaviours are dominant strategies, yielding an efficient production of public goods. The class of Groves mechanisms is the one of all dominant strategy mechanisms with efficient production, but, unfortunately, will not balance the budget. Hence, we can say that, in general, there is no dominant strategy mechanism that is both efficient and budget-balancing.

D'Aspremont and Gerard-Varet [1] considered Bayesian incentive compatible mechanisms in which the truth-telling behaviours are optimal Bayesian strategies. They constructed a Bayesian incentive compatible mechanism that is both efficient and budget-balancing, under the independence condition that the beliefs of all players are common knowledge. Moreover, d'Aspremont and Gerard-Varet [2,3] and Cremer and Riordan [4] tried to weaken this condition. The conditions introduced by the above authors, however, seem too restrictive in terms of information to be the representative of the general case.

The purpose of this paper is to show that a possibility result will hold in general even though the independence condition is removed; that is, it is shown in Theorem 1 that there exists a Bayesian incentive compatible mechanism which is both efficient and budget-balancing under a certain regularity condition. The regularity condition is very weak and will be typically satisfied when there are at least three players in the society.

When we construct mechanisms in the proof of Theorem 1, we require that all players, except player 1, play the truth-telling behaviours as dominant strategies. Cremer and Riordan [4] have pointed out that this treatment has several advantages. We will supplement their argument with another advantage that this treatment resolves the following drawback of the Bayesian approach to the collective choice problem; that is, in Bayesian incentive compatible mechanisms, telling the truth is not necessarily a unique Bayesian equilibrium. We will show in Theorem 2 that in mechanisms constructed in Theorem 1, telling the truth will be a unique Bayesian equilibrium.

2. Model Setting

We consider the following collective choice problem, which is the same as the one developed in d'Aspremont and Gerard-Varet [1,2,3], Myerson [7], Cremer and Riordan [4], and so on. $N = \{1,\ldots,n\}$ is the set of players. We assume:

A.1. There exist at least three players, i.e., $n \geq 3$.

Players have to choose amongst the set of all alternative public productions, X. We introduce a commodity called money in order to allow any kind of transfer amongst all players. A transfer is denoted by an element $t = (t_1,\ldots,t_n)$ of R^n, where t_i is the transfer-payment to player i. Thus, an outcome will be an element of (X,R^n). A_i is the set of possible types for player i. To avoid mathematical complication, we shall assume that A_i is a nonempty finite set. We denote $A := \times_{i\varepsilon N} A_i$. Player i has a von Neumann-Morgenstern utility function $U_i: X \times R \times A_i \rightarrow R$, where $U_i(x,t_i,a_i)$ is the payoff obtained by player i when (x,t) is chosen and a_i is the true type of player i. We assume:

A.2. For every $i\varepsilon N$, there exists a function u_i from $X \times A_i$ to R such that for every $a_i \varepsilon A_i$, every $x \varepsilon X$ and every $t_i \varepsilon R$,

$$U_i(x,t_i,a_i) = u_i(x,a_i) + t_i.$$

A.2 will amount to admit <u>unrestricted side-payments with full transferability</u>.

Player i knows, a priori, his type a_i, but not the other players' types a_{-i}. The belief of player i concerning the other players' types is represented by a family $P^i = \{p^i(.|a_i), a_i \varepsilon A_i\}$, where $p^i(.|a_i)$ is a probability over $A_{-i} := \underset{j\varepsilon N/(i)}{\times} A_j$. $\{P^i, i\varepsilon N\}$ is assumed to be common knowledge.

Players agree to delegate the choice of a particular outcomes to some central planning board according to some well-specified rules. The type a_i of player i is, a priori, unknown to the central planning board, but, will have to be taken into account by these rules in some way. Specifically, we assume that player i has to publicly announce some message $m_i \varepsilon A_i$ as being his own type. A <u>decision rule</u> is a function $g:A \rightarrow X$, and a <u>transfer rule</u> is a function $s = (s_i)_{i\varepsilon N}:A \rightarrow R^n$. $g(m)$ and $s_i(m)$ are the alternative and the tranfer-payment to player i respectively when the messages are $m := (m_i)_{i\varepsilon N}$. We shall confine attentions to transfer rules that are <u>budget-balancing</u>: A transfer rule s is <u>budget-balancing</u> if and only if for every $m\varepsilon A$,

$$\sum_{i\varepsilon N} s_i(m) = 0.$$

A pair of decision rule and a transfer rule, (g,s), is called a <u>mechanism</u>. Each player is supposed to know the mechanism ruling the behaviour of the central planning board.

3. Bayesian Incentive Compatibility

In a mechanism, players are assumed to simultaneously announce some messages. A message of player i is regarded as a strategy for player i. A strategy rule for player i is a function σ_i from A_i into A_i. $\sigma_i(a_i)$ is the message which player i announces when his true type is a_i. $\sigma := (\sigma_i)_{i \varepsilon N}$ denotes a strategy rule. Given a mechanism (g,s) and a strategy rule σ, the conditionally-expected payoff for player i given that his type is a_i is

$$V_i(g,s,\sigma|a_i) := \sum_{a_{-i} \varepsilon A_{-i}} \{u_i(g(\sigma(a)),a_i) + s_i(\sigma(a))\} p^i(a_{-i}|a_i).$$

A strategy rule σ is said to be a Bayesian equilibrium for a mechanism (g,s) if for every $i \varepsilon N$ and every $a_i \varepsilon A_i$,

$V_i(g,s,\sigma|a_i) \geq V_i(g,s,(\sigma_i', \sigma_{-i})|a_i)$ for all strategy rules for player i, σ_i'.

Let σ_i^* be the honest strategy rule for player i, such that for every $a_i \varepsilon A$,

$$\sigma_i^*(a_i) = a_i.$$

A mechanism (g,s) is said to be Bayesian incentive compatible if the honest strategy rule $\sigma^* = (\sigma_i^*)_{i \varepsilon N}$ is a Bayesian equilibrium for (g,s).

4. Results

g^* is a decision rule that is efficient; that is, for every $a \varepsilon A$,

$$\sum_{i \varepsilon N} u_i(g(a), a_i) \geq \sum_{i \varepsilon N} u_i(x, a_i) \text{ for all } x \varepsilon X.$$

Our purpose is to find a transfer rule $\hat{s}$ with budget-balancing such that $(g^*, \hat{s})$ is Bayesian incentive compatible.

We shall introduce the following regularity condition: Let

$$A_{-1-n} := A_2 \times \ldots \times A_{n-1},$$

where $a_{-1-n} = (a_2, \ldots, a_{n-1}) \varepsilon A_{-1-n}$. Let $m := |A_{-1-n}|$. We denote

$$A_{-1-n} = \{a^1_{-1-n}, \ldots, a^m_{-1-n}\}.$$

Moreover, let $w := |A_1|$, and we denote

$$A_1 = \{a^1_1, \ldots, a^w_1\}.$$

$p^1_{-n}(.|a_1)$ is a probability over A_{-1-n} induced by $p^1(.|a_1)$; that is, for every $a_{-n} \varepsilon A_{-n}$,

$$p^1_{-n}(a_{-1-n}|a_1) := \sum_{a_n \varepsilon A_n} p^1(a_{-1}|a_1),$$

where $A_{-n} := A_1 \times \ldots \times A_{n-1}$. B is a $[w \times m]$ matrix, such that the (r,k) component is $p^1_{-n}(a^k_{-1-n}|a^r_1)$.

C.1. rank B = w.

C.1 will hold in general, provided that A.1 holds and $w \leq m$. Under C.1, we can easily check from full utility-transferability that for any function h from A_1 into R, there exists a function $\theta: A_{-1-n} \to R$ that solves the following simultaneous equations: for each $a_1 \varepsilon A_1$,

$$\sum_{a_{-1-n} \varepsilon A_{-1-n}} \theta(a_{-1-n}) p_{-n}^{1}(a_{-1-n} | a_1) = h(a_1).$$

Define a transfer rule s^* in the following way: for every $i \varepsilon N$ and every $a \varepsilon A$,

$$s_i^*(a) = \sum_{j \varepsilon N/\{i\}} u_j(g^*(a), a_j).$$

The mechanism (g^*, s^*) is, so-called, a Groves mechanism, in which the truth-telling behaviours are dominant strategies, i.e., for every $i \varepsilon N$ and every $a \varepsilon A$,

$$u_i(g^*(m), a_i) + s_i^*(a) \geq u_i(g^*(m_i, a_{-i}), a_i) + s_i^*(m_i, a_{-i}) \text{ for all } m_i \varepsilon A_i.$$

Thus, (g^*, s^*) is Bayesian incentive compatible. However, s^* will not be budget-balancing, as Groves pointed out in [6]. In the following theorem, we will show that under C.1, we can construct a transfer rule will budget-balancing based on the Groves mechanism s^*.

Theorem 1. Under C.1, there exists a transfer rule $\hat{s}$ will budget-balancing such that $(g^*, \hat{s})$ is Bayesian incentive compatible.

<u>Proof</u>. We define a function $\beta: A_1 \times A_1 \to R$ as follows: for every $m_1 \varepsilon A_1$ and every $a_1 \varepsilon A_1$,

$$\beta(m_1, a_1) := \sum_{a_{-1} \varepsilon A_{-1}} \sum_{i \varepsilon N} s_i^*(m_1, a_{-1}) p^1(a_{-1} | a_1).$$

Moreover, we define a function $\mu: A_{-n} \to R$ to satisfy that for every $m_1 \varepsilon A_1$ and every $a \varepsilon A$,

$$\sum_{a_{-1-n} \varepsilon A_{-1-n}} \mu(m_1, a_{-1-n}) p^1_{-n}(a_{-1-n} | a_1) = \beta(m_1, a_1).$$

C.1, together with full utility-transferability, guarantees the existence of such μ. We construct a transfer rule $\hat{s}$ in the following way: for every $m \varepsilon A$,

$$\hat{s}_1(m) := \mu(m_{-n}) - \sum_{i \varepsilon N/\{1\}} s_i^*(m),$$

$$\hat{s}_i(m) := s_i^*(m) \text{ for all } i \varepsilon N/\{1,n\},$$

and

$$\hat{s}_n(m) := s_n^*(m) - \mu(m_{-n}).$$

Notice that $\hat{s}$ is budget-balancing. From the definition of μ, it is easy to check that for every $a_1 \varepsilon A_1$ and every strategy rule σ_1 for player 1,

$$V_1(g^*, \hat{s}, (\sigma_1, \sigma^*_{-1}) | a_1) = V_1(g^*, s^*, (\sigma_1, \sigma^*_{-1}) | a_1).$$

This implies that σ_1^* is the best response in the Bayesian sense in the mechanism $(g^*, \hat{s})$ whenever the others conform σ^*_{-1}. It is easy to check that for every $i \varepsilon N/\{1\}$ and every $a_i \varepsilon A_i$, $\sigma_i^*(a_i)$ is the dominant strategy for player i in the mechanism $(g^*,\hat{s})$. Therefore, σ^* is a Bayesian equilibrium for $(g^*,\hat{s})$.

Q.E.D.

It is essential in the proof to make μ independent of player n's message m_n. This guarantees that telling the truth is the dominant strategy for player n irrespective of μ. We can not apply the idea to two-player cases, because the independence of m_n means that μ depends on m_1 only: We may find such μ only if $P^i(.|a_i)$ is independent of a_i. However, such independency is too restrictive to be the representative of general cases (see D'Aspremont and Gérard-Varet [1,2,3] and Crémer and Riordan [4]).

In the mechanism constructed in Theorem 1, all players except player 1 play the truth-telling behaviours as dominant strategies. Such mechanism have, as pointed out by Crémer and Riordan [4], the following advantage: Telling the truth is always a Bayesian equilibrium for the mechanism irrespective of specification of P^i for all $i \in N/\{1\}$, or to put it in other words, P^i, $i \in N/\{1\}$, need not be common knowledge.

The drawback of the concept of Bayesian incentive compatibility is that the honest strategy rule σ^* is not necessarily a unique Bayesian equilibrium for the mechanism. However, we can show that σ^* will typically be a unique Bayesian equilibrium for the mechanism $(g^*,\hat{s})$: Additionally, we assume the following strict inequalities:

A.3. For every $a \varepsilon A$ and every $i \varepsilon N$,

$$\sum_{i \varepsilon N} u_i(g^*(a), a_i) > \sum_{i \varepsilon N} u_i(g^*(m_i, a_{-i}), a_i) \text{ for all } m_i \varepsilon A_i / \{a_i\}.$$

Theorem 2. σ^* is a unique Bayesian equilibrium for the mechanism $(g^*, \hat{s})$, where $\hat{s}$ is the transfer rule constructed in the proof of Theorem 1.

Proof. We show that there is no Bayesian equilibrium for $(g^*, \hat{s})$ constructed in Theorem 1 that is not the honest stratey rule. Notice from A.3 that, for every $i \varepsilon N/\{1\}$ and every $a_i \varepsilon A_i$, $\sigma_i^*(a_i) = a_i$ is the dominant strategy for player i with strict inequalities: for every $a \varepsilon A$,

$$u_i(g^*(a), a_i) + \hat{s}_i(a) > u_i(g^*(m_i, a_{-i}), a_i) + \hat{s}_i(m_i, a_{-i}) \text{ for all}$$
$$m_i \varepsilon A_i / \{a_i\}.$$

Thus, for every $i \varepsilon N/\{1\}$, σ_i^* is a unique optimal Bayesian-Nash strategy rule for player i irrespective of the other players' strategy rules, i.e., for every $i \varepsilon N/\{1\}$ and every $a_i \varepsilon A_i$,

$$V_i(g^*, \hat{s}, (\sigma_i^*, \sigma_{-i}) | a_i) > V_i(g^*, \hat{s}, \sigma | a_i) \text{ whenever } \sigma_i(a_i) \neq a_i.$$

Thus, if σ is a Bayesian equilibrium for $(g^*, \hat{s})$, then $\sigma_i = \sigma_i^*$ for all $i \varepsilon N/\{1\}$. Moreover, notice from A.3 that for every $a_1 \varepsilon A_1$,

$$V_1(g^*, \hat{s}, \sigma^* | a_1) > V_i(g^*, \hat{s}, (\sigma_1, \sigma_{-1}^*) | a_1) \text{ whenever } \sigma_1(a_1) \neq a_1.$$

Therefore, σ^* is a unique Bayesian equilibrium for $(g^*,\hat{s})$.

Q.E.D.

Throughout the argument presented above, we have assumed that each player's preference is described completely by his own type. When players' preferences are allowed to depend on the other players' types, $(g^*,\hat{s})$ will not be Bayesian incentive compatible. I'm now preparing the paper concerning the same problem in the general model-setting.

References

1 C. d'Aspremont and L.-A. Gerard-Varet, Incentive and Incomplete Information, Journal of Public Economics 11 (1979), 25-45.

2 ---, On Bayesian Incentive Compatible Mechanisms, in "Aggregation and Revelation of Preferences," (J.-J. Laffont, Ed.), 1979, pp. 289-308.

3 ---, Bayesian Incentive Compatible Beliefs, Journal of Mathematical Economics 10 (1982), 83-103.

4 J. Cremer and M. H. Riordan, A Sequential Solution to the Public Goods Problem, Econometrica 53 (1985), 77-84.

5 J. Green and J.-J. Laffont, "Incentive in Public Decision Making," Amsterdam, North-Holland, 1979.

6 T. Groves, Incentive in Teams, Econometrica 41 (1973), 617-663.

7 R. B. Myerson, Incentive Compatibility and the Bargaining Problem, Econometrica 47 (1979), 61-73.