

Virtual Nash implementation with admissible support*

Olivier Bochet[†] and François Maniquet[‡]

August 2006
CORE Discussion Paper 2006/84

Abstract

A social choice correspondence (SCC) is virtually implementable if it is ε -close (in the probability simplex) to some (exactly) implementable correspondence. Abreu and Sen (1991) proved that, without restriction on the set of alternatives receiving strictly positive probability at equilibrium, every SCC is virtually implementable in Nash Equilibrium. We study virtual implementation when the supports of equilibrium lotteries are restricted. We provide a necessary and sufficient condition, imposing joint restrictions on SCCs and admissible supports. Then, we discuss how to construct supports. Finally, we study virtual implementation when the support is restricted to the efficient or individually rational alternatives.

*We thank seminar participants from Universitat Autònoma de Barcelona, University of Rochester, Society for Economic Design conference 2004, Social Choice and Welfare conference 2004, GREQAM, University of Maastricht and University of Caen for their comments and suggestions. Corresponding author: o.bochet@algec.unimaas.nl. Mailing address: Maastricht University, department of economics, P.O. Box 616. 6200 MD Maastricht, The Netherlands. Financial support from Belgium Science Policy is gratefully acknowledged.

[†]Maastricht University and CORE (UCL, Belgium).

[‡]CORE (UCL, Belgium).

1 Introduction

The two key concepts of Nash implementation are those of a social choice correspondence and a mechanism. A social choice correspondence (henceforth an SCC) associates each state of the world (that is, typically, each profile of individual preferences) with a set of selected alternatives, which may be considered as what a benevolent planner views as optimal for society. A mechanism (or, game form) specifies the list of strategies available to each individual agent, and an outcome function associating each list of individual strategies with an outcome, that is, a social alternative. An SCC is said to be Nash implementable if there exists a mechanism such that, in each state of the world, the set of selected alternatives coincides with the set of outcomes which are supported by a Nash equilibrium of the mechanism.

In a seminal paper, Maskin (1997, published in 1999) identified a necessary and almost sufficient condition for Nash implementation, which he called *monotonicity*. In spite of this breakthrough, the news was bad: *monotonicity* imposes severe restrictions on the class of Nash implementable SCCs. Many SCCs of interest fail to satisfy it. To list just a few, the Walrasian and the egalitarian-equivalent correspondences, the plurality and any other voting rules based on scoring functions do not satisfy *monotonicity*.

How restrictive *monotonicity* can be is illustrated by the following example. There are two agents 1 and 2, social alternatives a, b, c and d , and two states of the world θ and ϕ . Preferences of agents over the alternatives in the different states are described in the table below, which reads: in state θ , agent 1 strictly prefers a to b to c to d , etc.

θ		ϕ	
1	2	1	2
a	d	a	d
b	b	b	b
c	c	c	a
d	a	d	c

Let us consider the SCC f defined as follows: $f(\theta) = \{b\}$ and $f(\phi) = \{a\}$. It cannot be Nash implemented. Suppose, indeed, that there is a mechanism where alternative b is supported by a Nash equilibrium at θ : by deviating, agent 1 cannot obtain any better alternative than b . But the alternatives which are worse than b at θ remain worse at ϕ . Consequently, the same

profile of strategies keeps the Nash stability property at ϕ , whereas b is no longer viewed as a desirable alternative.

This example can also be described as follows: the only difference in preferences between θ and ϕ is agent 2's preference reversal between alternatives a and c . It is, consequently, the only information the planner needs, to know whether b or a must be selected. But this information cannot be revealed if the requirement is that b be supported by a Nash equilibrium at θ .

A way out of the limitations imposed by *monotonicity* was proposed by Matsushima (1988) and Abreu and Sen (1991) under the name of virtual implementation. It consists of allowing the outcome of the mechanism to be a lottery on the social alternatives, and requiring that the alternatives selected by the SCC come out with a probability arbitrarily large but possibly smaller than 1.

What makes virtual implementation so powerful is that, in equilibrium, every alternative receives positive probability. Let us come back to the above example. At the Nash equilibrium supporting b with an arbitrarily large probability, alternatives a, c and d receive some strictly positive probability. If the mechanism is constructed in such a way that, by deviating, agent 2 can slightly increase the probability of a at the expense of c , then this agent will not deviate in state θ , whereas he would in state ϕ , and b —with arbitrarily large probability—is no longer supported as an equilibrium, the desired outcome.

The virtual Nash implementation approach may seem to be the solution to the implementation problem, but we would like to point out the following issue. Any alternative can now be the outcome of the game, as it receives a strictly positive probability in the equilibrium lottery. However, these alternatives may be arbitrarily inefficient, unfair, or, in any meaningful sense, far from the selected alternative. We ask the following question. What happens if one restricts the class of mechanisms that can be used and imposes that the support of equilibrium lotteries belong to some admissible subset of the set of alternatives? Let us look back at the example. There is one and only one inefficient alternative, c . What would happen if we wish to remove it from the set of alternatives likely to become the outcome of the game? Also, what if we want to minimize the set of alternatives which receive a strictly positive probability at equilibrium?

The example illustrates the two main lessons we will reach in this paper. First, it may turn out to be impossible to remove some undesirable alternatives from the set of alternatives which may end up being selected. In the

example, indeed, it is impossible to remove c from the equilibrium lottery. Because the only preference reversal when going from θ to ϕ occurs between c and a with c preferred to a at θ , the mechanism must provide the opportunity to agent 2 to decrease the probability of a at the expense of c , and this requires that c receives some strictly positive probability at equilibrium.

The second lesson is that admissible supports may indeed be very narrow. Let us look at the following mechanism implementing f , where b_c , for instance, means that the outcome is b with probability $1 - \varepsilon$ and c with probability ε (for some arbitrarily small positive ε). Agent 1, whose preferences don't change, announces θ if 2 announces θ , and ϕ in the other case. Agent 2 has now a dominant strategy to announce the true state. This shows that very few additional alternative need to be added. In particular, only c needs to be given strictly positive probability at equilibrium, and only when the state of the world is θ .

		2	
		θ	ϕ
1	θ	b_c	b_a
	ϕ	c	a

These two lessons will follow from our main formal result. It consists of the definition of a property, called *extended monotonicity*, which we prove to be necessary and sufficient for virtual implementation of an SCC with an admissible support. It therefore imposes joint restriction on the SCC and the support.

The paper is organized as follows. In Section 2, we present the model and the basic definitions. We introduce the property of *extended monotonicity* and give the main result in Section 3, where we also provide some examples of SCCs and their supports. In Section 4, we discuss how to construct supports and minimal supports, and we highlight a difficulty. In Section 5, we give other examples showing that restricting the support to the set of efficient or individually rational alternatives may not be sufficient, thereby proving that undesirable alternatives may need to be inserted in the admissible support. We give concluding remarks in Section 6.

2 The set-up

We will develop our basic result under the assumptions that there are at least three agents, the number of social alternatives is finite and preferences over alternatives are strict. None of these assumptions are necessary, and we explain at the end of Section 3 how the result generalizes.

There is a finite set of agents $N = \{1, \dots, n\}$, $n \geq 3$ and a finite set of alternatives $A = \{1, \dots, m\}$. The set of admissible preference profiles over A is denoted by Θ . For any $\theta \in \Theta, i \in N$, i 's weak preferences are represented by the ordering $R_i(\theta)$. We denote by $P_i(\theta)$ and $I_i(\theta)$ the associated strict preference and indifference relations, respectively. We assume that preferences are strict, that is, for each $\theta \in \Theta$, each $a, b \in A$ and each $i \in N$,

$$[a R_i(\theta) b] \Rightarrow [a P_i(\theta) b \text{ or } a = b].$$

Define by $LC_i(\theta, a) = \{b \in A : a R_i(\theta) b\}$, the lower contour set of agent $i \in N$ at profile $\theta \in \Theta$ and alternative $a \in A$.

A *social choice correspondence* (SCC) $f : \Theta \rightarrow A$ associates to each preference profile a non-empty subset of social alternatives.

Our problem is to implement SCCs. As it is well known after Maskin's (1977, 1999) seminal paper, the following *monotonicity* property is necessary for implementation in Nash equilibrium. It requires that an enlargement of any agent's lower contour set at a selected alternative do not remove this alternative from the selection.

Definition 1 *An SCC f satisfies monotonicity if and only if for all $\theta, \phi \in \Theta$ and all $a \in f(\theta)$,*

$$[LC_i(\theta, a) \subseteq LC_i(\phi, a) \text{ for each } i \in N] \Rightarrow [a \in f(\phi)].$$

Monotonicity is also sufficient for Nash implementation, when it is complemented with the requirement, called no-veto power, that an alternative which is top ranked by all but one agents must be selected by the SCC.

A lot of SCCs of interest fail to satisfy *monotonicity* and, therefore, to be Nash implementable. It is therefore necessary to modify the implementation requirement. The idea of virtual implementation is to allow outcomes to be lotteries over alternatives, and to simply require that the desired alternative receive an arbitrarily large but strictly lower than 1 probability of being selected. Furthermore, the idea of virtual implementation with an admissible

support is to restrict the set of alternatives which receive a strictly positive probability at the equilibrium lottery. We therefore need to introduce the following notations and terminology.

Let us begin with our central concept. An *admissible support* h for an SCC f is a correspondence defined for each $\theta \in \Theta$ and each $a \in f(\theta)$ and satisfying the property that $a \subseteq h(\theta, a) \subseteq A$. It represents the set of alternatives to which the mechanism designer agrees to give positive probability at equilibrium. Let us emphasize that the admissible support is allowed to depend on the targeted alternative. This possibility is not used, however, if, for instance, all efficient alternatives are admissible. But the formulation we adopt is more general and allows us to define the admissible support, for instance, like the set of alternatives in a neighborhood of the targeted alternative.

Let $\mathcal{L} = \Delta^{m-1}$ denote the set of lotteries over elements of A . In lottery $\ell = (\ell_a)_{a \in A} \in \mathcal{L}$ alternative a occurs with probability ℓ_a . With a slight abuse of notation, we write a both for the alternative $a \in A$ and the lottery $\ell \in \mathcal{L}$ such that $\ell_a = 1$. The support of $\ell \in \mathcal{L}$, denoted suppl , is the set of alternatives receiving a strictly positive probability in ℓ : $\text{suppl} = \{a \in A \mid \ell_a > 0\}$.

We now need to consider that agents have preferences over lotteries. The set of admissible preference profiles over \mathcal{L} is Γ . For any $\gamma \in \Gamma, i \in N$, $R_i(\gamma), P_i(\gamma)$ and $I_i(\gamma)$ denote weak preference, strict preference and indifference relations over lotteries. Given $\theta \in \Theta$, two elements γ and δ of Γ are said to be elements of the set $\Sigma(\theta)$ if they are consistent with the same preferences over elements of A . We assume that preferences over lotteries are *monotonic in probabilities*, that is, shifts in probabilities to strictly preferred alternatives yield strictly preferred lotteries. Formally, for any $i \in N, k \in \{1, \dots, m\}$, let $p_{ik} : \Theta \rightarrow A$ be defined by

$$p_{ik}(\theta) = a \Leftrightarrow |\{b \in A \mid b R_i(\theta) a\}| = k,$$

that is, $p_{i1}(\theta)$ is the preferred alternative of agent i in state θ , $p_{i2}(\theta)$ her second preferred, etc. Then, for all $\gamma \in \Sigma(\theta)$, if two lotteries $\ell \equiv (\ell_a)_{a \in A}$ and $\ell' \equiv (\ell'_a)_{a \in A}$ are such that for each $k^* \in \{1, \dots, m\}$, $\sum_{k \leq k^*} \ell_{p_{ik}(\theta)} \geq \sum_{k \leq k^*} \ell'_{p_{ik}(\theta)}$, then $\ell R_i(\gamma) \ell'$ and $\ell P_i(\gamma) \ell'$ if one inequality is strict. Let us note that this assumption is much weaker than that of expected utility.

A *mechanism (or game form)* is a pair $G = (S, g)$ with $S = \prod_i S_i$, where S_i the strategy space of agent $i \in N$, and $g : S \rightarrow \mathcal{L}$ is the outcome function that associates a lottery to every profile of strategies. A typical strategy

profile is $s = (s_i)_{i \in N} \in S$. A game for G is a pair (G, γ) for some $\gamma \in \Gamma$. Denote the set of Nash equilibria of game (G, γ) by $NE(G, \gamma)$.

A mechanism $G = (S, g)$ is *ordinal* if the set of Nash equilibria only depends on agents' preferences over pure alternatives, that is, for each $\theta \in \Theta$, each $s \in S$ and all $\gamma, \delta \in \Sigma(\theta)$, $NE(G, \gamma) = NE(G, \delta)$. For the sake of simplicity, we confine our attention to ordinal mechanisms.¹ Abusing notation, we sometimes let $NE(G, \theta)$ denote the Nash equilibria of game (G, γ) , for any $\gamma \in \Sigma(\theta)$.

Rather than implementing an SCC itself, we will implement a lottery correspondence that is close to it. A *lottery correspondence* $f^\ell : \Theta \rightarrow \mathcal{L}$ associates to each preference profile a non-empty subset of lotteries. A lottery correspondence f^ℓ is (*ordinally*) *Nash implementable* if there exists an ordinal mechanism $G = (S, g)$ such that the Nash equilibrium outcomes of each game coincides with the selection of f^ℓ , that is, for all $\theta \in \Theta$, $f^\ell(\theta) = g(NE(G, \theta))$.

Let $d(., .)$ denote the Euclidean distance (in \mathcal{L}). An SCC f is ε -close to a lottery correspondence f^ℓ if for all $\theta \in \Theta$, there exists a bijection $\tau_\theta : f(\theta) \rightarrow f^\ell(\theta)$ such that for all $a \in f(\theta) : d(a, \tau_\theta(a)) \leq \varepsilon$.

An SCC f is *virtually Nash implementable with admissible support* h if for all $\varepsilon > 0$, there exists a Nash implementable lottery correspondence f_ε^ℓ that is ε -close to f , with associated τ_θ for each $\theta \in \Theta$, and such that for each $\theta \in \Theta$, each $a \in f(\theta) : \text{supp} \tau_\theta(a) \subseteq h(\theta, a)$.

3 A necessary and sufficient condition

Characterizing the class of SCCs and supports which are compatible with virtual implementation turns out to be surprisingly simple. The required property is just an extension of *monotonicity*: It requires that an enlargement of any agent's lower contour set at a selected alternative *or at any alternative receiving a strictly positive probability when the former alternative is implemented* does not remove this alternative from the selection *and does not remove any of the latter alternatives from the admissible support*.

Definition 2 *An SCC f satisfies extended monotonicity with respect to an admissible support h if and only if there exists H , such that for all $\theta \in \Theta$,*

¹However, accepting to look at mechanisms where $NE(G, \gamma) \neq NE(G, \delta)$ would not allow us to enlarge the set of implementable SCCs.

all $a \in f(\theta)$, $a \subseteq H(\theta, a) \subseteq h(\theta, a)$ and for all $\theta, \phi \in \Theta$, all $a \in f(\theta)$, and all $i \in N$,

$$[LC_i(\theta, b) \subseteq LC_i(\phi, b) \ \forall b \in H(\theta, a)] \Rightarrow [a \in f(\phi) \text{ and } H(\theta, a) \subseteq h(\phi, a)]$$

Correspondence H gathers all the alternatives which do receive a strictly positive probability: it may therefore not always coincide with h . We are now ready to state the main theorem of the paper.

Theorem 1 *An SCC f is virtually implementable with admissible support h if and only if f satisfies extended monotonicity with respect to h .*

The proof is relegated to the appendix. Let us note that if there is no restriction on the admissible support (that is, for all $\theta \in \Theta$ and all $a \in f(\theta)$, $h(\theta, a) = A$) then *extended monotonicity* is trivially satisfied by any SCC, as an enlargement of the lower contour set at each and every alternative means that preferences do not change. We are then back to the Abreu-Sen result. If, at the other extreme, the support is restricted to be composed of the selected alternative only (that is, for all $\theta \in \Theta$ and all $a \in f(\theta)$, $h(\theta, a) = a$), then *extended monotonicity* boils down to *monotonicity*. In such a case, notice that allowing us to use lotteries out of equilibrium does not increase the class of implementable SCCs, except that we can get rid of no-veto power.²

Let us now review how the above result generalizes when some of the assumptions we made are dropped.

If the number of agents is restricted to 2, then, the difficulty is exactly the same as for virtual Nash implementation. *Extended monotonicity* complemented with the *non-empty lower contour intersection property* as in Abreu and Sen (1991) is sufficient.

If preferences are not strict, then restrictions on the domain of preferences are needed. Our proofs in the appendix work, for instance, if we assume that each agent has a uniquely best alternative. Alternatively, we could use that there are pairs of alternatives such that for all preferences in the domain, the first alternative is always strictly preferred to the second one. Those conditions are trivially satisfied in all economic environments under the assumption of strictly monotonic preferences.

²See Bochet (2005a) for a more detailed discussion on the equivalence between *monotonicity* and (exact) Nash implementability.

If the number of alternatives is countable but not necessarily finite, then our result generalizes, provided the same adjustment as the one proposed by Abreu and Sen is made. It consists in ordering the alternatives and assigning them a decreasing sequence of probabilities which sum up to ε . If the number of alternatives is uncountable but preferences are continuous on a Euclidean space, then, again, the same adjustment as the one proposed by Abreu and Sen can be used: it consists of using a dense and countable subset of the alternatives and applying the same method as in the countable case.

Let us now exemplify these facts.

Example 1 *The plurality correspondence.*

The plurality correspondence selects all alternatives which appear most frequently at the top of the agents' rankings. Let us introduce the following rank function. For all $a \in A, i \in N, \theta \in \Theta$, the rank of a in the preference of i at θ is the number of alternatives which i deems at least as good as a :

$$r(a, i, \theta) = \# \{b \in A \mid b R_i(\theta) a\}.$$

The plurality score of a at θ , $s(a, \theta)$ is the number of agents who rank a first:

$$s(a, \theta) = \# \{i \in N \mid r(a, i, \theta) = 1\}.$$

The plurality correspondence is defined by

$$P(\theta) = \{a \in A \mid s(a, \theta) \geq s(b, \theta), \forall b \in A\}.$$

The plurality correspondence is not monotonic. It is even not subgame perfect implementable (see Abreu and Sen (1990)).³ Let us assume that preferences are strict, and let us consider $a \in P(\theta) \setminus P(\phi)$. Then, in the admissible support at $h(a, \theta)$ there needs to be an alternative b which is involved in a preference reversal from θ to ϕ : $b R_i(\theta) c$ and $c P_i(\phi) b$. This alternative may be different from a (otherwise P would satisfy *monotonicity*) but must necessarily be top ranked at θ (otherwise, $s(b, \theta) = s(b, \phi)$ so that $P(\theta) = P(\phi)$). Let $T(\theta)$ denote the set of alternatives which are top ranked by at least one agent:

$$T(\theta) = \{a \in A : s(a, \theta) \geq 1\}.$$

³From Jackson, Palfrey and Srivastava (1994), we know it can be implemented in Undominated Nash Equilibrium using bounded mechanisms.

We have thus shown the following: *If preferences are strict, the plurality correspondence P satisfies extended monotonicity with respect to the top correspondence T .*

One remarkable consequence of this fact is that when all agents prefer the same alternative, then it is the only alternative in the admissible support, so that the mechanism assigns it a probability of 1.

Example 2 *The Walrasian correspondence*

We need first to give a precise definition of A and Θ . There are L goods and Θ is composed of all preferences, defined over \mathbb{R}_+^L , which are continuous, strongly monotone and convex. The endowment of agent $i \in N$ is $\omega_i > 0$, and $\bar{\omega} = \sum_{i \in N} \omega_i \gg 0$.⁴ The set of alternatives, or feasible allocations is,

$$A = \left\{ a \in \mathbb{R}_+^{Ln} : \sum_{i \in N} a_i = \bar{\omega} \right\}.$$

The Walrasian correspondence W is defined as follows. For $\theta \in \Theta$, $W(\theta)$ is composed of all allocations a such that a_i maximizes $R_i(\theta)$ over the budget hyperplane $B_i(p) \equiv \{b_i \in \mathbb{R}_+^L : p \cdot b_i = p \cdot \omega_i\}$ for some $p \in \mathbb{R}_+^L$. Let $B(p) = \{b \in \mathbb{R}_+^{Ln} : b_i \in B_i(p)\}$.

The Walrasian correspondence is known to violate *monotonicity* when Θ is not restricted in such a way that Walrasian allocations are all interior to A . It is even not implementable in any game-theoretical solution concept when preferences are not differentiable (see, e.g., Bochet (2005b)). Let us assume that Θ only contains differentiable preferences and let us solve the boundary problem by constructing an admissible support. Let H^W be defined as follows. For $\theta \in \Theta$, $a \in W(\theta)$ with supporting price vector p , if $a \in \text{int}A$, then $H^W(a, \theta) = a$; if $a \in A \setminus \text{int}A$, then $H^W(a, \theta)$ is a countable and dense subset of a closed neighborhood of a in $B(p) \cap A$.

The only case where W fails to satisfy *monotonicity* is when $W(\theta) \ni a \in A \setminus \text{int}A$, $LC_i(\theta, a) \subseteq LC_i(\phi, a)$ for each $i \in N$, whereas $a \notin W(\phi)$ as illustrated in a two-agent-two-good economy in Figure 1, where only the preferences of agent 1 change. By continuity of preferences, there exists $b \in H^W(\theta, a)$ and $c \in A$ such that,

$$b_1 R_1(\theta) c_1 \text{ and } c_1 P_1(\phi) b_1,$$

⁴Our conventions for vector inequalities are \gg , $>$ and \geq .




Figure 1

which is sufficient to guarantee virtual implementation of W .

We have thus shown the following: *If preferences are continuous, strongly monotonic, convex and differentiable, the Walrasian correspondence satisfies extended monotonicity with respect to H^W .*

The above example shows how to save on the size of the support: as long as we know that the selected allocation remains selected after monotonic transformations of preferences, it is not necessary to provide opportunities to reveal preference reversal and the support can be minimal. On the contrary, when the Walrasian allocation is on the boundary of the feasible set, one needs a larger support.

Example 3 *The egalitarian equivalent correspondence*

In the same model as above (except that we now disregard individual endowments to view $\bar{\omega}$ as a social endowment), an allocation $a \in A$ is *egalitarian-equivalent* if there exists a reference bundle $a_0(\theta, a) \leq \bar{\omega}$ such that each agent is indifferent between what she gets and this reference bundle, that is, for all $i \in N$, $a_i I_i(\theta) a_0(\theta, a)$ (see Pazner and Schmeidler (1974)). The *egalitarian equivalent correspondence* $EQ : \Theta \rightarrow A$ associates to each economy $\theta \in \Theta$, its set of egalitarian-equivalent allocations $EQ(\theta)$. This correspondence is not monotonic (cfr Thomson (1999)). On the other hand, it is subgame perfect implementable (cfr Demange (1994) and Maniquet

(2002)). Let us show that it can be virtually implemented with a small support: each support must contain no more than $n + 1$ allocations. Let $H^{EQ}(\theta, a) = \left\{ a, \cup_{i \in N} \left(a_0(\theta, a), \left(\frac{\bar{\omega} - a_0(\theta, a)}{n-1} \right)_{j \neq i} \right) \right\}$, that is, each support is composed of the desired allocation together with allocations where each agent precisely gets the reference bundle (and the other agents get an equal share of the remaining resources).

We have: *If preferences are continuous, strongly monotonic, and convex, the egalitarian equivalent correspondence satisfies extended monotonicity with respect to H^{EQ} .*

Indeed, let us consider $\theta \in \Theta$ and $a \in EQ(\theta)$ with corresponding a_0 . Let $\phi \neq \theta$. If $a \notin EQ(\phi)$, then for at least one agent it is no longer true that $a_i I_i(\phi) a_0$ and we have a preference reversal between a_i and a_0 (note that the reversal can go either way). If $a \in EQ(\phi)$, then either the reference bundle is the same, in which case $H^{EQ}(\theta, a) = H^{EQ}(\phi, a)$, or it is different, in which case we have again at least one preference reversal between some a_i and a_0 .

Interestingly, the game form proposed by Demange (1994) to implement egalitarian equivalent allocations in subgame perfect equilibrium has $n + 1$ stages, and out of equilibrium allocations are similar to the ones we need to have in the admissible support.

Example 4 *The efficient and egalitarian equivalent correspondence*

Under our assumptions on preferences, efficient egalitarian equivalent allocations always exist and the efficient correspondence (selecting all efficient allocations) satisfies *monotonicity*. One may therefore wonder whether the result above can be used to construct the support which would allow us to virtually implement the efficient and egalitarian equivalent correspondence.

Actually, if a SCC f_1 satisfies *extended monotonicity* with respect to h_1 , and another SCC f_2 satisfies *extended monotonicity* with respect to h_2 , then if the SCC $f_1 \cap f_2$ is non-empty, it satisfies *extended monotonicity* with respect to $h_1 \cup h_2$. Indeed, as h_1 contains all the alternatives necessary to reveal a preference reversal indicating that an alternative is no longer in f_1 , and h_2 does the same for f_2 , combining h_1 and h_2 also gives a sufficient support. Given that the efficient correspondence satisfies *extended monotonicity* with respect to the trivial support $h(a, \theta) = a$ (by *monotonicity*), we have: *If preferences are continuous, strongly monotonic, and convex, the efficient and egalitarian equivalent correspondence satisfies extended monotonicity with respect to H^{EQ} .*

4 Constructing the supports

The result presented in the previous section can be used to answer two questions: first, given a SCC, which admissible supports would allow us to implement it; second, given some admissible support, which SCCs can be implemented with this support. We answer the former question in this section, and leave the latter for the next section.

Let us rewrite *extended monotonicity* in the following way.

Definition 3 *A SCC f satisfies extended monotonicity with respect to an admissible support h if and only if there exists H such that for all $\theta \in \Theta$, all $a \in f(\theta)$, $a \subseteq H(\theta, a) \subseteq h(\theta, a)$ and :*

1) *For all $\theta, \phi \in \Theta$ such that $a \in f(\theta) \setminus f(\phi)$, there exist $i \in N$, $b \in H(\theta, a)$ and $c \in A$ such that,*

$$bR_i(\theta)c \text{ and } cP_i(\phi)b.$$

2) *For all $\theta, \phi \in \Theta$ such that $a \in f(\theta) \cap f(\phi)$, either*

a) *there exists $i \in N$, $b \in H(\theta, a)$ and $c \in A$ such that $bR_i(\theta)c$ and $cP_i(\phi)b$;*
or

b) *$H(\theta, a) \subseteq h(\phi, a)$.*

Part 1 of the above definition illustrates the specificity of virtual Nash implementation compared to other notions of implementation. The key problem of implementation is to give agents the opportunity to reveal preference reversals in such a way that alternatives which are no longer selected by the SCC after the reversal (like alternative a in the definition) are no longer supported by an equilibrium. For any such reversal, like the one involving b and c in the condition, it is sufficient to add b to the admissible support. The possibility to decrease the probability assigned to b and to increase the probability assigned to c is just what is needed for the mechanism designer to know whether a should be selected or not. That is what makes virtual implementation so easy and successful. The lesson to draw from this part of the definition is that only one alternative needs to be added to the support (that is, needs to receive strictly positive probability at equilibrium) for each pair of preference profiles.

Let us apply this idea to identify the minimal support necessary to virtually implement an arbitrary f in the case of a finite Θ . Let $>$ be an order on A . Let $\theta \in \Theta$ and $a \in f(\theta)$. For each $\phi \in \Theta$ such that $a \notin f(\phi)$, identify

the pair $(b_{\theta\phi}, c_{\theta\phi})$ such that for some $i \in N$, $b_{\theta\phi}R_i(\theta) c_{\theta\phi}$ and $c_{\theta\phi}P_i(\theta) b_{\theta\phi}$ and no other preference reversal involving a $b_{\theta\phi}$ higher in the order on A exists. Then

$$h(\theta, a) \supseteq \cup_{\phi \in \Theta; a \notin f(\phi)} b_{\theta\phi}.$$

A minimal support must contain at least those elements.⁵

Part 2 of the definition illustrates the trade-off one faces when constructing support. Indeed, if there were no two profiles selecting the same alternative (that is, no $\theta, \phi \in \Theta$ such that $a \in f(\theta) \cap f(\phi)$), then the support could be constructed as described above. But as soon as an alternative is in the selection at several profiles, a choice has to be made. The first option is to add an alternative to the support, so that the relevant preference reversal is revealed and the designer is able to distinguish between profiles θ and ϕ . But it is also possible to save on the number of alternatives in the support. It is indeed what part 2) b) of the definition tells us: there is no need to reveal a preference reversal from θ to ϕ as there is nothing wrong if alternative a is still supported by the same equilibrium at ϕ as at θ , provided all the alternatives which receive positive probability at the equilibrium supporting a are still admissible at ϕ . We highlight next the difficulties and the choices that have to be made in constructing the supports.

In the construction of minimal support we begun above, defining $h(\theta, a) \supseteq \cup_{\phi \in \Theta; a \notin f(\phi)} b_{\theta\phi}$ may, indeed, not be sufficient. Let $\theta, \phi, \varphi \in \Theta$ be defined as follows:

$R_1(\theta) = R_1(\phi) = R_1(\varphi)$	$R_2(\theta)$	$R_2(\phi)$	$R_2(\varphi)$
c	e	a	e
a	a	b	a
b	b	d	c
d	c	c	b
e	d	e	d

Let f be defined by $f(\theta) = f(\phi) = \{a\}$, $f(\varphi) = \{c\}$. Let us assume that the order is $a > b > c > d > e$. We must have $a, b \in h(\theta, a)$, as $b R_2(\theta) c$ and $c P_2(\varphi) b$ is the only reversal between θ and φ . Also, $a \in h(\phi, a)$ and $c \in h(\varphi, c)$. If h is not larger than that, we face the following difficulty:

⁵This set depends on the order on A we started with. If minimality is computed with respect to set cardinality and not set inclusion, then the lowest cardinality should be evaluated across all possible orders.

the equilibrium strategy profile supporting a at θ cannot but be also an equilibrium strategy profile at ϕ (there are no preferences reversal involving a or b between θ and ϕ), whereas b is not admissible at ϕ . Therefore, either we impose $b \in h(\phi, a)$, and the same strategy profile may be an equilibrium at θ as at ϕ , or we impose $c \in h(\theta, a)$, so that the reversal $c R_2(\theta) d$ and $d P_2(\phi) c$ is used to break the undesired equilibrium at ϕ .

This illustrates the different possible strategy to construct admissible supports. It also shows that minimal supports are not unique, and, more importantly, that there is no single algorithm to construct them.

5 Fixing the support

The second question about virtual implementation with admissible support is the following one: assuming that an admissible support has been chosen, which SCCs can be virtually implemented with that support. We find the question particularly relevant if the admissible support is either the set of efficient or individually rational (that is, Pareto dominating a status quo) alternatives. Our main result in this section is the observation that both the efficient and the individually rational correspondences are invariant with respect to extended monotonic transformations, that is, if ϕ is obtained from θ by a monotonic transformation of the preferences at each and every efficient (resp., individually rational) alternative, then the set of efficient (resp., individually rational) alternatives is the same at ϕ as at θ . As we show below, this considerably simplifies the condition an SCC must satisfy to be virtually implementable with efficient or individually rational support. We then give examples.

Let us begin by formally defining the efficient and individually rational (with respect to some status quo $c \in A$) correspondences respectively. For each $\theta \in \Theta$,

$$PE(\theta) = \{a \in A \mid \forall b \in A : [b R_i(\theta) a, \forall i \in N] \Rightarrow [b I_i(\theta) a, \forall i \in N]\}.$$

For each $\theta \in \Theta$, and for some $c \in A$,

$$IR^c(\theta) = \{a \in A \mid a R_i(\theta) c, \forall i \in N\}.$$

Theorem 2 *The PE and IR^c correspondences are invariant with respect to extended monotonic transformations, that is, for all $\theta, \phi \in \Theta$,*

$$[LC_i(\theta, b) \subseteq LC_i(\phi, b), \forall b \in PE(\theta)] \Rightarrow [PE(\theta) = PE(\phi)], \text{ and}$$

$$[LC_i(\theta, b) \subseteq LC_i(\phi, b), \forall b \in IR^c(\theta)] \Rightarrow [IR^c(\theta) = IR^c(\phi)].$$

The proof is relegated to the appendix. The consequence of Theorem 2 is the following. In order to check whether an SCC is virtually implementable with admissible support PE or IR , we need only check Part 1) of definition 3, that is, that all selected alternatives are still selected after a change in preferences which extends the lower contour sets at each Pareto efficient (resp. individually rational) alternatives. Formally, we have the following definition.

Definition 4 *An SCC f satisfies PE (resp. IR)-extended monotonicity if and only if, for all $\theta, \phi \in \Theta$,*

$$[\forall b \in PE(\theta) \text{ (resp. } IR^c(\theta)), \forall i \in N, LC_i(b, \theta) \subseteq LC_i(b, \phi)] \Rightarrow [f(\theta) \subseteq f(\phi)].$$

Example 5 *The Plurality correspondence and other scoring correspondences*

Given that the top correspondence is a subcorrespondence of the efficient correspondence, example 1 above shows that the Plurality correspondence can be virtually implemented with efficient support (it does not make sense to check whether or not it can be virtually implemented with individually rational support as the plurality correspondence is not a subcorrespondence of the individually rational correspondence). We raise the question whether other scoring correspondences, that is, correspondences based on the maximization of the overall score of an alternative are virtually implementable with efficient support. Let us recall that Vartiainen (2003) shows that no scoring rule can be implemented in subgame perfect equilibrium. Sjöström (1993) shows that the Borda rule is not implementable in trembling-hand perfect equilibrium. Here also, the answer is negative.

Let $\lambda \in R_+^{|A|}$ be a weight vector, such that $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{|A|} = 0$. Given λ , the score of a at θ is equal to

$$s_\lambda(a, \theta) = \sum_{k=1}^{|A|} \lambda_k \# \{i \in N \mid r(a, i, \theta) = k\}.$$

The scoring correspondence S_λ associated to λ is defined by

$$S_\lambda(\theta) = \{a \in A \mid s_\lambda(a, \theta) \geq s_\lambda(b, \theta), \forall b \in A\}.$$

Let us note that if $\lambda_2 = 0$, S_λ is the plurality correspondence, if $\lambda_{|A|-1} = 1$, S_λ is the anti-plurality correspondence, and if $\lambda_k = 1 - \frac{k-1}{|A|-1}$, then S_λ is the Borda correspondence. Let us assume that $S_\lambda \neq P$, that is, $\lambda_2 > 0$. Then there exist two integers n_1, n_2 , such that $n_2 > n_1$ and $\frac{n_2 - n_1}{n_2} < \lambda_2$. Let $\theta \in \Theta$ be described as below, that is, n_1 agents have preferences $a P_i(\theta) b P_i(\theta) \dots P_i(\theta) c$, etc.

θ		ϕ	
n_1	n_2	n_1	n_2
$a \dots a$	$c \dots c$	$a \dots a$	$c \dots c$
$b \dots b$	$a \dots a$	$c \dots c$	$a \dots a$
.....
$c \dots c$	$b \dots b$	$b \dots b$	$b \dots b$

We have

$$s_\lambda(a, \theta) = n_1 + \lambda_2 n_2 > s_\lambda(c, \theta) = n_2,$$

and, clearly, $s_\lambda(a, \theta) > s_\lambda(b, \theta)$ for all $b \neq a, c$. Therefore, $S_\lambda(\theta) = \{a\}$. Now, $PE(\theta) = \{a, c\}$. Let $\phi \in \Theta$ be described as in the table, that is, ϕ is obtained from θ by an extended monotonic transformation at each alternative in $PE(\theta)$. By *PE-extended monotonicity*, we should have: $\{a\} \in S_\lambda(\phi)$, whereas $s_\lambda(a, \phi) = n_1 + \lambda_2 n_2 < s_\lambda(c, \phi) = n_2 + \lambda_2 n_1$, the desired contradiction.

Consequently, we have proven: *the plurality correspondence is the only scoring correspondence satisfying PE-extended monotonicity.*

This result confirms the special status, among scoring correspondences, of the plurality correspondence, when one takes incentives into account. Indeed, Jackson, Palfrey and Srivastava (1994) show that the plurality correspondence is boundedly implementable in undominated Nash equilibrium, though no other scoring rule is.

Example 6 *The Walrasian correspondence*

If the support is the set of efficient allocations, can we solve the boundary problem? The answer is negative. Indeed, if the set of efficient allocations is itself on the boundary, then an extended monotonic change in preferences from θ to ϕ may break the Walrasian property of an allocation without affecting the set of efficient allocations.

6 Concluding remarks

In this paper, we have characterized the set of pairs of SCCs and admissible supports such that the former can be virtually implemented with respect to the latter. Using this result, we have shown that admissible support sufficient to implement well-known correspondences may be surprisingly small, but we have also shown that undesirable allocations may have to belong to the support necessary to implement some SCCs.

Our conclusion is that, in spite of the Abreu and Sen general possibility result, virtual Nash implementation should not be viewed as the end of the story, and it would be interesting to explore which (especially small) supports can be used to implement well-known desirable SCCs.

References

1. Abreu D., Sen A., 1990 “Subgame Perfect Implementation: a Necessary and Almost Sufficient Condition” *Journal of Economic Theory*, 50, 285-299.
2. Abreu D., Sen A., 1991 “Virtual Implementation in Nash Equilibrium” *Econometrica*, 59, 997-1021.
3. Bochet O., 2005a “Nash Implementation with Lottery Mechanisms” *Forthcoming in Social Choice and Welfare*.
4. Bochet O., 2005b “Implementation of the Walrasian Correspondence: The Boundary Problem” *Forthcoming in International Journal of Game Theory*.
5. Demange G., 1994 “Implementing efficient egalitarian equivalent allocations.” *Econometrica*, 54, 1167–1177.
6. Hurwicz L., Maskin E. and Postlewaite A., 1995 “Feasible Nash Implementation” in Ledyard J.O. (ed), *The Economics of Informational Decentralization. Complexity, Efficiency and stability*. Kluwer Academic Publisher.
7. Jackson M., 1992 “Implementation in Undominated Strategies: a Look at Bounded Mechanisms” *Review of Economic Studies*, 59, 757-775.
8. Jackson M., Palfrey T. and Srivastava S., 1994 “Undominated Nash Implementation in Bounded Mechanisms” *Games and Economic Behavior*, 6, 474-501.
9. Maniquet F., 2002 “A study of proportionality and robustness in economies with a commonly owned technology” *Review of Economic Design*, 7, 1-15.

10. Maskin E., 1977 “Nash Equilibrium and Welfare Optimality” mimeo, published in *Review of Economic Studies*, 66 (1999), 23-38.
11. Matsushima H., 1988 “A New Approach to the Implementation Problem” *Journal of Economic Theory*, 45, 128-144.
12. Moore J. and Repullo R., 1988 “Subgame Perfect Implementation” *Econometrica*, 56, 1191-1120.
13. Sjöström T., 1993 “Implementation in Perfect Equilibria” *Social Choice and Welfare*, 10, 97-106.
14. Thomson W., 1999 “Monotonic extensions on economic domains” *Review of Economic Design*, 4, 13-33.
15. Vartiainen H., 2003 “Subgame Perfect Implementation: A Full Characterization” *Forthcoming in Journal of Economic Theory*.

7 Appendix

Proof of Theorem 1:

First part: *if f is virtually implementable with admissible support h , then f satisfies extended monotonicity with respect to h .*

Since f is virtually implemented with admissible support h , for all $\varepsilon > 0$, there exist a Nash implementable lottery correspondence f_ε^ℓ , a bijection τ_θ for each $\theta \in \Theta$, and a mechanism $G = (S, g)$ such that for all $\theta \in \Theta$, all $a \in f(\theta)$, there exists $s \in S$ such that $d(a, \tau_\theta(a)) \leq \varepsilon$, $g(s) = \tau_\theta(a)$ and $s \in NE(G, \theta)$. Let H be defined by: for all $\theta \in \Theta$, all $a \in f(\theta)$, $H(\theta, a) = \text{supp}\tau_\theta(a)$. As f is virtually implemented with admissible support h by G , $H(\theta, a) \subseteq h(\theta, a)$. Let us fix some $\theta \in \Theta$ and $a \in f(\theta)$. Let $\phi \in \Theta$ be derived from θ, a and $H(\theta, a)$ like in the definition of *extended monotonicity*. We need to prove that $s \in NE(G, \phi)$. Since $s \in NE(G, \theta)$, for all $\gamma \in \Sigma(\theta)$, all $i \in N$, all $s'_i \in S_i$, we have $g(s) R_i(\gamma) g(s'_i, s_{-i})$. Let us choose some s'_i arbitrarily. Let $z' = g(s'_i, s_{-i})$. By the way ϕ is derived from θ , we have that,

$$LC_i(\theta, a) \subseteq LC_i(\phi, a) \text{ for each } i \in N \text{ and each } a \in \text{supp}g(s)$$

These expansions imply that, in the probability simplex,

$$LC_i(\gamma, g(s)) \subseteq LC_i(\lambda, g(s)) \text{ for each } i \in N, \text{ each } \gamma \in \Sigma(\theta) \text{ and each } \lambda \in \Sigma(\phi).$$

Therefore,

$$g(s)R_i(\gamma)g(s') \Rightarrow g(s)R_i(\lambda)g(s') \text{ for each } i \in N, \text{ each } \gamma \in \Sigma(\theta) \text{ and each } \lambda \in \Sigma(\phi).$$

As $z' = g(s'_i, s_{-i})$ and s'_i has been chosen arbitrarily, this proves that $s \in NE(G, \phi)$. Therefore, there is $b \in f(\phi)$ such that $d(b, \tau_\theta(a)) \leq \varepsilon$. It is clear that for small enough ε , this implies $b = a$, so that $a \in f(\phi)$. Finally, since $g(s) \in f_\varepsilon^\ell(\phi) \subseteq h(\phi, a)$, $H(\theta, a) \subseteq h(\phi, a)$.

Second part: *if f satisfies extended monotonicity with respect to h , then it is virtually Nash implementable with admissible support h .*

First, given $\theta \in \Theta$ and $a \in f(\theta)$, let $\tilde{\ell}(H(\theta, a)) \in \mathcal{L}$ be defined by: for all $b \in H(\theta, a)$, $\tilde{\ell}_b = \frac{1}{|H(\theta, a)|}$. We construct $G = (S, g)$ as follows.

For each $i \in N$, $S_i = \Theta \times A \times \{(c, b) \in A^2 : c \neq b\} \times \mathbb{N}_{++}$. A typical strategy is $s_i = (\theta^i, a^i, c^i, b^i, n^i)$. Denote by i^* the agent with the lowest index among those who announces the highest integer, i.e. $i^* \equiv \min \{i \in N : n^i \geq n^j \forall j \in N\}$. For $s = (s_i)_{i \in N}$, $g(s)$ is defined as follows.

Rule 1: If for each $i \in N$, $s_i = (\theta^i, a^i, \dots, \dots) = (\theta, a, \dots, \dots)$ and $a \in f(\theta)$, then

$$g(s) = (1 - \varepsilon)a + \varepsilon\tilde{\ell}(H(\theta, a)).$$

Rule 2: If there exists $i \in N$ such that for all $j \neq i$, $s_j = (\theta^j, a^j, \dots, \dots) = (\theta, a, \dots, \dots)$, $a \in f(\theta)$, then,

2a) $g(s) = (1 - \varepsilon)a + \varepsilon\tilde{\ell}(H(\theta, a)) + \frac{\varepsilon}{2|H(\theta, a)|+1}(c^i - b^i)$ if $c^i \in LC_i(\theta, b^i)$ and $b^i \in H(\theta, a)$, and

2b) $g(s) = (1 - \varepsilon)a + \varepsilon\tilde{\ell}(H(\theta, a))$ otherwise.

Rule 3: In all other cases, $g(s) = (1 - \frac{1}{n^{i^*}+1})c^{i^*} + \frac{1}{n^{i^*}+1}b^{i^*}$.

We show that this game form implements the lottery correspondence f_ε^ℓ defined by:

$$\ell \in f_\varepsilon^\ell(\theta) \Leftrightarrow \exists a \in f(\theta) : \ell = (1 - \varepsilon)a + \varepsilon\tilde{\ell}(H(\theta, a)).$$

Given that f_ε^ℓ is ε -close to f , this will prove the claim. The proof is divided in two steps.

Step 1: For all $\phi \in \Theta$, $\ell \in f_\varepsilon^\ell(\phi)$, there exists $s^* \in NE(G, \phi)$ such that $g(s^*) = \ell$ and $\text{supp}g(s^*) \subseteq h(\phi, \tau_\phi^{-1}(\ell))$. Suppose the true profile is ϕ and

$a \in f(\phi)$. Each agent $i \in N$ announcing $s_i^* = (\phi, a, c, b, 1)$ —for a arbitrary $c \neq b$ —is a Nash Equilibrium of G . By deviating to $s_i' \neq s_i^*$, an agent $i \in N$ can only trigger rule 2, decrease the probability of an alternative $b \in H(\phi, x)$ and increase the probability of a worse alternative c . Therefore, $s^* \in NE(G, \phi)$ and $\text{supp}g(s^*) = H(\phi, a) \subseteq h(\phi, a)$.

Step 2: For all $\phi \in \Theta, s^* \in NE(G, \phi), g(s^*) \in f_\varepsilon^\ell(\phi)$ and $\text{supp}g(s^*) \subseteq h(\phi, \tau_\phi^{-1}(g(s^*)))$. Suppose the true profile is $\phi \in \Theta$.

1) There is no equilibria under rule 2a or 3.

At any strategy profile under either rule 2 or 3, two different alternatives receive strictly positive probability (even if b^i or $c^i = a$). Under the assumption that there is a uniquely best alternative for each agent, the outcome lottery cannot be the preferred lottery of any agent. Some agents, therefore, can profitably deviate.

2) Assume $s^* \in NE(G, \phi)$ under rule 2b and for all $j \neq i : s_j^* = (\theta, a, \dots)$. By the same reasoning as above, that cannot happen if two alternatives or more get strictly positive probability. Therefore, $g(s^*) = a$, so that $H(\theta, a) = \{a\}$. Any $j \neq i$ could have deviated under rule 3. Therefore, $p_{j1}(\phi) = a$. Also, i does not deviate, so that $LC_i(\phi, a) \subseteq LC_i(\theta, a)$. Given that $a \in f(\theta)$, by *extended monotonicity*, $a \in f(\phi)$.

3) Assume $s^* \in NE(G, \phi)$ under rule 1 and for all $i \in N : s_i^* = (\theta, a, \dots)$. By announcing $s^i = (\theta^i, x^i, c^i, b^i, n^i)$ such that $\theta \neq \theta^i, c^i \in LC_i(\theta, b_j), b^i \in H(\bar{\theta}, \bar{x})$, any i could have obtained (under rule 2a) $(1 - \varepsilon)a + \varepsilon \tilde{\ell}(H(\theta, a)) + \frac{\varepsilon}{2|H(\theta, a)|+1}(c^i - b^i)$. This implies that $LC_i(\theta, b) \subseteq LC_i(\phi, b) \forall b \in H(\theta, a)$. Since f satisfies *extended monotonicity*, we obtain that $a \in f(\phi)$ and $H(\theta, a) \subseteq h(\phi, a)$.

Q.E.D.

Proof of Theorem 2:

Consider the *PE* correspondence. The argument runs in the same way for *IR*. Let $\theta \in \Theta$. Let $\phi \in \Theta$ be such that for all $b \in PE(\theta)$, all $i \in N, LC_i(b, \theta) \subseteq LC_i(b, \phi)$. Therefore, for all $b, c \in PE(\theta)$, all $i \in N, b R_i(\theta) c \Leftrightarrow b R_i(\phi) c$, whereas for all $b \in PE(\theta), c \in A \setminus PE(\theta), b R_i(\theta) c \Rightarrow b R_i(\phi) c$. Consequently, $PE(\theta) \subseteq PE(\phi)$. Finally, if $b \in A \setminus PE(\theta)$, then there exists $c \in PE(\theta)$ such that for all $i \in N, c R_i(\theta) b$ and $c P_j(\theta) b$ for some $j \in N$. Consequently, for all $i \in N, c R_i(\phi) b$ and $c P_j(\phi) b$ so that $b \in A \setminus PE(\phi)$.

Q.E.D.