Coalitionally strategyproof functions depend only on the most-preferred alternatives.*

H. Reiju Mihara†
reiju@ec.kagawa-u.ac.jp
Economics, Kagawa University, Takamatsu, 760-8523, Japan

April, 1999
[Social Choice and Welfare (2000) 17: 393–402]

Abstract

In a framework allowing infinitely many individuals, I prove that coalitionally strategyproof social choice functions satisfy "tops only." That is, they depend only on which alternative each individual prefers the most, not on which alternative she prefers the second most, the third, ..., or the least. The functions are defined on the domain of profiles measurable with respect to a Boolean algebra of coalitions. The *unrestricted* domain of profiles is an example of such a domain. I also prove an extension theorem.

[*Journal of Economic Literature* Classifications: D71, C72, C71.]

[*Keywords:* Gibbard-Satterthwaite theorem, dominant strategy implementation, social choice functions, infinitely large societies, tops only.]

*Earlier versions of the paper were circulated under the title "Coalitionally strategyproof functions depend only on which alternative each individual prefers the most." [The running title is "Tops only."]

†I benefited from comments by Yoshikatsu Tatamitani and two anonymous referees. I also thank Kaori Hasegawa, Mitsunobu Miyake, and Dao-Zhi Zeng for valuable discussions. [Fax: +81-87-832-1820. Phone: +81-87-832-1831.]

1 Introduction

Pazner and Wesley [14] showed that when there are infinitely many individuals, there exists a nondictatorial, coalitionally strategyproof social choice function on the unrestricted domain of profiles of individual preferences. In their proof, using a non-constructive mathematical technique (the existence of a free ultrafilter over an arbitrary infinite set), they defined a nondictatorial function and showed that it is coalitionally strategyproof. In Mihara [11], I proved the same result for countably infinite societies; I did so by constructing a concrete example of a nondictatorial, coalitionally strategyproof social choice function. The functions that these authors defined satisfy the property (called "tops only") that only the most-preferred alternatives of individuals matter: the functions depend only on which alternative each individual prefers the most, not on which alternative she prefers the second most, the third, . . . , or the least.

While the property of "tops only" may be desirable in terms of informational simplicity, it may not be necessarily so in other aspects. For example, suppose that the set of individuals consists of all natural numbers and that the set of alternatives consists of three elements a, b, and c. Suppose that all even-numbered individuals prefer a to the other two, and that all odd-numbered individuals prefer b to the other two. In such a case, to choose one alternative from among the three, the planner might want to base her choice on each individual's worst alternative: if for example a is the least-preferred alternative of significantly many individuals but b is not, then she might want to choose b in this case.

The main goal of this paper is to see whether the "tops only" property is necessary for coalitional strategyproofness. The answer is yes. Proposition 1 establishes in a very general framework that when there are *infinitely* many individuals, every coalitionally strategyproof function satisfies "tops only." (Example 1 shows that *coalitional* strategyproofness in the proposition cannot be replaced by *individual* strategyproofness.) On the other hand, when there are only *finitely* many individuals, then a coalitionally (actually, individually) strategyproof function is necessarily dictatorial, according to the Gibbard-Satterthwaite theorem [7, 15]. Hence the property is trivially satisfied since such functions must choose the best alternative of the dictator.

The first proof of Proposition 1 goes as follows: Given are two preference profiles $\mathbf{p}$ and $\mathbf{p}'$ that are "top equivalent" (1). This means that for each alternative, the coalition that prefers it the most at $\mathbf{p}$ is the same as that at $\mathbf{p}'$. Partition the set of individuals into classes so that those in the same class have an identical preference at $\mathbf{p}$ and another identical preference at $\mathbf{p}'$. The partition is finite since the set of alternatives (hence the set of preferences) is finite. Treat each class as an individual and apply the Gibbard-Satterthwaite theorem to find a dictatorial class. What this class prefers the most at $\mathbf{p}$ is what it prefers the most at $\mathbf{p}'$.

In the context of provision of public goods, several authors show the proposition that *individually* strategyproof functions satisfy "tops only." They do so by *restricting* preferences to those satisfying a generalized "single-peakedness" condition or the related condition of "separability." For example, Barberà and Jackson [4, Theorem 1] show a variant of the proposition for the case of a public good, available at different levels expressed as a real number. Barberà, Sonnenschein, and Zhou [5, pp. 601–2] give another variant for the case of one or more public goods, each available at two levels ("accept" or "reject"). Finally, Barberà, Gul, and Stacchetti [3, Theorem 1] prove the proposition for one or more public goods, each available at a finite number of different levels. All these papers consider finite sets of individuals.

In contrast to those papers, this paper considers an abstract social choice setting and allows the *un*restricted domain of profiles. It also sticks to the *un*modified notion of coalitional strategyproofness. One reason for doing that is that there are some contexts in which neither restriction on profiles nor modifications (such as introducing "counter-threats" [13] by non-members of the manipulating coalition) of the notion make much sense. For instance, it often happens in the political arena that all alternatives for a voting are so similar that no voter outside can (bother to) predict what coalition, if any, is secretly forming. No preference can be ruled out in advance, and no "counter-threats" are likely to be realized there. Another reason is that it is natural to begin our investigation with a simple setting—the unrestricted domain and the unmodified notion of coalitional strategyproofness. (As I discussed above, however, investigations of restricted domains have preceded the present one.)

I investigate the case of *infinitely* many individuals. (Mihara [10, 11] gives an interpretation of an infinite society of individuals. See also Remark 2.) That is the only case in which the question of finding a coalitionally strategyproof function violating "tops only" is interesting, since I allow the unrestricted domain and stick to the unmodified notion of coalitional strategyproofness. The case of finitely many individuals is trivially answered by the Gibbard-Satterthwaite theorem [7, 15].

The functions in this paper are defined on the domain of profiles *mea-*

surable [1] with respect to a Boolean algebra. These are the profiles such that for each pair of alternatives x and y, the coalition of individuals that prefer x to y is "observable." (Which coalitions are "observable" depends on the choice of a Boolean algebra. If the Boolean algebra is chosen so that all coalitions are "observable," then I have, as a special case, the unrestricted domain.) But if a social choice function satisfies "tops only," why should one care about observability of all those coalitions, many of which are irrelevant? This suggests extending the domain. Applying the main result, I show (Proposition 2) that any coalitionally strategyproof social choice function on the domain of measurable profiles can be extended to a coalitionally strategyproof social choice function on the larger domain of "admissible" profiles [11].

2 Framework

The framework is adapted from that of Armstrong [1, 2] for Arrow's social welfare functions. (Armstrong's framework has found application in my studies [10, 9, 12] of aspects of information processing in social choice. It has also found application in studies [6] of "describable" social welfare functions.)

Let I be an infinite set of individuals. Let X be a *finite* set of alternatives, which has at least three elements. Let $\mathcal{S}$ be the set of (strict) preferences, i.e., total (if $x \neq y$, then either $x \succ y$ or $y \succ x$), asymmetric ($x \succ y$ implies that $y \succ x$ is false), and transitive binary relations $\succ$ on X. (For simplicity, indifference is not allowed.)

Let $\mathcal{B}$ be a *Boolean algebra* consisting of subsets of I. By definition, it satisfies the following: (i) $\emptyset$, $I \in \mathcal{B}$; (ii) $A \cup B$, $A \cap B$, $A^c \in \mathcal{B}$ if A, $B \in \mathcal{B}$ (where A^c denotes the complement of A). Intuitively, an element of the Boolean algebra is a coalition observable to the planner.

Remark 1. Under a different interpretation, the case of *finitely* many individuals can be incorporated into the present framework. For this purpose, consider a finite partition of I and consider the Boolean algebra consisting of all unions of elements in the partition, and reinterpret each partition element as an "individual." ‖

A *profile* is a list $\mathbf{p} = (\succ_i^{\mathbf{p}})_{i \in I} \in \mathcal{S}^I$ of individual preferences $\succ_i^{\mathbf{p}}$, for $i \in I$. A profile $(\succ_i^{\mathbf{p}})_{i \in I}$ is $\mathcal{B}$-*measurable* if $\{\, i \in I : x \succ_i^{\mathbf{p}} y \,\} \in \mathcal{B}$ for all x, $y \in X$. Denote by $\mathcal{S}_{\mathcal{B}}^I$ the set of all $\mathcal{B}$-measurable profiles. A $\mathcal{B}$-*social choice function* is a function g from $\mathcal{S}_{\mathcal{B}}^I$ *onto* X. (The assumption that it is onto is usually included in the statement of the Gibbard-Satterthwaite theorem, ruling out uninteresting strategyproof functions such as constant ones. The assumption will be used almost everywhere I resort to the Gibbard-Satterthwaite theorem. Remark 3 shows that the assumption is indispensable.) For a preference $\succ$, denote by $\max \succ$ the greatest (maximal) element of X with respect to $\succ$. A profile $(\succ_i^{\mathbf{p}})_{i \in I}$ is $\mathcal{B}$-*admissible* if $\{\, i \in I : x = \max \succ_i^{\mathbf{p}} \,\} \in \mathcal{B}$ for all $x \in X$. Denote by $\mathcal{S}^I(\mathcal{B})$ the set of all $\mathcal{B}$-admissible profiles. (Note that when X is finite, $\mathcal{S}_{\mathcal{B}}^I \subseteq \mathcal{S}^I(\mathcal{B})$.) A ($\mathcal{B}$)-*social choice function* is a function G from $\mathcal{S}^I(\mathcal{B})$ onto X. I say two profiles $\mathbf{p} = (\succ_i^{\mathbf{p}})_{i \in I}$ and $\mathbf{p}' = (\succ_i^{\mathbf{p}'})_{i \in I}$ are *top equivalent* ($\mathbf{p} \simeq \mathbf{p}'$) if

$$\{\, i \in I : x = \max \succ_i^{\mathbf{p}} \,\} = \{\, i \in I : x = \max \succ_i^{\mathbf{p}'} \,\} \tag{1}$$

for all $x \in X$.

Let g be a $\mathcal{B}$-social choice function. A coalition E *can manipulate* g *at* a profile $\mathbf{p} = (\succ_i^{\mathbf{p}})_{i \in I}$ *by reporting* another profile $\mathbf{p}' = (\succ_i^{\mathbf{p}'})_{i \in I}$ if $\succ_i^{\mathbf{p}'} = \succ_i^{\mathbf{p}}$ for all $i \notin E$ and $g(\mathbf{p}') \succ_i^{\mathbf{p}} g(\mathbf{p})$ for all $i \in E$. g is said to be *coalitionally manipulable* if for some $\mathcal{B}$-measurable profiles $\mathbf{p}$, $\mathbf{p}'$ and for some nonempty coalition $E \subseteq I$, E can manipulate g at $\mathbf{p}$ by reporting $\mathbf{p}'$. g is *coalitionally strategyproof* if it is not coalitionally manipulable. *Coalitionally strategyproof* ($\mathcal{B}$)-social choice functions are defined similarly; in this case, only admissibility is required of the profiles $\mathbf{p}$ and $\mathbf{p}'$.

Remark 2. For $\mathcal{B}$-social choice functions, restricting the manipulating coalitions to those in $\mathcal{B}$ will not alter the notion of strategyproofness: If E can manipulate g at $\mathbf{p}$ by reporting $\mathbf{p}'$, then

$$E' = \{\, i \in I : g(\mathbf{p}') \succ_i^{\mathbf{p}} g(\mathbf{p}) \,\} \in \mathcal{B}$$

can manipulate g at $\mathbf{p}$ by reporting $\mathbf{p}'$. ‖

g is said to be *individually manipulable* if for some $\mathcal{B}$-measurable profiles $\mathbf{p}$, $\mathbf{p}'$ and for some one-individual coalition $\{i\} \subseteq I$, $\{i\}$ can manipulate g at $\mathbf{p}$ by reporting $\mathbf{p}'$. g is *individually strategyproof* if it is not individually manipulable.

Let $\mathcal{Q} \subseteq \mathcal{S}_{\mathcal{B}}^I$. A coalition S *dictates the profiles in* $\mathcal{Q}$ if for all $\mathbf{p} \in \mathcal{Q}$ and for all $x \in X$, whenever x is most preferred by all individuals in S, $g(\mathbf{p}) = x$.

3 The main result

I begin with an example of an *individually* strategyproof social choice function (on the unrestricted domain) that violates “tops only.” Though the function is obviously nondictatorial, it is *not* coalitionally strategyproof.

Example 1. Let $I = \mathbf{N}$, the set of natural numbers. Let $X = \{a, b, c\}$, where a is indexed by 1, b by 2, and c by 3. For each $x \in X$ and each $\mathbf{p} \in \mathcal{S}^I$, partition the set of individuals into $r_1^{\mathbf{p}}(x)$, $r_2^{\mathbf{p}}(x)$, and $r_3^{\mathbf{p}}(x)$—the sets of individuals that prefer x the most, second most, and the least (at profile $\mathbf{p}$) respectively. Define a social choice function g as follows: (i) if there is exactly one x such that $r_1^{\mathbf{p}}(x)$ is infinite and $r_3^{\mathbf{p}}(x)$ is finite, then let $g(\mathbf{p}) = x$; (ii) otherwise, let $g(\mathbf{p})$ be the first-indexed x such that $r_1^{\mathbf{p}}(x)$ is infinite. (The case (i) means that the following condition is satisfied by only one x: infinitely many individuals rank x first, and only finitely many individuals rank x third. For example, if all but finitely many individuals rank x first, the condition is satisfied. If finitely many individuals rank x first and all the others rank x second, the condition is violated.)

It is easy to see that the function is individually strategyproof: If $\mathbf{p}$ and $\mathbf{p}'$ are different with respect to only one individual's preference, then $g(\mathbf{p}) = g(\mathbf{p}')$ since finiteness and infiniteness of the set of individuals that rank an alternative first (second, third) are unaffected. No individual can thus manipulate g by reporting $\mathbf{p}'$ at $\mathbf{p}$.

To see that the function violates "tops only," and that it is not coalitionally strategyproof, consider the following profile. All even-numbered individuals rank a first, b second, and c third; all odd-numbered individuals rank b first, a second, and c third. In this case, g assigns a to the profile (case (ii) above applies). Now suppose that all the odd-numbered individuals report that they rank b first, c second, and a third. (Note that the first-ranked alternative is the same as before.) Then, g assigns b to the reported profile (case (i) above applies). But b is the alternative that the odd-numbered individuals prefer the most. They can thereby manipulate the function. $\|$

On the other hand, if *coalitional* strategyproofness is required, no social choice function violates "tops only":

Proposition 1 *Suppose that $g\colon \mathcal{S}_{\mathcal{B}}^I \to X$ is a coalitionally strategyproof $\mathcal{B}$-social choice function. Then for any $\mathcal{B}$-measurable profiles $\mathbf{p}$ and $\mathbf{p}'$ that are top equivalent (i.e., $\mathbf{p} \simeq \mathbf{p}'$; see (1) above), we have $g(\mathbf{p}) = g(\mathbf{p}')$.*

Proof. Suppose that g, $\mathbf{p}$, and $\mathbf{p}'$ satisfy the assumptions. Partition the set I of individuals into coalitions $C_j = \{\, i : \succ_i^{\mathbf{p}} = \succ_j^{\mathbf{p}} \,\}$ (where j ranges over I) so that all individuals in the same coalition have the same preference at $\mathbf{p}$. Since X is finite, there are only finitely many possible preferences. Hence the partition $\{\, C_j : j \in I \,\}$ is finite. For profile $\mathbf{p}'$, partition the set I

into coalitions $C'_k = \{\, i \in I : \succ_i^{\mathbf{p}'} = \succ_k^{\mathbf{p}'} \,\}$ (where k ranges over I) in a similar way.

From these partitions, form the refinement

$$\{\, C_j \cap C'_k : j, k \in I \,\},$$

which is also a partition of I. Note that the individuals belonging to the same partition element $C_j \cap C'_k$ have an identical preference at $\mathbf{p}$. Likewise they have an identical preference at $\mathbf{p}'$.

Now, consider the class of profiles $\mathbf{q}$ such that for each j and k in I, all individuals within the same coalition $C_j \cap C'_k$ have an identical preference at $\mathbf{q}$. In particular, $\mathbf{p}$ and $\mathbf{p}'$ belong to this class of profiles. Regarding each of these (finitely many) coalitions as an "individual," we can obviously derive from g a social choice function for finitely many individuals, which becomes *individually* strategyproof. Then the Gibbard-Satterthwaite theorem for finitely many individuals implies that one of the coalitions—say, $C_{j'} \cap C'_{k'}$—dictates the profiles in the class.

The members of this dictatorial coalition $C_{j'} \cap C'_{k'}$ prefer a certain x the most at $\mathbf{p}$ (that is, if $C_{j'} \cap C'_{k'} \subseteq \{\, i \in I : x = \max \succ_i^{\mathbf{p}} \,\}$). We thus have $g(\mathbf{p}) = x$. But since $\mathbf{p}$ and $\mathbf{p}'$ are top equivalent, the same coalition members prefer x the most at $\mathbf{p}'$ (that is, $C_{j'} \cap C'_{k'} \subseteq \{\, i : x = \max \succ_i^{\mathbf{p}'} \,\}$), with the result that $g(\mathbf{p}') = x$. ∎

Remark 3. One might be tempted to modify the proof by considering the partition consisting of the coalitions of individuals having the same *maximal element*, instead of the same *preference*. Such an argument does not work since a coalition cannot be identified with an "individual" unless the members have the same preference. ‖

I now give an alternative proof, using Lemma 1 below. The lemma, which will be used again in an alternative proof of the extension theorem (Proposition 2) later, is of some interest in itself, as its corollary (Proposition 4) illustrates.

Before stating the lemma, I introduce two definitions: A *filter* $\mathcal{F}$ on the Boolean algebra $\mathcal{B}$ is a collection of sets in $\mathcal{B}$ satisfying: (i) $\emptyset \notin \mathcal{F}$; (ii) if $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$; (iii) if A, $B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$. We may think of a filter as a collection of "large" sets. An *ultrafilter* is a filter $\mathcal{U}$ that satisfies (iv): if $A \notin \mathcal{U}$, then $A^c \in \mathcal{U}$.

Lemma 1 *Suppose that* $g\colon \mathcal{S}^I_{\mathcal{B}} \to X$ *is a coalitionally strategyproof* $\mathcal{B}$*-social choice function. Let* W_g *be the set of all* winning *coalitions:* $S \in W_g$ *iff* $S \in \mathcal{B}$ *dictates the profiles in* $\mathcal{S}^I_{\mathcal{B}}$*. Then* W_g *is an ultrafilter on* $\mathcal{B}$*.*

Remark 4. Recall that the social choice function is onto by definition. It is easy to see that the assumption is necessary for the conclusion of the lemma. If g is not onto, then the set W_g of winning coalitions is empty. Then W_g cannot be an ultrafilter. ‖

The lemma extends Corollary 7.1 of Ishikawa and Nakamura [8]. They showed the result for social choice functions on the specific domain where every profile is measurable. Their indirect proof using a result for simple games, essentially applies to my lemma. Accordingly, I give a different proof below (after the proof of the main result). Like the above proof of Proposition 1, the proof of the lemma uses the Gibbard-Satterthwaite theorem for finitely many individuals. The proof of the lemma is more complex, however. In the proof of the proposition, deriving in one step that a certain coalition dictates a certain restricted class of profiles was enough. In the proof of the lemma, it is shown (item (iii) and (iv)) in several steps that a certain coalition dictates successively larger classes of profiles.

Alternative Proof of Proposition 1. Suppose that g, $\mathbf{p}$, and $\mathbf{p}'$ satisfy the assumptions. Then (by a well-known property of an ultrafilter) exactly one of the partition elements $\{\, i \in I : x = \max \succ_i^{\mathbf{p}} \}$, where x ranges over the finite X, belongs to the ultrafilter W_g (of Lemma 1). Assume $\{\, i \in I : x = \max \succ_i^{\mathbf{p}} \} \in W_g$ for a particular x. Then since $\mathbf{p}$ and $\mathbf{p}'$ are top equivalent (see (1) above), we have $\{\, i \in I : x = \max \succ_i^{\mathbf{p}'} \} \in W_g$. By the definition of W_g it follows from these two that $g(\mathbf{p}) = x$ and $g(\mathbf{p}') = x$. ∎

Proof of Lemma 1. We check the four properties that defines an ultrafilter.

(i) Suppose $\emptyset \in W_g$. Then for all $\mathbf{p}$ and x, $g(\mathbf{p}) = x$, which is impossible since X has at least three elements.

(ii) Obvious.

(iii) Suppose A, $B \in W_g$. Let $C = A \cap B$. We show $C \in W_g$. Consider four disjoint coalitions C, $A \setminus B$, $B \setminus A$, and $(A \cup B)^c$. Consider the profiles such that all individuals within the same coalition have an identical preference. Regarding each of these four coalitions as an "individual," we can obviously derive from g a social choice function for four individuals, which becomes individually strategyproof. Then the Gibbard-Satterthwaite theorem for finitely many individuals implies that one of the four coalitions dictates these profiles. We show in three steps that C dictates successively larger classes of profiles.

We first claim that C dictates the profiles described in the previous paragraph. To see this, Suppose for example $A \setminus B$ dictates these profiles. But since $B \in W_g$, B dictates these profiles in particular. So, we have two disjoint dictatorial coalitions, which is a contradiction. Hence $A \setminus B$ cannot dictate these profiles. Similarly, $B \setminus A$ cannot dictate. Also, $(A \cup B)^c$ cannot dictate since $A \cup B \in W_g$ by (ii) above.

Next, we show that C dictates those profiles such that different individuals belonging to the same one of the three other coalitions may have different preferences but all individuals in C must have an identical preference. Consider a profile $\mathbf{p}$ such that all individuals in C have the same preference preferring an a the most and all individuals in the same coalition other than C have the same preferences preferring a the least. Since C dictates this profile, $g(\mathbf{p}) = a$. Now suppose for a particular profile $\mathbf{p}'$, all individuals in C have the same preferences as in $\mathbf{p}$ (hence they prefer a the most) but $g(\mathbf{p}') \neq a$. Then at $\mathbf{p}$, C^c can manipulate g by reporting $\mathbf{p}'$.

To conclude that $C \in W_g$ (that is, C dictates all profiles in the domain $\mathcal{S}^I_{\mathcal{B}}$), suppose that for a $\mathbf{p}$ such that individuals in C prefer an a the most (but otherwise may have different preferences), $g(\mathbf{p}) \neq a$. Consider a $\mathbf{p}'$ such that all individuals in C have an identical preference preferring a the most and individuals not in C have the same preferences as in $\mathbf{p}$. From what has been shown above, $g(\mathbf{p}') = a$. Then at $\mathbf{p}$, C can manipulate g by reporting $\mathbf{p}'$.

(iv) Suppose $A \notin W_g$ and $A \in \mathcal{B}$. We must show $A^c \in W_g$. The proof is similar to (iii). Defining a similar social choice function for two individuals, we see either A or $C = A^c$ dictates the restricted profiles. If C dictates the profiles, then $C \in W_g$ as we showed above. If A dictates the profiles, then $A \in W_g$, contradicting the assumption. ∎

4 Extension theorem

I defined a $\mathcal{B}$-social choice function on the domain $\mathcal{S}^I_{\mathcal{B}}$ of $\mathcal{B}$-measurable profiles $\mathbf{p}$. The measurability condition requires that the coalitions $\{ i \in I : x \succ^{\mathbf{p}}_i y \}$ that prefer x to y be "observable," for *all* x and y in X. (Here I am taking the "observability" interpretation as given: "observable" equals "belonging to the Boolean algebra." Questioning the interpretation would require specifying a particular Boolean algebra, as I did elsewhere [10, 12].) Proposition 1, however, implies that a coalitionally strategyproof function uses the information about those coalitions of the form $\{ i \in I : x = \max \succ_i^{\mathbf{p}} \}$ only, ignoring some of the coalitions $\{ i \in I : x \succ_i^{\mathbf{p}} y \}$. But if a coalition $\{ i \in I : x \succ_i^{\mathbf{p}} y \}$ is ignored (irrelevant), requiring it to be "observable" seems too demanding from the informational viewpoint (consider the cognitive burden of the planner).

The following extension theorem states that (i) the function is in fact extendable to the domain $\mathcal{S}^I(\mathcal{B})$ of admissible profiles, which domain does not require the irrelevant coalitions to be "observable" and (ii) the extension is also coalitionally strategyproof. Note that coalitional manipulability in a larger domain does not generally imply coalitional manipulability in a smaller domain. Thus the key to the proof is in suitably constructing profiles $\mathbf{q}$ and $\mathbf{q}'$ in the *smaller* domain such that a certain coalition can manipulate g at $\mathbf{q}$ by reporting $\mathbf{q}'$.

Proposition 2 *Suppose that* $g\colon \mathcal{S}_{\mathcal{B}}^I \to X$ *is a coalitionally strategyproof* $\mathcal{B}$*-social choice function. Then* g *can be extended to a coalitionally strategyproof* $(\mathcal{B})$*-social choice function* $G\colon \mathcal{S}^I(\mathcal{B}) \to X$*.*

Proof. Given a g satisfying the assumption, define G by

$$G(\mathbf{p}) = g(\mathbf{q})$$

for all $\mathbf{p} \in \mathcal{S}^I(\mathcal{B})$, where $\mathbf{q}$ is an arbitrary profile in $\mathcal{S}_{\mathcal{B}}^I$ that is top equivalent to $\mathbf{p}$ (i.e., $\mathbf{q} \simeq \mathbf{p}$). (The existence of such a $\mathbf{q}$ will become evident later.) Note that G is well-defined since Proposition 1 implies that if $\mathbf{q} \simeq \mathbf{p}$ and $\mathbf{q}' \simeq \mathbf{p}$ (hence $\mathbf{q} \simeq \mathbf{q}'$), then $g(\mathbf{q}) = g(\mathbf{q}')$.

Now, to show that G is coalitionally strategyproof, suppose that it is coalitionally manipulable. Then there exist $\mathbf{p}$, $\mathbf{p}' \in \mathcal{S}^I(\mathcal{B})$ and nonempty $E \subseteq I$ such that (a) $\succ_i^{\mathbf{p}'} = \succ_i^{\mathbf{p}}$ for all $i \notin E$ and, (b) $G(\mathbf{p}') \succ_i^{\mathbf{p}} G(\mathbf{p})$ for all $i \in E$.

Fix an arbitrary preference $\succ$ (the "prototype order"). We *obtain* the $\mathbf{q}$ *from* the "prototype order" $\succ$ *by pushing up maximal elements and by pushing down the nonmaximal* $G(\mathbf{p})$ *at* $\mathbf{p}$ as follows: for each i and for each x, we transform $\succ$ in two steps, first from $\succ$ into R_i and second from R_i into $\succ_i^{\mathbf{q}}$:
(i) If $x = \max \succ_i^{\mathbf{p}}$, push x up to the top to get R_i (i.e, for all y, $z \in X$ different from x, xR_iy, and $yR_iz \iff y \succ z$).
(ii) If $G(\mathbf{p}) \neq \max \succ_i^{\mathbf{p}}$, push $G(\mathbf{p})$ down to the bottom to get $\succ_i^{\mathbf{q}}$.
Since all individuals having the same maximal element at $\mathbf{p}$ have an identical

preference at $\mathbf{q}$, we have $\mathbf{q} \in \mathcal{S}^I_\mathcal{B}$ and $\mathbf{q} \simeq \mathbf{p}$. Similarly, we can obtain the $\mathbf{q}' \in \mathcal{S}^I_\mathcal{B}$ from $\succ$ by pushing up maximal elements and by pushing down the nonmaximal $G(\mathbf{p})$ at $\mathbf{p}'$ such that $\mathbf{q}' \simeq \mathbf{p}'$.

Let $i \notin E$. Then, since $\mathbf{q}$ and $\mathbf{q}'$ are both obtained from $\succ$, the preferences $\succ^{\mathbf{q}}_i$ and $\succ^{\mathbf{q}'}_i$ are identical, except possibly for the maximal element and for the least element. But since $\mathbf{q}$ and $\mathbf{q}'$ are obtained by pushing up maximal elements at $\mathbf{p}$ and at $\mathbf{p}'$ (and by pushing down the nonmaximal $G(\mathbf{p})$), under which i's preferences are identical (by (a) above), we have in fact: $\succ^{\mathbf{q}}_i = \succ^{\mathbf{q}'}_i$.

Let $i \in E$. Since $G(\mathbf{p}) \neq \max \succ^{\mathbf{p}}_i$ by (b) above, $G(\mathbf{p})$ is the least element with respect to $\succ^{\mathbf{q}}_i$. But $G(\mathbf{p}) = g(\mathbf{q})$ since $\mathbf{q} \simeq \mathbf{p}$ and $\mathbf{q} \in \mathcal{S}^I_\mathcal{B}$. Hence $g(\mathbf{q})$ is the least element with respect to $\succ^{\mathbf{q}}_i$. But since $g(\mathbf{q}') \neq g(\mathbf{q})$ (because $g(\mathbf{q}') = G(\mathbf{p}') \neq G(\mathbf{p}) = g(\mathbf{q})$), it follows that $g(\mathbf{q}') \succ^{\mathbf{q}}_i g(\mathbf{q})$.

We have established in the preceding two paragraphs that g is coalitionally manipulable, contrary to the assumption. ∎

Corollary 3 *Suppose that $g: \mathcal{S}^I_\mathcal{B} \to X$ is a nondictatorial, coalitionally strategyproof $\mathcal{B}$-social choice function. Then g can be extended to a nondictatorial, coalitionally strategyproof $(\mathcal{B})$-social choice function $G: \mathcal{S}^I(\mathcal{B}) \to X$.*

Proof. If an extension of a social choice function g is dictatorial, then so is g. ∎

I conclude the paper with a sketch of an alternative proof of the extension theorem (Proposition 2). The proof uses the following proposition, which is immediate from Lemma 1. The proposition characterizes the graph of a coalitionally strategyproof $\mathcal{B}$-social choice function, whose domain consists of the measurable profiles of strict preferences.

Proposition 4 *Suppose that $g: \mathcal{S}^I_\mathcal{B} \to X$ is a coalitionally strategyproof $\mathcal{B}$-social choice function. Then there exists an ultrafilter W_g on $\mathcal{B}$ such that for all $\mathbf{p} = (\succ^{\mathbf{p}}_i)_{i \in I} \in \mathcal{S}^I_\mathcal{B}$ and for all $x \in X$,*

$$g(\mathbf{p}) = x \iff \{\, i \in I : x = \max \succ^{\mathbf{p}}_i \,\} \in W_g.$$

To prove the extension theorem using this proposition, define G from $\mathcal{S}^I(\mathcal{B})$ to X by

$$G(\mathbf{p}) = x \iff \{\, i \in I : x = \max \succ^{\mathbf{p}}_i \,\} \in W_g.$$

By a property of an ultrafilter, it is easy to see that G is a well-defined ($\mathcal{B}$)-social choice function extending g. To prove that G is coalitionally strategyproof, use the argument that Pazner and Wesley [14, p.255] give in the second paragraph of the proof of their Theorem.

References

[1] Armstrong TE (1980) Arrow's Theorem with restricted coalition algebras. J Math Econ 7:55–75

[2] Armstrong TE (1985) Precisely dictatorial social welfare functions: Erratum and addendum to 'Arrow's Theorem with restricted coalition algebras'. J Math Econ 14:57–59

[3] Barberà S, Gul F, and Stacchetti E (1993) Generalized median voter schemes and committees. J Econ Theory 61:262–289

[4] Barberà S and Jackson M (1994) A characterization of strategy-proof social choice functions for economies with pure public goods. Soc Choice Welfare 11:241–252

[5] Barberà S, Sonnenschein H, and Zhou L (1991) Voting by committees. Econometrica 59:595–609

[6] Brunner N and Mihara HR (1999) Arrow's theorem, Weglorz' models and the axiom of choice. To appear in Math Logic Quarterly. A version available as ewp-pe/9902001 from the EconWPA

[7] Gibbard A (1973) Manipulation of voting schemes: A general result. Econometrica 41:587–601

[8] Ishikawa S and Nakamura K (1979) The strategy-proof social choice functions. J Math Econ 6:283–295

[9] Mihara HR (1997) Anonymity and neutrality in Arrow's Theorem with restricted coalition algebras. Soc Choice Welfare 14:503–12

[10] Mihara HR (1997) Arrow's Theorem and Turing computability. Econ Theory 10:257–76

[11] Mihara HR (1998) Existence of a coalitionally strategyproof social choice function: A constructive proof. Available as ewp-pe/9604002 from the EconWPA

[12] Mihara HR (1998) Arrow's Theorem, countably many agents, and more visible invisible dictators. To appear in J Math Econ. A version available as ewp-pe/9705001 from the EconWPA

[13] Pattanaik PK (1976) Counter-threats and strategic manipulation under voting schemes. Rev Econ Studies 43:11–18

[14] Pazner EA and Wesley E (1977) Stability of social choices in infinitely large societies. J Econ Theory 14:252–262

[15] Satterthwaite MA (1975) Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. J Econ Theory 10:187–217