

The Scholarship Assignment Problem¹



Pablo Amorós

*Departamento de Teoría e Historia Económica, Universidad de Málaga,
29013 Málaga, Spain
E-mail: pag@uma.es*

Luis C. Corchón

*Departamento de Economía, Universidad Carlos III de Madrid,
28903 Getafe (Madrid), Spain*

and

Bernardo Moreno

*Departamento de Teoría e Historia Económica, Universidad de Málaga,
29013 Málaga, Spain*

Received December 8, 1998

There are n graduate students and n faculty members. Each student will be assigned a scholarship by the joint faculty. The socially optimal outcome is that the best student should get the most prestigious scholarship, the second-best student should get the second most prestigious scholarship, and so on. The socially optimal outcome is common knowledge among all faculty members. Each professor wants one particular student to get the most prestigious scholarship and wants the remaining scholarships to be assigned according to the socially optimal outcome. We consider the problem of finding a mechanism such that in equilibrium, all scholarships are assigned according to the socially optimal outcome. *Journal of Economic Literature* Classification Numbers: D70, D78. © 2002 Elsevier Science

Key Words: implementation; dominant strategies; Nash equilibrium.

¹We thank J. Alcalde, S. Barberà, A. Mas-Colell, W. Thomson, M. Walker, and two anonymous referees for helpful remarks. Financial assistance from Analistas Económicos de Andalucía and CAICYT under Project PB98-0024 is gratefully acknowledged. A previous version of this paper was titled “The Truth, the Whole Truth and Nothing but the Truth.”

1. INTRODUCTION

A group of students apply for a scholarship. A ranking of students must be provided. The social objective is to assign the most prestigious scholarship to the best student, the second most prestigious scholarship to the second best student, and so on. We assume that each student has a different adviser and that there is one scholarship per student, so that the number of applicants, advisers, and scholarships is the same. The ranking of students will be provided by a jury composed of all advisers. There is a true ranking of students, which is known by all advisers. The preferences of advisers are influenced by the true ranking.²

We want to design procedures under which the true ranking is obtained, even when advisers behave strategically to favor their most preferred students.³ This is precisely the aim of implementation theory, which we can describe briefly as follows. There is a set of alternatives and a space of preferences defined on these alternatives. The planner wants to implement a social choice rule mapping preferences into alternatives. To do so, he sets up a mechanism, i.e., a list of message spaces (one for each agent) and an outcome function mapping messages into alternatives. In this paper we focus on a specific implementation problem with the following characteristics:

I. The set of alternatives is the set of possible rankings (permutations) of a given set of agents. Each ranking is interpreted as an assignment of scholarships.

II. The social choice rule is such that, given a true ranking, it yields this ranking as the outcome. We call this rule the **socially optimal choice function** (SOCF).

III. The true ranking is observed perfectly by all deciders. The set of deciders is identical to the set of agents. This may be interpreted as saying that each agent has a decider who is on his side.

IV. Preferences are such that the following two conditions are met: (1) Each professor wants a particular student to get the most prestigious scholarship, and (2) each professor wants the remaining scholarships to be assigned according to the true ranking. We call these preferences “moderately selfish.” We feel that these preferences may be a reasonable approx-

²Other situations with a similar structure include rankings of Ph.D. programs, wine tasting, gymnastic competitions, and choosing fellows for a society.

³Most procedures used in real life are designed to avoid blatant manipulation; i.e., the higher and lower scores received by any agent do not count, people cannot vote for themselves, and so on.

imation of situations in which professors judge the performance of other students unbiasedly.

Despite the enormous literature generated by implementation theory in the last three decades, little attention has been paid to the particular problem described above (see Sect. 5 for a discussion of this point).

It is easy to show that with two agents, the SOCF cannot be implemented in any equilibrium concept when preferences are moderately selfish. Thus, we study implementation of the SOCF in the case of three or more agents.

We first focus our attention on implementation in dominant strategies. Since preference profiles do not have a Cartesian product structure, the standard revelation principle cannot be applied to our problem. Moreover, the Gibbard–Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975) is of no direct application here, since the domain of admissible preferences is restricted to moderately selfish preferences.⁴ Unfortunately, this restriction is not powerful enough to bring positive results: If there are three agents, then the SOCF cannot be implemented in dominant strategies (Theorem 1).

Next, we consider Nash implementation. We first notice that the SOCF can be implemented by the “canonical” Maskin mechanism as described by Williams (1984), Repullo (1987), Saijo (1988), and McKelvey (1989). This follows from the fact that the SOCF is Maskin monotonic and satisfies no veto power. However, the canonical mechanism has been subject to a fair amount of criticism (see Jackson, 1992), and thus we are led to the study of “nice” mechanisms. To obtain some intuition about what “nice” may mean in this context, we first study mechanisms used in real life, including the Borda count, the plurality rule, and others. We show that none of these mechanisms implements the SOCF in Nash equilibrium (most of them create equilibria yielding undesired outcomes). However, all of these mechanisms use similar message spaces: Messages are rankings and/or real numbers that reflect these rankings. Since real numbers create a kind of integer game, given the latter’s controversial nature (see Jackson, 1992), we avoid real numbers.

We present two mechanisms in which the message space for each agent is the space of all possible rankings. The first mechanism implements the SOCF in Nash equilibrium when there are three agents (Theorem 2), the second implements this rule in Nash equilibrium when the number of agents is greater than three (Theorem 3). We do not know whether there is a

⁴There are many proofs of the Gibbard–Satterthwaite theorem in a restricted domain (see, e.g., Aswal and Sen 1997). The main difference between the domain used in those proofs and our domain is that in our case, the set of admissible preferences profiles does not have a Cartesian product structure.

single “nice” mechanism implementing the SOCF for any number of agents. Since the number of agents involved in situations considered in this paper is public knowledge, we do not regard this as very important.

The rest of the paper goes as follows. Section 2 presents the model, and Section 3 studies dominant strategies. Section 4 gathers our results on Nash implementation. Finally, Section 5 discusses possible extensions of our work.

2. THE MODEL

Let N be a set of n students applying for scholarships. A social alternative, π , is an **assignment of scholarships** to the students, i.e., a ranking (permutation) of the elements of N , so that the student in the first position gets the most prestigious scholarship, the student in the second position gets the second most prestigious scholarship, and so on. Let Π be the **set of assignments of scholarships** (i.e., the set of all rankings of students in N). For all $\pi \in \Pi$ and $i \in N$, we denote by P_i^π the position of student i in ranking π .

Each student, $i \in N$, has a different adviser. The final assignment of scholarships will be decided by the group of students’ advisers. In the sequel we will write agent i to denote both student i and his adviser. We assume that there exists a **true ranking** of the students, $\pi_i \in \Pi$, known by all agents. The socially optimal alternative is that the scholarships should be assigned according to the true ranking. We assume, however, that the true ranking is not verifiable. Thus, to elicit the socially optimal assignment of scholarships, we must rely on announcements made by agents. This is the idea behind the concept of a mechanism.

DEFINITION 1. A **mechanism** Γ is a tuple (S, g) , where $S = \times_{i \in N} S_i$ is a list of message spaces (one for each agent), and $g : S \rightarrow \Pi$ is an outcome function.

A profile of messages is denoted by $s \in S$. For all agents $i \in N$ and all profiles of messages $s \in S$, let s_i denote the message of agent i and $s_{-i} \in S_{-i} = \times_{j \in N \setminus \{i\}} S_j$ the messages of all agents except i .

The description of preferences is slightly more complicated in our case than in the standard case. The complication arises because here, the preference relation of each agent depends on the true ranking.

Let \mathfrak{R} denote the class of preference relations defined over Π satisfying reflexivity, transitivity, and completeness. Each agent $i \in N$ has a **preference function** $\succeq_i : \Pi \rightarrow \mathfrak{R}$, which associates with each true ranking, $\pi_i \in \Pi$, a preference relation $\succeq_i^{\pi_i} \in \mathfrak{R}$, where $\succ_i^{\pi_i}$ denotes the strict preference relation associated with $\succeq_i^{\pi_i}$. For instance, let $N = \{a, b\}$. Then

$\Pi = \{(a, b), (b, a)\}$. A possible preference function for agent a is $(a, b) \succ_a^{\pi_i} (b, a)$ if $\pi_i = (a, b)$ and $(b, a) \succ_a^{\pi_i} (a, b)$ if $\pi_i = (b, a)$.

We make two assumptions about the agents' preference functions. First, we assume that agents are selfish in the sense that, when comparing two different assignments, each agent prefers the one in which he is in a better position, whatever the true ranking.

DEFINITION 2. Agent i 's preference function, $\succeq_i: \Pi \rightarrow \mathfrak{R}$, is **selfish** if for all $\pi_i \in \Pi$ and all $\pi, \hat{\pi} \in \Pi$ with $P_i^\pi < P_i^{\hat{\pi}}$, we have $\pi \succ_i^{\pi_i} \hat{\pi}$.

Second, we also assume that, given a fixed position for an agent, the agent prefers the rest of the agents to be arranged as close as possible to the true ranking. Let $N(\pi, \hat{\pi}) = \{i \in N: P_i^\pi = P_i^{\hat{\pi}}\}$. In words, $N(\pi, \hat{\pi})$ is the set of agents who are in the same position in assignments π and $\hat{\pi}$.

DEFINITION 3. Agent i 's preference function, $\succeq_i: \Pi \rightarrow \mathfrak{R}$, is **unprejudiced** with respect to the other agents when, for all $\pi_i \in \Pi$ and all $\pi, \hat{\pi} \in \Pi$ satisfying the following three conditions:

- (1) $\pi \neq \hat{\pi}$,
 - (2) $i \in N(\pi, \hat{\pi})$, and
 - (3) for all $j, k \notin N(\pi, \hat{\pi})$, if $P_j^{\pi_i} < P_k^{\pi_i}$, then $P_j^\pi < P_k^\pi$,
- we have $\pi \succ_i^{\pi_i} \hat{\pi}$.

We say that a preference function is **moderately selfish** when it is selfish and unprejudiced. The following examples may clarify this concept.

EXAMPLE 1. Let $N = \{a, b, c, d\}$. Suppose that the preference function of agent d , $\succeq_d: \Pi \rightarrow \mathfrak{R}$, is moderately selfish. Then selfishness implies that $(c, d, b, a) \succ_d^{\pi_i} (a, b, d, c)$ for all $\pi_i \in \Pi$.

Suppose now that the true ranking is $\pi_i = (a, b, c, d)$. Consider assignments $\pi = (b, d, c, a)$ and $\hat{\pi} = (c, d, b, a)$. Notice that $N(\pi, \hat{\pi}) = \{a, d\}$. Since agents who are not in $N(\pi, \hat{\pi})$ are arranged among them in accordance with the truth in π , the unprejudiced condition implies that $\pi \succ_d^{\pi_i} \hat{\pi}$. However, the moderately selfish condition is not strong enough to determine whether $(c, d, a, b) \succeq_d^{\pi_i} (b, d, c, a)$ or $(b, d, c, a) \succeq_d^{\pi_i} (c, d, a, b)$.

Moderate selfishness is a severe restriction on the domain of admissible preference functions. Actually, in the three-agents case there is a unique preference function for each agent that satisfies this condition.

EXAMPLE 2. Let $N = \{a, b, c\}$. The set of all possible assignments is $\Pi = \{(a, b, c), (a, c, b), (b, a, c), (b, c, a), (c, a, b), (c, b, a)\}$. So the only preference function of agent a satisfying moderate selfishness is the one defined in Table I. Here each column represents the (strict) preference relation associated with a different true ranking.

TABLE I
Moderately Selfish Preference Function of Agent a when $N = \{a, b, c\}$

$\succeq_a^{(a,b,c)}$	$\succeq_a^{(b,a,c)}$	$\succeq_a^{(b,c,a)}$	$\succeq_a^{(a,c,b)}$	$\succeq_a^{(c,a,b)}$	$\succeq_a^{(c,b,a)}$
(a, b, c)	(a, b, c)	(a, b, c)	(a, c, b)	(a, c, b)	(a, c, b)
(a, c, b)	(a, c, b)	(a, c, b)	(a, b, c)	(a, b, c)	(a, b, c)
(b, a, c)	(b, a, c)	(b, a, c)	(c, a, b)	(c, a, b)	(c, a, b)
(c, a, b)	(c, a, b)	(c, a, b)	(b, a, c)	(b, a, c)	(b, a, c)
(b, c, a)	(b, c, a)	(b, c, a)	(c, b, a)	(c, b, a)	(c, b, a)
(c, b, a)	(c, b, a)	(c, b, a)	(b, c, a)	(b, c, a)	(b, c, a)

Let F_{ms} denote the class of moderately selfish preference functions. A profile of preference functions is denoted by $\succeq \in F_{ms}^n$. A **state of the world** is a list of preference functions and a true ranking observed by all agents, i.e., $(\succeq, \pi_t) \in F_{ms}^n \times \Pi$. Let ξ_{ms} be the class of states of the world. A **profile of preference relations is admissible** if there exists some state of the world $(\succeq, \pi_t) \in \xi_{ms}$ such that \succeq^{π_t} coincides with these preference relations. Let R be the set of admissible profiles of preference relations.

Given a state of the world and a mechanism, agents have to make decisions about the message to be sent. We follow the standard procedure in implementation theory of capturing this decision by means of a game-theoretical equilibrium concept. Two examples of these concepts follow.

DEFINITION 4. Let $\Gamma = (S, g)$ be a mechanism. We say that $s \in S$ is a **dominant strategy equilibrium** of Γ at state of the world $(\succeq, \pi_t) \in \xi_{ms}$ if for all $i \in N$, $\hat{s}_i \in S_i$ and $\hat{s}_{-i} \in S_{-i}$, $g(s_i, \hat{s}_{-i}) \succeq_i^{\pi_t} g(\hat{s}_i, \hat{s}_{-i})$.

DEFINITION 5. Let $\Gamma = (S, g)$ be a mechanism. We say that $s \in S$ is a **Nash equilibrium** of Γ at state of the world $(\succeq, \pi_t) \in \xi_{ms}$ if for all $i \in N$ and $\hat{s}_i \in S_i$, $g(s) \succeq_i^{\pi_t} g(\hat{s}_i, s_{-i})$.

Given a mechanism Γ and a state of the world $(\succeq, \pi_t) \in \xi_{ms}$, let $D(\Gamma, \succeq^{\pi_t})$ and $N(\Gamma, \succeq^{\pi_t})$ denote the sets of dominant strategy and Nash equilibria of Γ at (\succeq, π_t) , respectively.

Our objective in this paper is to implement the SOCF, $\Phi : \xi_{ms} \rightarrow \Pi$. This function associates with each state of the world the assignment of scholarships corresponding to the true ranking for that state of the world, i.e., $\Phi(\succeq, \pi_t) = \pi_t$ for all $(\succeq, \pi_t) \in \xi_{ms}$.

DEFINITION 6. A mechanism $\Gamma = (S, g)$ **implements** the SOCF in dominant strategy equilibrium (resp., in Nash equilibrium) if for all $(\succeq, \pi_t) \in \xi_{ms}$, $g(D(\Gamma, \succeq^{\pi_t})) = \{\pi_t\}$ (resp., $g(N(\Gamma, \succeq^{\pi_t})) = \{\pi_t\}$).

When $n = 2$, if preference functions are moderately selfish, then preference relations do not change with the true ranking. Therefore, the SOCF

cannot be implemented in dominant strategies or in Nash equilibrium.⁵ From now on, we assume that $n \geq 3$.

3. IMPLEMENTATION IN DOMINANT STRATEGIES

As we noted in the previous section, moderate selfishness is a serious restriction on the domain of admissible preference functions. Therefore, the domain of admissible preference relations is severely restricted as well. To see this, consider the three-agents case examined in Example 2. While under the unrestricted domain, there are $6! = 720$ different (strict) preference relations over Π , only 2 of them are compatible with the moderate-selfishness condition. (In Example 2, the only two preference relations compatible with moderate selfishness for agent a are $\succeq_a^{(a,b,c)} = \succeq_a^{(b,a,c)} = \succeq_a^{(b,c,a)}$ and $\succeq_a^{(a,c,b)} = \succeq_a^{(c,a,b)} = \succeq_a^{(c,b,a)}$, as defined in Table I).

The drastic reduction in the number of admissible preference relations suggests that we might avoid the usual impossibility results on dominant strategies implementation. Unfortunately, this is not so, at least in the three-agents case.

THEOREM 1. *If $n = 3$, then the SOCF cannot be implemented in dominant strategies.*

Proof. Let $N = \{a, b, c\}$. As we have seen in Example 2, in this case there is a unique preference function for each agent satisfying moderate selfishness. Let $\succeq_a, \succeq_b, \succeq_c \in F_{ms}$ be the moderate-selfish preference functions of agents a, b , and c , respectively. Then, abusing notation, let $\succeq_a^1 \equiv \succeq_a^{(a,b,c)} = \succeq_a^{(b,a,c)} = \succeq_a^{(b,c,a)}$ and $\succeq_a^2 \equiv \succeq_a^{(a,c,b)} = \succeq_a^{(c,a,b)} = \succeq_a^{(c,b,a)}$. Similarly, for agents b and c , let $\succeq_b^1 \equiv \succeq_b^{(b,a,c)} = \succeq_b^{(a,b,c)} = \succeq_b^{(a,c,b)}$, $\succeq_b^2 \equiv \succeq_b^{(b,c,a)} = \succeq_b^{(c,b,a)} = \succeq_b^{(c,a,b)}$ and $\succeq_c^1 \equiv \succeq_c^{(c,a,b)} = \succeq_c^{(a,c,b)} = \succeq_c^{(a,b,c)}$, $\succeq_c^2 \equiv \succeq_c^{(c,b,a)} = \succeq_c^{(b,a,c)} = \succeq_c^{(b,c,a)}$.

Notice that the set of admissible profiles of preference relations does not have a Cartesian product structure, since they are partly determined by the true ranking (e.g., $\succeq_a^1, \succeq_b^2, \succeq_c^1 \in \mathfrak{R}$, but $(\succeq_a^1, \succeq_b^2, \succeq_c^1) \notin \mathfrak{R}$). Therefore, the standard revelation principle cannot be applied here.⁶

Suppose that there exists a mechanism $\Gamma_D(S, g)$ implementing Φ in dominant strategies. Let $s_i^1 \in S_i$ ($s_i^2 \in S_i$) be a dominant strategy for agent i

⁵In this case the SOCF cannot be implemented in any type of equilibrium.

⁶The revelation principle states the following necessary condition for a choice function to be implementable in dominant strategies: In the manipulation game associated with the choice function, to announce the true profile of preference relations must be a dominant strategy equilibrium.

TABLE II
Relevant Part of Mechanism Γ_D

		Players b and c			
		$s_b^1 s_c^1$	$s_b^2 s_c^1$	$s_b^1 s_c^2$	$s_b^2 s_c^2$
Player a	s_a^1	(a, b, c)	?	(b, a, c)	(b, c, a)
	s_a^2	(a, c, b)	(c, a, b)	?	(c, b, a)

when his preference relation is \succeq_i^1 (resp. \succeq_i^2). Since Γ_D implements Φ in dominant strategies, the relevant part of Γ_D is given in Table II.

Notice that the profiles of strategies $(s_a^1, s_b^2, s_c^1) \in \mathcal{S}$ and $(s_a^2, s_b^1, s_c^2) \in \mathcal{S}$ correspond with the profiles of preference relations $(\succeq_a^1, \succeq_b^2, \succeq_c^1) \in \mathfrak{N}^n$ and $(\succeq_a^2, \succeq_b^1, \succeq_c^2) \in \mathfrak{N}^n$, respectively. However, these profiles do not correspond with any possible true ranking, i.e., $(\succeq_a^1, \succeq_b^2, \succeq_c^1) \notin R$ and $(\succeq_a^2, \succeq_b^1, \succeq_c^2) \notin R$. Now we have to specify $g(s_a^1, s_b^2, s_c^1)$.

Claim 1. Since $(s_a^2, s_b^2, s_c^1) \in D(\Gamma_D, \succeq^{(c, a, b)})$ and $g(s_a^2, s_b^2, s_c^1) = (c, a, b)$, we have $g(s_a^1, s_b^2, s_c^1) \notin \{(a, b, c), (a, c, b)\}$; otherwise, $g(s_a^1, s_b^2, s_c^1) \succ_a^{(c, a, b)} g(s_a^2, s_b^2, s_c^1)$, which is a contradiction.

Claim 2. Since $(s_a^1, s_b^1, s_c^1) \in D(\Gamma_D, \succeq^{(a, b, c)})$ and $g(s_a^1, s_b^1, s_c^1) = (a, b, c)$, we have $g(s_a^1, s_b^2, s_c^1) \notin \{(b, a, c), (b, c, a)\}$; otherwise, $g(s_a^1, s_b^2, s_c^1) \succ_b^{(a, b, c)} g(s_a^1, s_b^1, s_c^1)$, which is a contradiction.

Claim 3. Since $(s_a^1, s_b^2, s_c^2) \in D(\Gamma_D, \succeq^{(b, c, a)})$ and $g(s_a^1, s_b^2, s_c^2) = (b, c, a)$, we have $g(s_a^1, s_b^2, s_c^1) \notin \{(c, a, b), (c, b, a)\}$; otherwise, $g(s_a^1, s_b^2, s_c^1) \succ_c^{(b, c, a)} g(s_a^1, s_b^2, s_c^2)$, which is a contradiction.

Claims 1, 2, and 3 contradict the definition of an outcome function. ■

4. IMPLEMENTATION IN NASH EQUILIBRIUM

The impossibility result shown earlier leads us to study the implementation of the SOCF in Nash equilibrium. A standard result in the theory of Nash implementation (see, e.g., Repullo, 1987) shows that, if there are three or more agents, any choice function satisfying Maskin monotonicity⁷ and no veto power⁸ is implementable in Nash equilibrium. It is easy to

⁷Roughly speaking, this condition says that if the choice function selects some outcome for some preferences profile, then it must select the same outcome if it becomes more preferred by all agents.

⁸No veto power means that when $n \geq 3$, if there exists an outcome which is the most preferred for at least $n - 1$ agents, then it must be selected by the choice function.

check that, under the moderate-selfishness assumption, the SOCF satisfies both conditions.

The canonical mechanism used in the proof of the result invoked above has been subject to a fair amount of criticism (see, e.g., Jackson, 1992). Customary complaints include that the message spaces are too complex and that rules of the game include integer games. Many researchers feel that when implementing a particular choice function, message spaces must be simple and the rules of the game must be as natural as possible. A difficulty with this argument is that in general, it is not easy to define what “simple” and “natural” mean. However, in our case, hints about what these properties mean may be found in the study of mechanisms used in the real world and the literature on voting. Therefore, we first review some of those mechanisms.

Borda Mechanism (Γ_B). This is the natural mechanism associated with the Borda rule. There is a common set of $n!$ different scores. Agents identify each assignment in Π with one of these scores (with the restriction that each agent cannot give the same score to two different assignments). The assignment that receives the maximum score is chosen.

Modified Borda Mechanism (Γ_{MB}). In this variation of the Borda mechanism, there is a common set of n different scores. Agents identify each agent with one of these scores (with the restriction that each agent cannot give the same score to two different agents). The total score received by each agent determines the final assignment.

Plurality Mechanism (Γ_P). In this mechanism, each agent has to vote for one possible assignment. The assignment receiving the most number of votes is chosen.⁹

Scoring Mechanism (Γ_{SC}). In this mechanism, each agent has to announce a score in a given interval for each agent, and agents are arranged according to the total scores received.

Modified Scoring Mechanism (Γ_{MSC}). This is a variation of the scoring mechanism that tries to avoid the incentive that each agent has to give the maximum score to himself and the minimum score to other agents. Thus, the highest and the lowest scores received by the agent do not count.

Unfortunately, as we show in the following examples, all of these mechanisms may fail to implement the SOCF in Nash equilibrium.

EXAMPLE 3 (Borda Mechanism). Let $N = \{a, b, c\}$. Then $\Pi = \{(a, b, c), (a, c, b), (b, a, c), (b, c, a), (c, a, b), (c, b, a)\}$. Let $\Gamma_B = (S, g)$

⁹It is easy to see that under our assumptions, no Condorcet winner exists.

TABLE III
Unwanted Nash Equilibrium of Mechanism Γ_B

	Score for each assignment					
	(a, b, c)	(a, c, b)	(b, a, c)	(b, c, a)	(c, a, b)	(c, b, a)
s_a	5	4	3	6	2	1
s_b	4	2	5	6	1	3
s_c	2	3	1	6	5	4
s^π	11	9	9	18	8	8

be as follows. For all $i \in N$, $S_i = \{(s_i^\pi)_{\pi \in \Pi} \in \{1, 2, 3, 4, 5, 6\}^{3!} : \text{for all } \pi \neq \hat{\pi}, s_i^\pi \neq s_i^{\hat{\pi}}\}$. For all $s \in S$ and $\pi \in \Pi$, let $s^\pi = \sum_{i \in N} s_i^\pi$. The outcome function $g : S \rightarrow \Pi$ is such that, if for some $s \in S$ and $\pi \in \Pi$, $\pi = \arg \max_{\hat{\pi} \in \Pi} \{s^{\hat{\pi}}\}$, then $g(s) = \pi$. Let $\succeq \in F_{ms}^3$ and $\pi_t = (a, b, c)$. Table III gives a profile of strategies $s \in S$ such that $s \in N(\Gamma_B, \succeq^{\pi_t})$ but $g(s) = (b, c, a) \neq \pi_t$.

EXAMPLE 4 (Modified Borda Mechanism). Let $N = \{a, b, c, d, e\}$. Let $\Gamma_{MB} = (S, g)$ be as follows. For all $i \in N$, let $S_i = \{(s_i^j)_{j \in N} \in \{1, 2, 3, 4, 5\}^5 : \text{for all } j \neq k, s_i^j \neq s_i^k\}$. For all $s \in S$ and $i \in N$, let $s^i = \sum_{j \in N} s_j^i$. The outcome function $g : S \rightarrow \Pi$ is such that, for all $s \in S$ and $i, j \in N$, if $s^i > s^j$, then $P_i^{g(s)} < P_j^{g(s)}$. Let $\succeq \in F_{ms}^5$ and $\pi_t = (a, b, c, d, e)$. Table IV gives a profile of strategies $s \in S$ such that $s \in N(\Gamma_{MB}, \succeq^{\pi_t})$ but $g(s) = (b, a, c, d, e) \neq \pi_t$.

EXAMPLE 5 (Plurality Mechanism). Let N be such that $n \geq 3$. Let $\Gamma_P = (S, g)$ be as follows. For all $i \in N$, let $S_i = \{(s_i^\pi)_{\pi \in \Pi} \in \{0, 1\}^{n!} : \text{there is a unique } \pi \in \Pi \text{ with } s_i^\pi = 1\}$. For all $s \in S$ and $\pi \in \Pi$, let $s^\pi = \sum_{i \in N} s_i^\pi$. The outcome function $g : S \rightarrow \Pi$ is such that, if for some $s \in S$ and $\pi \in \Pi$, $\pi = \arg \max_{\hat{\pi} \in \Pi} \{s^{\hat{\pi}}\}$, then $g(s) = \pi$. Let $\succeq \in F_{ms}^n$ and $\pi_t \in \Pi$. Let $s \in S$ be such that, for some $\pi \in \Pi$ with $\pi \neq \pi_t$ and all $i \in N$, $s_i^\pi = 1$. Notice that $s \in N(\Gamma_P, \succeq^{\pi_t})$ but $g(s) = \pi \neq \pi_t$.

EXAMPLE 6 (Scoring and Modified Scoring Mechanisms). Let $N = \{a, b, c, d\}$. Let $\Gamma_{SC} = (S, g)$ be as follows. For all $i \in N$, let $S_i = \{(s_i^j)_{j \in N} \in [0, 10]^4\}$. For all $s \in S$ and $i \in N$, let $s^i = \sum_{j \in N} s_j^i$. The outcome function $g : S \rightarrow \Pi$ is such that, for all $s \in S$ and $i, j \in N$, if $s^i > s^j$, then $P_i^{g(s)} < P_j^{g(s)}$. Let $\succeq \in F_{ms}^4$ and $\pi_t = (a, b, c, d)$. Table V gives a profile of strategies $s \in S$ such that $s \in N(\Gamma_{SC}, \succeq^{\pi_t})$ but $g(s) = (b, a, c, d) \neq \pi_t$. A similar example can be used to show that the Modified Scoring mechanism presents identical shortcomings.

TABLE IV
Unwanted Nash Equilibrium of Mechanism Γ_{MB}

	Score for each agent				
	a	b	c	d	e
s_a	4	5	3	2	1
s_b	4	5	3	2	1
s_c	3	5	4	2	1
s_d	3	5	1	4	2
s_e	4	5	3	2	1
s^i	18	25	14	12	6

All of the foregoing mechanisms use similar message spaces that are rankings of students and/or real numbers. The rules of the game reflect the fact that the actual assignments provided by the mechanisms must be positively associated with the rankings provided by agents.

Next, we provide two mechanisms that implement the SOCF in Nash equilibrium. In both mechanisms the message spaces are sets of possible rankings of students. The first mechanism works for $n = 3$, and the second one works for $n \geq 4$. We do not know whether there is a single, natural mechanism for implementing the SOCF for all $n \geq 3$.

In the mechanism for $n = 3$, each agent announces a position for the rest of the agents. The position of an agent $i \in N$ in the final assignment is the lowest number among the announcements of all the other agents. In case of a tie between two agents, the relative ordering among the agents involved in the tie is decided by the agent not involved in the tie. If all agents are tied, then an arbitrary assignment occurs.

Mechanism 1 (Γ_1). Let $n = 3$. Let $\Gamma_1 = (S, g)$ be as follows. For all $i \in N$, $S_i = \{(s_i^j)_{j \in N \setminus \{i\}} \in \{1, 2, 3\}^2 : \text{for all } j \neq k, s_i^j \neq s_i^k\}$. We interpret s_i^j as the position that agent i announces for agent j . Let $S = \times_{i \in N} S_i$.

TABLE V
Unwanted Nash Equilibrium of Mechanism Γ_{SC}

	Score for each agent			
	a	b	c	d
s_a	10	10	1	0
s_b	0	10	0	0
s_c	5	10	10	0
s_d	5	10	3	10
s^i	20	40	14	10

For all $s \in S$, let $s_{\min}^i = \min_{j \in N \setminus \{i\}} \{s_j^i\}$ (i.e., s_{\min}^i is the best position for agent i announced by the other two agents). Let $\pi_r \in \Pi$ be an arbitrary assignment known by all agents. The outcome function $g : S \rightarrow \Pi$ consists of the following three rules. For all $i, j, k \in N$,

- (1) if $s_{\min}^i < s_{\min}^j$, then $P_i^{g(s)} < P_j^{g(s)}$;
- (2) if $s_{\min}^i = s_{\min}^j \neq s_{\min}^k$, then $P_i^{g(s)} < P_j^{g(s)}$ iff $s_k^i < s_k^j$; and
- (3) if $s_{\min}^i = s_{\min}^j = s_{\min}^k$, then $g(s) = \pi_r$.

As the following theorem states, Mechanism 1 implements the SOCF in Nash equilibrium when $n = 3$. We omit the proof of this result in the interests of brevity (the proof is obtainable by request). The intuition is that (1) truth-telling is a Nash equilibrium, and (2) if the assignment selected by the outcome function is not the corresponding to the true ranking, then there are two agents that are not arranged in accordance with the true ranking. In the latter case, the third agent can deviate and arrange them according to the truth, without changing his position.

THEOREM 2. *If $n = 3$, then Mechanism 1 implements the SOCF in Nash equilibrium.*

Unfortunately, extension of Mechanism 1 to the case of more than three agents is not straightforward. The problem is to select the agent who breaks ties. We now present a mechanism for implementing the SOCF in Nash equilibrium when there are four or more agents. This mechanism sidesteps the previous difficulty at the cost of making the outcome function less transparent. In this mechanism only four agents are strategically active. One of these (say agent d) states a relative order for the rest of agents, and the other three, say a, b , and c , determine the place of agent d .

Mechanism 2 (Γ_2). Let $n \geq 4$. Let $\Gamma_2 = (S, g)$ be as follows. Only four agents, say $a, b, c, d \in N$, have to send messages. Let $\Pi^{N \setminus \{d\}}$ be the set of all possible rankings of the agents in $N \setminus \{d\}$. Then $S_d = \Pi^{N \setminus \{d\}}$ is the message space for agent d . Let $S_a = S_b = S_c = \{1, 2, \dots, n\}$ be the message spaces of the other three agents. For all $i \in \{a, b, c\}$, $s_i \in S_i$ can be interpreted as the position that agent i announces for agent d . Let $S = S_a \times S_b \times S_c \times S_d$. The outcome function $g : S \rightarrow \Pi$ is defined as follows. For all $s \in S$, $g(s)$ is such that,

- (1) agents in $N \setminus \{d\}$ are arranged among them in accordance with s_d ,

(2) the position of agent d is such that

$$P_d^{g(s)} = \begin{cases} n & \text{if } s_a = s_b = s_c \\ \min\{s_a, s_b, s_c\} & \text{if for some } i \in \{a, b, c\}, s_i = n \\ \max\{s_a, s_b, s_c\} & \text{if for all } i \in \{a, b, c\}, s_i \neq n, \text{ and} \\ & \text{for some } i, j, k \in \{a, b, c\}, s_i = s_j \neq s_k \\ \text{med}\{s_a, s_b, s_c\} & \text{if for all } i \in \{a, b, c\}, s_i \neq n, \text{ and} \\ & \text{for all } i, j \in \{a, b, c\}, s_i \neq s_j, \end{cases} \quad (1)$$

where $\text{med}\{\cdot\}$ denotes the median element of the corresponding set.

In Mechanism 2 we need only four agents whose preferences are moderately selfish, whatever the number of agents.

Mechanism 2 implements the SOCF in Nash equilibrium when $n \geq 4$. The intuition behind the proof of this result is as follows. On the one hand, agent d decides the relative order of the rest of the agents, and he cannot influence his own position. Since his preference function is moderately selfish, he has a strictly dominant strategy—namely, to tell the truth. On the other hand, if the position of agent d is not the true one, then some agent has an incentive to deviate and to change the position of d .

The rule used to determine the final position of agent d is not very intuitive. Since agent d can be trusted to reveal truthfully the relative order of all other agents, one might think that it should be easier to use the messages of the remaining agents to place agent d in the right position. The complexity of the former rule, however, arises from the need to rule out unwanted Nash equilibria.

THEOREM 3. *If $n \geq 4$, then Mechanism 2 implements the SOCF in Nash equilibrium.*

Proof. Let $\{a, b, c, d\} \subseteq N$. Let $\Gamma_2 = (S, g)$ be as defined above. Let $\succeq = (\succeq_a, \succeq_b, \succeq_c, \succeq_d) \in F_{ms}^4$ and $\pi_t \in \Pi$. Suppose, without loss of generality, that $P_a^{\pi_t} < P_b^{\pi_t} < P_c^{\pi_t}$.

Notice that for all $s = (s_d, s_{-d})$, $\hat{s} = (\hat{s}_d, s_{-d}) \in S$, $P_d^{g(s)} = P_d^{g(\hat{s})}$ (i.e., agent d cannot change his position given s_{-d}). Then, since $\succeq_d \in F_{ms}$, agent d has a strictly dominant message, $\tilde{s}_d \in S_d$, such that for all $i, j \in N \setminus \{d\}$, $P_i^{\tilde{s}_d} < P_j^{\tilde{s}_d}$ iff $P_i^{\pi_t} < P_j^{\pi_t}$. Now let $\tilde{s}_{-d} = (\tilde{s}_a, \tilde{s}_b, \tilde{s}_c) \in S_{-d}$ be as follows:

$$(\tilde{s}_a, \tilde{s}_b, \tilde{s}_c) = \begin{cases} (n, n, n) & \text{if } P_c^{\pi_t} < P_d^{\pi_t} = n \\ (P_d^{\pi_t}, P_d^{\pi_t}, n) & \text{if } P_c^{\pi_t} < P_d^{\pi_t} < n \\ (P_d^{\pi_t}, n, n) & \text{if } P_a^{\pi_t} < P_d^{\pi_t} < P_c^{\pi_t} \\ (P_a^{\pi_t}, P_d^{\pi_t}, n) & \text{if } P_d^{\pi_t} < P_a^{\pi_t}. \end{cases} \quad (2)$$

It is easy to see that $\tilde{s} = (\tilde{s}_d, \tilde{s}_{-d}) \in N(\Gamma_2, \succeq^{\pi_t})$ and $g(\tilde{s}) = \pi_t$.

Now we show that all $s \in N(\Gamma_2, \succeq^{\pi_t})$ is such that $g(s) = \pi_t$. Since $\tilde{s}_d \in S_d$ is a strictly dominant message for agent d , this is equivalent to showing that there is no $s_{-d} \in S_{-d}$ such that $(\tilde{s}_d, s_{-d}) \in N(\Gamma_2, \succeq^{\pi_t})$ and $P_d^{g(\tilde{s}_d, s_{-d})} \neq P_d^{\pi_t}$. Suppose on the contrary that there is some $s_{-d} = (s_a, s_b, s_c) \in S_{-d}$ like that. We make four claims.

Claim 1. For some $i, j \in \{a, b, c\}$, $s_i \neq s_j$.

Suppose, on the contrary that $s_a = s_b = s_c$. Then $P_d^{g(\tilde{s}_d, s_{-d})} = n \neq P_d^{\pi_t}$. Let $\hat{s}_a \in S_a$ be as follows:

$$\hat{s}_a = \begin{cases} n & \text{if } \max\{P_a^{\pi_t}, P_d^{\pi_t}\} \leq s_a \neq n \\ \max\{P_a^{\pi_t}, P_d^{\pi_t}\} & \text{otherwise.} \end{cases} \quad (3)$$

It is easy to see that since $\succeq_a \in \Gamma_{ms}$, $g(\hat{s}_a, s_{-a}) \succ_a^{\pi_t} g(s)$, which is a contradiction.

Claim 2. For all $i \in \{a, b, c\}$, $s_i \neq n$.

Suppose, on the contrary that for some $i \in \{a, b, c\}$, $s_i = n$. Let $j \in \{a, b, c\}$ be such that $s_j = \min\{s_a, s_b, s_c\}$. Then $P_d^{g(\tilde{s}_d, s_{-d})} = s_j \neq P_d^{\pi_t}$, and by Claim 1, $s_j < n$. Let $k \in \{a, b, c\}$ be such that $i \neq k \neq j$. We distinguish two cases.

Case 2.1. Suppose that $P_j^{\pi_t} < P_d^{\pi_t}$.

Subcase 2.1.1. Suppose that $P_d^{\pi_t} < s_j$. Then agent j prefers to deviate sending $\hat{s}_j = P_d^{\pi_t}$.

Subcase 2.1.2. Suppose that $s_j < P_d^{\pi_t}$. If $s_j \neq s_k$, then agent j prefers to send $\hat{s}_j = s_j + 1$. If $s_j = s_k$, then agent i prefers to deviate sending either $\hat{s}_i = P_d^{\pi_t}$ (if $P_d^{\pi_t} \neq n$) or $\hat{s}_i = s_j$ (if $P_d^{\pi_t} = n$).

Case 2.2. Suppose that $P_d^{\pi_t} < P_j^{\pi_t}$.

Subcase 2.2.1. Suppose that $P_j^{\pi_t} < s_j$. Then agent j prefers to deviate sending $\hat{s}_j = P_j^{\pi_t}$.

Subcase 2.2.2. Suppose that $s_j < P_j^{\pi_t}$. If $P_d^{\pi_t} < s_j$, then agent j prefers to deviate, sending $\hat{s}_j = P_d^{\pi_t}$. If $s_j < P_d^{\pi_t}$ and $s_j \neq s_k$, then agent j prefers to deviate sending $\hat{s}_j = s_j + 1$. If $s_j < P_d^{\pi_t}$ and $s_j = s_k$, then agent i prefers to deviate sending $\hat{s}_i = P_d^{\pi_t}$.

Subcase 2.2.3. Suppose that $s_j = P_j^{\pi_t}$. Then agent k prefers to deviate sending either $\hat{s}_k = P_d^{\pi_t}$ (if $s_j < P_k^{\pi_t}$) or $\hat{s}_k = s_j - 1$ (if $P_k^{\pi_t} < s_j$).

Claim 3. For all $i, j \in \{a, b, c\}$, $s_i \neq s_j$.

Suppose, on the contrary, that for some $i, j \in \{a, b, c\}$, $s_i = s_j$. Let $k \in \{a, b, c\}$ be such that $k \notin \{i, j\}$. By Claim 1, $s_i \neq s_k$, and by Claim 2, $s_i \neq n$ and $s_k \neq n$. We distinguish two cases.

Case 3.1. Suppose that $s_k < s_i$. Then $P_d^{g(\hat{s}_d, s-d)} = s_i \neq P_d^{\pi_i}$.

Subcase 3.1.1. Suppose that $P_k^{\pi_i} < P_d^{\pi_i}$. We distinguish two possibilities:

(3.1.1.1) Suppose that $s_i < P_d^{\pi_i}$. Then agent k prefers to deviate sending either $\hat{s}_k = P_d^{\pi_i}$ (if $P_d^{\pi_i} \neq n$) or $\hat{s}_k = s_i$ (if $P_d^{\pi_i} = n$).

(3.1.1.2) Suppose that $P_d^{\pi_i} < s_i$. If $P_i^{\pi_i} \neq s_i$, then agent i prefers to deviate sending either $\hat{s}_i = s_i - 1$ (if $s_i - 1 \neq s_k$) or $\hat{s}_i = n$ (if $s_i - 1 = s_k$). If $P_i^{\pi_i} = s_i$, then agent j prefers to deviate sending either $\hat{s}_j = s_j - 1$ (if $s_j - 1 \neq s_k$) or $\hat{s}_j = n$ (if $s_j - 1 = s_k$).

Subcase 3.1.2. Suppose that $P_d^{\pi_i} < P_k^{\pi_i}$. We distinguish two possibilities:

(3.1.2.1) Suppose that $s_i < P_k^{\pi_i}$. Then agent k prefers to deviate sending either $\hat{s}_k = P_k^{\pi_i}$ (if $P_k^{\pi_i} \neq n$) or $\hat{s}_k = s_i$ (if $P_k^{\pi_i} = n$).

(3.1.2.2) Suppose that $P_k^{\pi_i} \leq s_i$. Then agent i or agent j (depending on whether or not $P_i^{\pi_i} \neq s_i$) prefers to deviate in the same way as in (3.1.1.2).

Case 3.2. Suppose that $s_i < s_k$. Then $P_d^{g(\hat{s}_d, s-d)} = s_k \neq P_d^{\pi_i}$.

Subcase 3.2.1. Suppose that $P_k^{\pi_i} < P_d^{\pi_i}$. We distinguish two possibilities:

(3.2.1.1) Suppose that $s_k < P_d^{\pi_i}$. Then agent k prefers to deviate as in (3.1.1.1).

(3.2.1.2) Suppose that $P_d^{\pi_i} < s_k$. Then agent k prefers to deviate sending either $\hat{s}_k = s_k - 1$ (if $s_k - 1 \neq s_i$) or $\hat{s}_k = n$ (if $s_k - 1 = s_i$).

Subcase 3.2.2. Suppose that $P_d^{\pi_i} < P_k^{\pi_i}$. We distinguish three possibilities:

(3.2.2.1) Suppose that $P_k^{\pi_i} < s_k$. Then agent k prefers to deviate in the same way as in (3.2.1.2).

(3.2.2.2) Suppose that $s_k < P_k^{\pi_i}$. If $P_d^{\pi_i} < s_k$, then agent k prefers to deviate in the same way as in (3.2.1.2). If $s_k < P_d^{\pi_i}$, agent k prefers to deviate sending $\hat{s}_k = P_d^{\pi_i}$.

(3.2.2.3) Suppose that $s_k = P_k^{\pi_i}$. Then agent i prefers to deviate sending either $\hat{s}_i = s_k - 1$ (if $s_k - 1 \neq s_i$) or $\hat{s}_i = n$ (if $s_k - 1 = s_i$).

Claim 4. Suppose that for all $i, j \in \{a, b, c\}$, $s_i \neq s_j$. Then there is some $i \in \{a, b, c\}$ with $s_i = n$.

Suppose, on the contrary, that for some $i, j, k \in \{a, b, c\}$, we have $s_i < s_j < s_k < n$. Then $P_d^{g(\hat{s}_d, s-d)} = s_j \neq P_d^{\pi_i}$.

Case 4.1. Suppose that $P_j^{\pi_i} < P_d^{\pi_i}$.

Subcase 4.1.1. Suppose that $P_d^{\pi_i} < s_j$. Then agent j prefers to deviate sending either $\hat{s}_j = s_j - 1$ (if $s_j - 1 \neq s_i$) or $\hat{s}_j = n$ (if $s_j - 1 = s_i$).

Subcase 4.1.2. Suppose that $s_j < P_d^{\pi_i}$. Then agent j prefers to deviate sending either $\hat{s}_j = s_j + 1$ (if $s_j + 1 \neq s_k$) or $\hat{s}_j = s_i$ (if $s_j + 1 = s_k$).

Case 4.2. Suppose that $P_d^{\pi_i} < P_j^{\pi_i}$.

Subcase 4.2.1. Suppose that $P_j^{\pi_i} < s_j$. Then agent j prefers to deviate in the same way as in Subcase 4.1.1.

Subcase 4.2.2. Suppose that $s_j < P_j^{\pi_i}$. Then agent j prefers to deviate sending either $\hat{s}_j = P_d^{\pi_i}$ (if $P_d^{\pi_i} \neq s_i$) or $\hat{s}_j = n$ (if $P_d^{\pi_i} = s_i$).

Subcase 4.2.3. Suppose that $s_j = P_j^{\pi_i}$. If $s_j < P_k^{\pi_i}$, then agent k prefers to deviate sending either $\hat{s}_k = P_d^{\pi_i}$ (if $P_d^{\pi_i} \neq s_i$) or $\hat{s}_k = n$ (if $P_d^{\pi_i} = s_i$). If $P_k^{\pi_i} < s_j$, then agent k prefers to deviate sending either $\hat{s}_k = s_j - 1$ (when $s_j - 1 \neq s_i$) or $\hat{s}_k = n$ (when $s_j - 1 = s_i$).

Obviously, Claims 2 and 3 contradict Claim 4. ■

5. FINAL REMARKS

In this paper we have studied a class of problems in which a given society wants to elicit the truth from its agents. We have proved that obtaining the true ranking from dominant strategies is impossible and that implementation in Nash equilibrium can be done with simple mechanisms.

Relationship with the Literature. There are two papers related to ours: Balinski and Sönmez (1999) and Duggan and Martinelli (1998). The paper by Balinski and Sönmez analyzes a class of matching problems in which students are matched according to their preferences. A difference between our paper and the matching literature is that in the former professors have preferences over all student's scholarships, but in the latter each agent cares only about his own matching. In the paper by Balinski and Sönmez, students are assigned according to their preferences. In our paper, each professor wants his own student to get the most prestigious scholarship and the other students to be assigned according to the socially optimal outcome, so preferences alone cannot be decisive in assigning students.

The paper by Duggan and Martinelli analyzes voting by jury members. They analyze the Bayesian equilibrium of several voting rules in which the possible outcomes are to convict or not to convict. They show that the unanimity rule that is common to many judicial systems does not implement the optimal conviction policy. There are several differences with our paper: In our case, preferences are restricted to be moderately selfish, whereas in their case, the allocation space includes only two alternatives (to convict or not to convict), but there are several voters. As a consequence of these differences, the socially optimal choice correspondence is not implementable in our model in any equilibrium concept in the case of two agents (which corresponds to the case of two alternatives in Duggan and Martinelli).

Extensions. We note some extensions of our work that might be fruitful.

1. Suppose that the true ranking is observed imperfectly by the agents. For instance, some agents may be more able than others to discern particular characteristics of the true ranking. This topic has been studied in psychology in the area of aggregation of expert's opinions (see Mongin, 1984 and references there in). The paper by Duggan and Martinelli analyzes a special case of this when there are only two alternatives. A recent entry in this area is a paper by Krishna and Morgan (1999).

2. Consider the following class of preferences. Each agent classifies all agents in three groups: friends, indifferent, and enemies. Each agent always prefers an alternative in which friends are higher in the ranking. Among those alternatives giving the same position to the group of friends, he prefers those in which enemies are worse off in the ranking. Finally, among all alternatives that are indifferent according to the above criteria, he prefers the alternative in which the indifferent agents are placed according to the true ranking. Moderately selfish preferences are a special case of this class where the group of friends includes only one agent and there are no enemies. We do not know the conditions on this class of preferences under which the SOCF is implementable.

3. Finally, in some problems the alternative is not a ranking, but rather a list of scores, one for each agent (i.e., gymnastics, skating, etc.).

We hope that further work will clarify the possibilities for implementing the socially optimal alternative (whatever it is) in the situations outlined above. We also hope that this work may be useful in designing mechanisms that can be applied to replace the existing ones.

REFERENCES

- Aswal, N. and Sen, A. (1997). "Restricted Domains and the Gibbard-Satterthwaite Theorem," mimeo.
- Balinski, M. and Sönmez, T. (1999). "A Tale of Two Mechanisms: Student Placement," *J. Econ. Theory* **84**, 73-94.
- Duggan, J. and Martinelli, C. (1998). "A Bayesian Model of Voting in Juries," mimeo.
- Gibbard, A. (1973). "Manipulation of Voting Schemes: A General Result," *Econometrica* **41**, 587-601.
- Jackson, M. O. (1992). "Implementation in Undominated Strategies: A Look at Bounded Mechanisms," *Rev. Econ. Stud.* **59**, 757-775.
- Krishna, V. and Morgan, J. (1999). "A Model of Expertise," mimeo.
- McKelvey, R. D. (1989). "Game Forms for Nash Implementation of General Social Choice Correspondences," *Soc. Choice Welfare* **6**, 139-156.
- Mongin, P. (1984). "Consistent Bayesian Aggregation," *J. Econ. Theory* **66**, 313-351.

- Repullo, R. (1987). "A Simple Proof of Maskin Theorem on Nash Implementation," *Soc. Choice Welfare* **4**, 39-41.
- Saijo, T. (1988). "Strategy Space Reduction in Maskin's Theorem: Sufficient Conditions for Nash Implementation," *Econometrica* **56**, 693-700.
- Satterthwaite, M. A. (1975). "Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions," *J. Econ. Theory* **10**, 187-217.
- Williams, S. (1984). "Sufficient Conditions for Nash Implementation," mimeo, IMA, Minneapolis.