CIRJE-F-149

Direct Mechanisms, Virtual Implementation, and Majority-Proofness

Hitoshi Matsushima
The University of Tokyo

April 2002

Direct Mechanisms, Virtual Implementation, and Majority-Proofness

Hitoshi Matsushima*

Faculty of Economics, University of Tokyo

March 20, 2002
This Version: April 20, 2002

Abstract

We investigate implementation in the complete information environments, where a social choice function is defined as a mapping from states to lotteries, and there exist four or more agents. We assume that for every agent, any pair of distinct states induces distinct strict preference orderings over all pure alternatives. In contrast to the previous works, we construct only direct mechanisms. Without any help of mechanism complexity, we can show that every social choice function is virtually implementable, provided that the set of states is restricted in ways that there always exist a majority of agents who dislike a particular agent's dictatorial choice the worst.

Keywords: Direct Mechanisms, Pure Strategy Nash Equilibrium, Virtual Implementation, Majority-Proofness, Domain Restrictions.

Journal of Economic Literature Classification Numbers: C72, D71, D78, H41.

---

* Faculty of Economics, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113, Japan. E-mail: hitoshi@e.u-tokyo.ac.jp

1. Introduction

We investigate implementation of a social choice function where there exist four or more agents. A social choice function is defined as a mapping from states to lotteries. We confine our attention to the complete information environments, in which all agents are assumed to know what the true state is. We restrict the set of possible states in ways that for every agent, any pair of distinct states induces distinct preference orderings, and that for every agent, any state induces a strict preference ordering over all pure alternatives. We assume that with the restrictions above, the set of states is inclusive enough.

We consider only *direct* mechanisms where each agent makes a single announcement about the state. We use pure strategy Nash equilibrium as the solution concept. We require a direct mechanism to virtually implement the social choice function in that at every state, truth telling is a pure strategy Nash equilibrium, and every pure strategy Nash equilibrium virtually induces the allocation that is suggested by the social choice function. Here, the definition of the virtualness above is originated by Matsushima (1988), and well cultivated by Abreu and Sen (1990). The purposes of the paper are to provide a sufficient condition for virtual implementation via direct mechanisms, and to argue that the class of social choice functions that are virtually implementable via direct mechanisms is large.

A social choice function is said to be *majority-proof* if at every state, there exist a majority of agents who never prefer a particular agent's dictatorial choice to the choice that is suggested by this social choice function. We show that if a social choice function is majority-proof, then it is virtually implementable via direct mechanisms. Based on this sufficiency result, we show as the main result of the paper that *every* social choice function is virtually implementable via direct mechanisms, provided that the set of states is restricted in ways that at every possible state, there exist a majority of agents who dislike a particular agent's dictatorial choice the worst.

Several earlier works in the implementation literature commonly indicate that it might be inevitable to make mechanisms *complicated* in order to implement a wide variety of social choice functions. See Moore (1992) and its references.[1] These works have constructed their respective mechanisms that are complicated in the sense that each agent has redundant, *slack* messages that she never announces as long as playing equilibrium behavior.[2] By announcing slack messages, each agent could deviate from any unwanted message profile. In contrast, direct mechanisms are *simple* in that they have *no* slack messages.

Several authors such as Maskin (1999) have constructed so-called 'modulo' mechanisms, in which each agent announces not only all agents' preferences but also an element of a large enough subset of integers. By combining the modulo mechanism design with the virtualness, Matsushima (1988), and Abreu and Sen (1990), showed that every social choice function is virtually implementable in pure strategy Nash equilibrium.[3]

Modulo mechanisms are complicated because all messages with the announcements of integers except '0' are slack. Since such complexity is a serious obstacle to implement a social choice function for practical reasons, we must need an alternative possibility result by using only simple mechanisms such as direct mechanisms. The present paper implies that in a class of environments, the idea of virtualness *alone* is enough for eliminating all unwanted pure strategy equilibria, and therefore, we need to ask *no* help to the use of slack messages.

The organization of the paper is as follows. Section 2 provides the model. Section 3 provides the sufficiency theorem and its corollaries. Section 4 provides the complete proof of the sufficiency theorem. Section 5 concludes.

---

[1] Most papers in this literature assumed that a state is defined as a profile of agents' preferences. Hence, the domain of a social choice function is restricted in ways that any pair of distinct states induces distinct preference profiles.

[2] Jackson (1992) discussed that it might be plausible for a mechanism implementing a social choice function to be bounded in some sense. Our simplicity is more restrictive than the boundedness, because a bounded mechanism may have slack messages.

[3] In the modulo mechanisms constructed by Matsushima (1988), every agent never announces the other agents' preferences except her two neighbors' preferences.

2. The Model

Let $N = \{1,...,n\}$, Ω, A and Δ denote the finite set of agents, the finite set of states, the finite set of pure alternatives and the set of lotteries over A, respectively, where we assume

$$n \geq 4 .$$

Agent $i's$ preference over lotteries is given by $u_i : \Delta \times \Omega \to R$ and satisfies the expected utility hypothesis. A *social choice function* is defined by a mapping from states to lotteries $f : \Omega \to \Delta$, where its range is denoted by $\Delta(f)$. When $f(\omega)(a) = 1$, we will simply denote $f(\omega) = a$.

We provide two assumptions on the profile of utility functions $(u_i)_{i\in N}$ as follows.

Assumption 1: For every $i \in N$, every $\omega \in \Omega$, and every $\omega' \in \Omega / \{\omega\}$, there exist $\alpha \in \Delta$ and $\alpha' \in \Delta / \{\alpha\}$ such that

$$u_i(\alpha,\omega) > u_i(\alpha',\omega) \text{ and } u_i(\alpha,\omega') < u_i(\alpha',\omega') .$$

Assumption 1 implies that *for every agent* $i \in N$, *any pair of distinct states induces distinct preference orderings over* Δ. Hence, by knowing a single agent's preference ordering, the central planner can know all *other* agents' preference orderings.

Assumption 2: For every $i \in N$, and every $\omega \in \Omega$,

$$u_i(a,\omega) \neq u_i(a',\omega) \text{ for all } a \in A \text{ and all } a' \in A / \{a\} .$$

Assumption 2 implies that for every agent $i \in N$, every state induces *strict* preference ordering over all pure alternatives.

We define a *direct mechanism* by (M, g) where $M_i = \Omega$ is the set of agent $i's$ messages, $M = \prod_{i\in N} M_i$, and $g : M \to \Delta$. When all agents announce the message profile $m \in M$, the central planner chooses the lottery $g(m)$. The *honest message rule* for each agent i is defined by $\mu_i : \Omega \to \Omega$ where

$$\mu_i(\omega) = \omega \text{ for all } \omega \in \Omega .$$

We denote $\mu = (\mu_i)_{i\in N}$ and $\mu(\omega) = (\mu_i(\omega))_{i\in N}$. A direct mechanism (M, g) will be simply denoted by g. For every $\omega \in \Omega$, (g,ω) defines a *game*. A message profile $m \in M$ is said to be a (pure strategy) *Nash equilibrium* in (g,ω) if

$$u_i(g(m),\omega) \geq u_i(g(m / m_i'),\omega) \text{ for all } i \in N \text{ and all } m_i' \in M_i .$$

A social choice function f is said to be *virtually implementable* if for every $\varepsilon \in (0,1)$, there exists a direct mechanism g^{ε} such that for every $\omega \in \Omega$, there exists a Nash equilibrium in $(g^{\varepsilon}, \omega)$, and for every Nash equilibrium m in $(g^{\varepsilon}, \omega)$, $g^{\varepsilon}(m)$ is $\varepsilon-$*close to* $f(\omega)$ in that

$$\left| g^{\varepsilon}(m)(a) - f(\omega)(a) \right| \leq \varepsilon \quad \text{for all} \quad a \in A.$$

We define *agent 1's dictatorial social choice function* $d : \Omega \to A$ by

$$u_1(d(\omega), \omega) \geq u_1(a, \omega) \quad \text{for all} \quad \omega \in \Omega \quad \text{and all} \quad a \in A.$$

At every state $\omega \in \Omega$, agent 1 prefers $d(\omega)$ the best among A. The following assumption on d is likely to hold when the set of states is *inclusive* enough.

Assumption 3: For every $\omega \in \Omega$, there exist $\omega' \in \Omega / \{\omega\}$ and $\omega'' \in \Omega / \{\omega, \omega'\}$ such that

$$d(\omega) = d(\omega') = d(\omega'').$$

Assumption 3 implies that for every $\omega \in \Omega$, there exist at least three distinct states at which agent 1 prefers $d(\omega)$ the best.

3. The Results

A social choice function f is said to be *majority-proof* if for every $\omega \in \Omega$,

$$\#\{i \in N/\{1\} \mid u_i(f(\omega),\omega) \geq u_i(d(\omega),\omega)\} > \frac{n}{2}.$$

Hence, at every state $\omega \in \Omega$, there exist a *majority* of agents who never prefer agent 1's dictatorial choice $d(\omega)$ to the socially desired choice $f(\omega)$. The following theorem states that majority-proofness is sufficient for virtual implementation via direct mechanisms.

Theorem 1: *A social choice function f is virtually implementable if it is majority-proof.*

In order to prove Theorem 1, we will construct the following direct mechanism. Fix a message profile $m \in M$ arbitrarily. With a small probability, each agent $i \in N$ is randomly chosen as a constrained dictator and the central planner chooses $l_i(m_i)$. Otherwise, the central planner chooses as follows. Fix two distinct states $\tilde{\omega} \in \Omega$ and $\hat{\omega} \in \Omega/\{\tilde{\omega}\}$ arbitrarily. The central planner chooses $f(\tilde{\omega})$ if a majority of agents except agent 1 announce $\tilde{\omega}$. She chooses $f(\tilde{\omega})$ if all agents announce either $\tilde{\omega}$ or $\hat{\omega}$, agent 1 announces $\hat{\omega}$, and the number of agents announcing $\tilde{\omega}$ is almost as large as the number of agents announcing $\hat{\omega}$. She chooses $d(m_1)$ if there exist three agents who announce distinct opinions.

Let $\omega \in \Omega$ denote the correct state. It is straightforward from the specification above that the honest message profile $\mu(\omega)$ is a strict Nash equilibrium and virtually induces $f(\omega)$. Whenever a majority of agents except agent 1 announce the same but incorrect state, there exists an agent who does not tell the truth but has incentive to tell the truth. From Assumption 3, it follows that whenever there exists no opinion about the state that a majority of agents except agent 1 commonly announce, agent 1 can make the outcome virtually equal $d(\omega)$ by changing her message approximately. From the majority-proofness, it follows that whenever the central planner chooses $d(\omega)$, a majority of agents have incentive to announce ω, and therefore, the message profile will be eventually switched into a message profile that virtually induces $f(\omega)$. Based on the outline above, we can prove that a majority-proof social choice function is virtually implementable. The full proof of Theorem 1 will be provided in the next section.

We provide the main result of the paper as a corollary of Theorem 1, which provides a restriction on the profile of utility functions $(u_i)_{i\in N}$ that guarantees every social choice function to be virtually implementable.

Corollary 2: *Suppose that for every* $\omega \in \Omega$,

$$\#\{i \in N/\{1\} \mid u_i(a,\omega) \geq u_i(d(\omega),\omega) \textit{ for all } a \in A\} > \frac{n}{2}.$$

Then, every social choice function is virtually implementable.

Proof: From the supposition, it follows that every social choice function is majority-proof. Theorem 1 implies that it is virtually implementable.

Q.E.D.

Corollary 2 states that *every* social choice function is virtually implementable, provided that the set of states is restricted in ways that at every possible state, there exist a majority of agents who dislike agent 1's dictatorial choice *the worst*.

Fix a social choice function f arbitrarily. We specify another social choice function f^* in ways that for every $\omega \in \Omega$,

$$f^*(\omega) = f(\omega) \quad \text{if} \quad \#\{i \in N/\{1\} \mid u_i(f(\omega),\omega) \geq u_i(d(\omega),\omega)\} > \frac{n}{2},$$

and

$$f^*(\omega) = d(\omega) \quad \text{otherwise.}$$

The definition of f^* implies that at every state ω, if a majority of agents prefer $d(\omega)$ to $f(\omega)$, then $f(\omega)$ must be replaced with $d(\omega)$. The following Corollary states that this modified social choice function f^* is virtually implementable.

Corollary 3: *For every social choice function* f, *its modified social choice function* f^* *is virtually implementable.*

Proof: It follows from the definition of f^* that f^* is majority-proof. Theorem 1 implies that it is virtually implementable.

Q.E.D.

4. Proof of Theorem 1

For every $\omega \in \Omega$, and every $k \in \{1,...,|A|\}$, let $\gamma_i(\omega,k) \in A$ denote the pure alternative that agent i prefers in the k-th place among A. We define $l_i : \Omega \to \Delta$ by

$$l_i(\omega)(\gamma_i(\omega,k)) \equiv \frac{|A|-k+1}{\sum_{\rho=1}^{|A|}\rho} \text{ for all } \omega \in \Omega \text{ and all } k \in \{1,...,|A|\}.$$

Since $l_i(\omega)(\gamma_i(\omega,k))$ is decreasing with respect to $k \in \{1,...,|A|\}$, it follows from Assumptions 1 and 2 that for every $\omega \in \Omega$, $\alpha = l_i(\omega)$ is the unique maximizer of $u_i(\alpha,\omega)$ among the range of l_i, i.e.,

$$u_i(l_i(\omega),\omega) > u_i(l_i(\omega'),\omega) \text{ for all } \omega \in \Omega \text{ and all } \omega' \in \Omega/\{\omega\}.$$

For every $m \in M$, and every $\omega \in \Omega$, let $n(m,\omega) \in \{0,...,n\}$ denote the number of agents $i \in N/\{1\}$ who announce $m_i = \omega$, i.e,

$$n(m,\omega) = \#\{i \in N/\{1\} \mid m_i = \omega\}.$$

We specify $x : M \to \Delta$ in ways that for every $m \in M$, and every $\omega \in \Omega$,

$$x(m) = f(\omega) \text{ whenever } n(m,\omega) > \frac{n}{2},$$

$$x(m) = f(\omega) \text{ whenever } \frac{n-1}{2} \le n(m,\omega) \le \frac{n}{2} \text{ and there exists } \omega' \in \Omega/\{\omega\} \text{ such that } m_i \in \{\omega,\omega'\} \text{ for all } i \in N, \text{ and either } [n(m,\omega) > n(m,\omega')] \text{ or } [n(m,\omega) = n(m,\omega') \text{ and } m_1 = \omega'],$$

and

$$x(m) = d(m_1) \text{ otherwise.}$$

For every $\varepsilon \in (0,1)$, we specify g^ε in ways that for every $m \in M$,

$$g^\varepsilon(m) = (1-\varepsilon)x(m) + \frac{\varepsilon}{n}\sum_{i \in N} l_i(m_i).$$

We will prove Theorem 1 by showing that for every $\varepsilon \in (0,1)$, and every $\omega \in \Omega$, the honest message profile $\mu(\omega)$ is a Nash equilibrium in (g^ε,ω), and that for every $\varepsilon \in (0,1)$ that is close enough to zero, every $\omega \in \Omega$, and every Nash equilibrium m in (g^ε,ω), it follows that

$$x(m) = f(\omega),$$

and therefore, $g^\varepsilon(m)$ is $\varepsilon-$close to $f(\omega)$.

Fix $\omega \in \Omega$ arbitrarily. Note that $x(\mu(\omega)) = f(\omega)$. Since $n \ge 3$, it follows that for every $i \in N$, and every $m_i \in M_i$,

$$x(\mu(\omega)/m_i) = f(\omega),$$

and therefore,

$$u_i(g^\varepsilon(\mu(\omega)/m_i),\omega) - u_i(g^\varepsilon(\mu(\omega)),\omega) = \frac{\varepsilon}{n}\{l_i(m_i) + l_i(\omega)\} \le 0.$$

Hence, $\mu(\omega)$ is a Nash equilibrium in (g^ε,ω).

We divide the set of message profiles M into the following three cases. Fix $m \in M$ arbitrarily.

Case 1: There exists $\tilde{\omega} \in \Omega$ such that

$$n(m,\tilde{\omega}) > \frac{n}{2}.$$

Case 2: There exist $\tilde{\omega} \in \Omega$ and $\hat{\omega} \in \Omega/\{\tilde{\omega}\}$ such that

$$m_i \in \{\tilde{\omega},\hat{\omega}\} \text{ for all } i \in N,$$

$$\frac{n-1}{2} \le n(m,\tilde{\omega}) \le \frac{n}{2},$$

and

either [$n(m,\tilde{\omega}) > n(m,\hat{\omega})$] or [$n(m,\tilde{\omega}) = n(m,\hat{\omega})$ and $m_1 = \hat{\omega}$].

Case 3: The message profile m does not belong to either of Case 1 and Case 2.

Fix $\varepsilon \in (0,1)$ arbitrarily, which is close enough to zero. Fix $m \in M/\{\mu(\omega)\}$ arbitrarily, where $x(m) \ne f(\omega)$. Assume that m is a Nash equilibrium in (g^ε,ω).

Suppose that m belongs to Case 1. Then, it follows that $\tilde{\omega} \ne \omega$, and every agent $i \in N$ who announces $m_i \ne \tilde{\omega}$ has strict incentive to announce $\mu_i(\omega)$ instead of m_i, because for every $m_i' \in M_i$,

$$x(m/\mu_i(\omega)) = x(m) = f(\tilde{\omega}),$$

and therefore,

$$u_i(g^\varepsilon(m/\mu_i(\omega)),\omega) - u_i(g^\varepsilon(m),\omega) = \frac{\varepsilon}{n}\{l_i(\omega) + l_i(m_i)\} > 0 \quad \text{if} \quad m_i \ne \omega.$$

Hence, it must hold that

$$m_i \in \{\tilde{\omega},\omega\} \text{ for all } i \in N.$$

Note that there exists an agent $i \in N$ such that $m_i = \tilde{\omega}$, $m/\mu_i(\omega)$ belongs to either Case 1 or Case 2, and

$$x(m/\mu_i(\omega)) = x(m) = f(\tilde{\omega}).$$

This agent has strict incentive to announce $\mu_i(\omega)$ instead of m_i, because

$$u_i(g^\varepsilon(m/\mu_i(\omega)),\omega) - u_i(g^\varepsilon(m),\omega) = \frac{\varepsilon}{n}\{l_i(\omega) + l_i(\tilde{\omega})\} > 0 .$$

However, this contradicts the Nash equilibrium property. Hence, it follows that m cannot belong to Case 1.

Suppose that m belongs to Case 2. Then, agent 1 has strict incentive to announce a message $m_1' \notin \{\omega, \tilde{\omega}\}$ satisfying $d(m_1') = d(\omega)$ instead of m_1, because m/m_1' belongs to Case 3,

$$x(m/m_1') = d(\omega) ,$$

and therefore,

$$u_1(g^\varepsilon(m/m_1'),\omega) - u_1(g^\varepsilon(m),\omega) = (1-\varepsilon)\{u_1(d(\omega),\omega) - u_1(f(\tilde{\omega}),\omega)\}$$
$$+\frac{\varepsilon}{n}\{u_1(l_1(m_1'),\omega) - u_1(l_1(m_1),\omega)\} > 0 ,$$

where Assumption 3 implies that such m_1' exists. this contradicts the Nash equilibrium property. Hence, it follows that m cannot belong to Case 2.

Suppose that m belongs to Case 3. Then,

$$n(m,\omega) \le \frac{n}{2} \text{ for all } \omega \in \Omega ,$$

and there exist $\tilde{\omega} \in \Omega/\{m_1\}$ and $\hat{\omega} \in \Omega/\{m_1, \tilde{\omega}\}$ such that

$$n(m,\omega) \ge 1 \text{ for all } \omega \in \{\tilde{\omega}, \hat{\omega}\} .$$

If $d(m_1) \ne d(\omega)$, then agent 1 has strict incentive to announce a message $m_1' \notin \{\tilde{\omega}, \hat{\omega}\}$ satisfying $d(m_1') = d(\omega)$ instead of m_1, because m/m_1' belongs to Case 3, $x(m/m_i') = d(\omega)$, and therefore, it follows from Assumption 2 and the fact that $\varepsilon \in (0,1)$ is close to zero that

$$u_1(g^\varepsilon(m/m_1'),\omega) - u_1(g^\varepsilon(m),\omega) = (1-\varepsilon)\{u_1(d(\omega),\omega) - u_1(d(m_1),\omega)\}$$
$$+\frac{\varepsilon}{n}\{u_1(l_1(m_1'),\omega) - u_1(l_1(m_1),\omega)\} > 0 .$$

However, this contradicts the Nash equilibrium property. Hence, it must hold that $d(m_1) = d(\omega)$. If $n(m,\omega) < \frac{n}{2}$, then it follows from inequality $n \ge 4$ that there exists an agent $i \in N$ such that $m_i \ne \omega$ and $m/\mu_i(\omega)$ belongs to Case 3. Such an agent has strict incentive to announce $\mu_i(\omega)$ instead of m_i, because

$$x(m/\mu_i(\omega)) = x(m) = d(\omega) ,$$

and therefore,

$$u_i(g^\varepsilon(m/\mu_i(\omega)),\omega) - u_i(g^\varepsilon(m),\omega) = \frac{\varepsilon}{n}\{u_i(l_i(\omega),\omega) - u_i(l_i(m_i),\omega)\} > 0 .$$

However, this contradicts the Nash equilibrium property. Hence, it must hold that $n(m,\omega)=\frac{n}{2}$. Since f is majority-proof, it follows that there exists an agent $i\in N/\{1\}$ who announces $m_i\neq\omega$ and never prefers $d(\omega)$ to $f(\omega)$. Such an agent has strict incentive to announce $\mu_i(\omega)$ instead of m_i, because $m/\mu_i(\omega)$ belongs to Case 1,

$$x(m/\mu_i(\omega))=f(\omega),$$

and therefore,

$$u_i(g^{\varepsilon}(m/\mu_i(\omega)),\omega)-u_i(g^{\varepsilon}(m),\omega)=(1-\varepsilon)\{u_i(f(\omega),\omega)-u_i(d(\omega),\omega)\}$$

$$+\frac{\varepsilon}{n}\{u_i(l_i(\omega),\omega)-u_i(l_i(m_i),\omega)\}>0.$$

However, this contradicts the Nash equilibrium property. Hence, it follows that m cannot belong to Case 3.

From the above arguments, we have proved that for every Nash equilibrium $m\in M$ in (g^{ε},ω), $x(m)=f(\omega)$, and therefore, $g^{\varepsilon}(m)$ is $\varepsilon-$close to $f(\omega)$. Hence, we have completed the proof of the Theorem.

Q.E.D.

5. Concluding Remarks

The present paper has investigated implementation of a social choice function as a mapping from states to lotteries. We have shown that with some domain-restrictions, majority-proofness is sufficient for virtual implementation in pure strategy Nash equilibrium via direct mechanisms. We have shown also that whenever there always exist a majority of agents who dislike a particular agent’s dictatorial choice the worst, then every social choice function is virtually implementable in pure strategy Nash equilibrium via direct mechanisms.

The paper did not check whether there exists unwanted mixed strategy Nash equilibria. Abreu and Matsushima (1992) showed that every social choice function might be virtually implementable in mixed strategy Nash equilibrium, when we can construct so-called Abreu-Matsushima mechanisms, which are more complicated than direct mechanisms. In the Abreu-Matsushima mechanisms, each agent makes multiple announcements about the state, and the message profile that induces every agent to make multiple honest announcements is the unique mixed strategy Nash equilibrium.[4] This positive result crucially depends on the assumption that each agent can be individually levied by a small amount. In contrast, the preset paper does not depend on this assumption.

The recent works by the author such as Matsushima (2002a, 2002b) are closely related, because both considered direct mechanisms or its variant in the implementation literature. Matsushima (2002b) investigated a particular class of direct mechanisms named virtual plurality mechanisms. A virtual plurality mechanism is defined on the basis of a plausible decision making procedure in that with a high probability the central planner enforces the allocation that the social choice function assigns to the opinion announced by a largest number of agents. Matsushima (2002b) showed that with three or more agents and some domain-restrictions, a social choice function is virtually implementable via virtual plurality mechanisms if and only if it always enforces the Condorcet winner. In contrast, the present paper showed that when we take into account a wider class of direct mechanisms than virtual plurality mechanisms, much wider variety of social choice functions, including inefficient ones, are all virtually implementable.[5]

---

[4] More precisely, this message profile is the unique iteratively undominated strategy profile.
[5] However, the characterization of Matsushima (2002b) does not depend on the assumption that there exists four or more agents.

Moreover, Matsushima (2002a) investigated much simpler mechanisms than direct mechanisms, named local direct mechanisms, in which each agent makes only a single announcement about her own and two neighbors' preferences. It must be noted that a social choice function is unlikely to be virtually implementable in pure strategy Nash equilibrium when we consider local direct mechanisms instead of direct mechanisms. However, Matsushima (2002a) showed that with a very minor restriction, every social choice function could be virtually implementable, provided that every agent is assumed to be boundedly rational in a naïve sense that she may announce any best reply, including disequilibrium messages. This result does not depend on Assumptions 1, 2 and 3 in the present paper, but depends on the assumption that each agent can be individually levied by a small amount.

References

Abreu, D. and H. Matsushima (1992): "Virtual Implementation in Iteratively Undominated Strategies: Complete Information," *Econometrica* 60, 993-1008.

Abreu, D. and A. Sen (1991): "Virtual Implementation in Nash Equilibrium," *Econometrica* 59, 997-1021.

Jackson, M. (1992): "Implementation in Undominated.Strategies: A Look at Bounded Mechanisms," *Review of Economic Studies* 59, 757-75.

Maskin (1999): "Nash equilibrium and Welfare Optimality," *Review of Economic Studies* 66, 23-38

Matsushima, H. (1988): "A New Approach to the Implementation Problem," *Journal of Economic Theory* 45, 128-144.

Matsushima, H. (2002a): "Stability and Implementation via Simple Mechanisms in the Complete Information Environments," Discussion Paper CIRJE-F-147, Faculty of Economics, University of Tokyo.

Matsushima, H. (2002b): "Plurality Mechanisms, Virtual Implementation, and Condorcet Decisiveness," Discussion Paper CIRJE-F-148, Faculty of Economics, University of Tokyo.

Moore, J. (1992): "Implementation in Environments with Complete Information," in *Advances in Economic Theory: Sixth World Congress*, ed. by J.-J. Laffont. Cambridge University Press.