

Bhaskar Dutta · Hans Peters · Arunava Sen

Strateg -proof cardinal decision schemes

Received: 26 January 2005 / Accepted: 17 February 2006 / Published online: 17 May 2006
© Springer-Verlag 2006

Abstract This paper analyses strategy-proof mechanisms or decision schemes which map profiles of cardinal utility functions to lotteries over a finite set of outcomes. We provide a new proof of Hylland's theorem which shows that the only strategy-proof cardinal decision scheme satisfying a weak unanimity property is the random dictatorship. Our proof technique assumes a framework where individuals can discern utility differences only if the difference is at least some fixed number which we call the *grid size*. We also prove a limit random dictatorship result which shows that any sequence of strategy-proof and unanimous decision schemes defined on a sequence of decreasing grid sizes approaching zero must converge to a random dictatorship.

1 Introduction

The classic results of Gibbard (1973) and Satterthwaite (1975) have shown that unless preferences are restricted, the only decentralized mechanism which induces truth-telling behaviour by individual agents is the dictatorial one. This impossibility

We are most grateful to an Associate Editor and two referees for very helpful comments on an earlier version of the paper.

B. Dutta (✉)
University of Warwick, Coventry, UK
E-mail: b.dutta@warwick.ac.uk

H. Peters
Department of Quantitative Economics, University of Maastricht,
P.O. Box 616, 6200 MD Maastricht, The Netherlands
E-mail: h.peters@ke.unimaas.nl

A. Sen
Indian Statistical Institute, New Delhi, India
E-mail: asen@isid.ac.in

result has induced a huge literature which analyzes the possibility of constructing *strategy-proof* mechanisms under various alternative frameworks. One variant, due to Gibbard (1977, 1978), which is the main focus of this paper is the extension of the original impossibility result to mechanisms which assign a *probability distribution* over the set of feasible outcomes for each profile of preferences. Gibbard (1977) characterized the class of such strategy-proof probabilistic mechanisms or *decision schemes*. He showed that a strategy-proof decision scheme must be a convex combination of *duples* and *unilaterals*. A duple is a mechanism which assigns positive probability to at most two alternatives, the pair of alternatives being independent of the profile of preferences, while a unilateral is one where the preference ordering of a single individual dictates the social lottery over feasible alternatives.¹

Such mechanisms need not satisfy even a weak form of efficiency. That is, even if all individuals unanimously prefer an alternative a to all other alternatives, the mechanism need not assign a probability of one to a . The only strategy-proof mechanisms satisfying even this weak form of efficiency are *random dictatorships*, in which each individual is assigned a fixed probability of being a dictator – fixed in the sense that these probabilities are independent of the preference profile. Duggan (1996) and Nandebam (1998) provide alternative proofs of the random dictatorship result, while Dutta et al. (2002) show that the random dictatorship result holds even if the feasible set of alternatives is some convex set in \mathfrak{R}^k (with $k > 1$), and preferences are strictly convex and continuous with a unique peak.²

In the original Gibbard (1977, 1978) framework, the decision scheme used only *ordinal* information about individual preferences. However, Gibbard assumed that individual preferences were represented by von Neumann-Morgenstern utility functions since these functions were used to rank alternative probability distributions. Thus, the assumption that the decision scheme can use only ordinal information about preferences imposed a strong invariance requirement on the aggregation rule. In order to appreciate the strength of the invariance requirement, we point out that strategy-proof *ordinal* decision schemes must satisfy a “local” property. That is, suppose that a voter changes her preference by “switching” two contiguous alternatives. In the ordinal context, strategy-proofness will immediately imply that only the probabilities of the two alternatives being switched are affected. This is a property with strong implications and considerably simplifies the task of characterizing strategy-proof ordinal decision schemes. In contrast, if a strategy-proof decision scheme utilizes *cardinal* information, then a change in the utility of a single alternative for a voter could in principle, have a “global” impact, that is, the probability of *all* alternatives could be affected. This makes the analysis in the cardinal model far more difficult.

Despite this difficulty, Hylland (1980) in an important and regrettably unpublished paper, showed that the random dictatorship result holds even if the decision scheme is allowed to use *cardinal* information. In this paper, we have two main objectives. First, we provide an alternative and more transparent proof of Hylland’s theorem. Second, we consider a framework where essentially individuals cannot discern infinitesimally small differences in utility. In particular, we assume that if an alternative a is strictly preferred to another alternative b , then the utility difference

¹ See Barberà (1978, 1979) for related characterizations of strategy-proof probabilistic mechanisms.

² See also Ehlers et al. (2002).

between a and b is *at least* some fixed number which we refer to as the grid size. We construct an example to show that the random dictatorship result no longer holds when individual utility functions satisfy this additional restriction. We then analyze the consequences of gradually reducing the grid size. That is, we consider an arbitrary sequence of strategy-proof and unanimous decision schemes defined on a sequence of decreasing grid sizes approaching zero. We obtain a ‘limit’ random dictatorship result in the sense that the sequence of such decision schemes must converge to a random dictatorship for all profiles for which the limit exists.

Recently and independently of our work, Nandeibam (2004) has provided another proof of the Hylland result. We compare our proof with those of Hylland and Nandeibam towards the end of Sect. 3.

2 The model

Let $A = \{a_1, a_2, \dots, a_M\}$ be a finite set of alternatives, with $M \geq 3$. A *lottery* λ is a probability distribution over the set A , and can be identified with an M -vector whose j th component λ_j denotes the probability that λ assigns to $a_j \in A$. Clearly every component of λ is non-negative and the sum of the components is 1. The set of lotteries is denoted by \mathcal{L} .

The set of voters will be denoted by $I = \{1, 2, \dots, N\}$. Each voter i has a preference ordering R_i over the elements of the set A . The ordering R_i is represented by an *admissible* utility function u_i , which is unique up to affine transformations. We normalize utility functions by assuming that the utility of the maximal element, which is assumed to be *unique*,³ is one, while the utility of the worst element is zero. We do not require distinct alternatives to have distinct utility levels (i.e. it is not required that R_i is a strict ordering).

Let \mathcal{U} denote the set of admissible utility functions. We will use $\tau(u_i)$ to refer to the maximal element of utility function u_i .⁴

A utility profile is an N -tuple $(u_1, u_2, \dots, u_N) \in \mathcal{U}^N$. Let u denote the utility profile (u_1, \dots, u_N) , and (u'_i, u_{-i}) denote the profile $(u_1, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_N)$.

De nition 1 A *cardinal decision scheme (CDS)* is a mapping $\phi: \mathcal{U}^N \rightarrow \mathcal{L}$.

A CDS utilizes cardinal information in individuals’ utility functions and specifies a probability distribution over the set of alternatives for each profile of utility functions. We let $\phi_j(u)$, $j = 1, 2, \dots, M$ denote the probability on alternative a_j in the lottery $\phi(u)$.

A CDS which only utilizes ordinal information about individual utility functions will be called an ordinal decision scheme (ODS).

Two admissible utility functions u_i, u'_i are *ordinally equivalent* if for all $a_k, a_j \in A$, $u_i(a_j) \geq u_i(a_k)$ iff $u'_i(a_j) \geq u'_i(a_k)$. Similarly, two utility profiles u and u' are *ordinally equivalent* if each pair u_i, u'_i is *ordinally equivalent*.

De nition 2 An *ODS* is a CDS ϕ with the property that $\phi(u) = \phi(u')$ whenever u and u' are *ordinally equivalent*.

³ See Remark 1 in the next section.

⁴ In Sect. 4, we will impose an additional restriction on admissible utility functions – we will assume that the minimal difference in utility levels of alternatives which have different utilities is at least some $\eta > 0$.

Different concepts of efficiency can be associated with decision schemes. One concept which has been used is that of *ex post efficiency*.⁵

De nition 3 A CDS ϕ is *ex post efficient* if for all $a_j, a_k \in A$ and for all admissible utility profiles u , $\phi_k(u) = 0$ if $u_i(a_j) > u_i(a_k)$ for all $i \in I$.

An *ex post efficient* CDS ensures that a Pareto non-optimal alternative is never assigned positive probability. A considerably weaker condition is that of Unanimity.

De nition 4 A CDS ϕ satisfies *Unanimity* if for all $a_j \in A$ and for all admissible utility profiles u , $\phi_j(u) = 1$ if $\tau(u_i) = a_j$ for all $i \in I$.

Unanimity simply requires that if an alternative is best for all individuals, then it should be assigned probability one.

Random dictatorships are an important class of ordinal decision schemes. These are rules in which each individual has a fixed probability (that is, independent of the utility profile) of being a dictator. More formally,

De nition 5 The CDS is a *random dictatorship* if there exist non-negative real numbers $\beta_1, \beta_2, \dots, \beta_N$ with $\sum_i \beta_i = 1$ such that for all $u \in \mathcal{U}^N$ and $a_j \in A$,

$$\phi_j(u) = \sum_{\{i|\tau(u_i)=a_j\}} \beta_i$$

We assume that individuals rank alternative lotteries in terms of expected utility.

De nition 6 A CDS is *manipulable* by an individual $i \in I$ at $u \in \mathcal{U}^N$ via $u'_i \in \mathcal{U}$ if $\sum_{j=1}^M u(a_j)\phi_j(u'_i, u_{-i}) > \sum_{j=1}^M u(a_j)\phi_j(u)$.

De nition 7 A CDS is *strategy-proof (SP)* if it is not manipulable by any voter at any profile.

Thus, a CDS is *strategy-proof* if no voter can strictly gain in terms of expected utility by misrepresenting her true preferences.

3 The Hlland theorem

An example of a *strategy-proof* decision scheme is the *random dictatorship*. If i is the dictator, then the alternative which is first in i 's preference ordering is chosen with probability one. Since the probability of any voter i being a dictator is independent of the profile of preferences, it is easy to see that no individual has an incentive to misreveal preferences.

The random dictatorship in which each individual has an equal chance of being the dictator is obviously anonymous and efficient – the voting scheme only puts positive weight on alternatives which are Pareto optimal. This might seem to suggest that this random dictatorship provides a positive resolution of the dilemma posed by the Gibbard–Satterthwaite result – an equal distribution of power is consistent with efficiency and truthful revelation of preferences. Unfortunately, random dictatorships possess an undesirable property, as shown in the following example.

⁵ See for instance Gibbard (1977), Duggan (1996), Nandeibam (1998).

Example 1 Let $|I| = 1000$, and $A = \{a_1, \dots, a_{1001}\}$. Consider a profile P such that for each individual i , $a_i P_i a_{1001} P_i a$ for all $a \in A \setminus \{a_i, a_{1001}\}$. Although every individual considers a_{1001} as the second most preferred alternative, and no two individuals agree on what is the best alternative in A , a random dictatorship must assign zero probability to a_{1001} .

This example provides a motivation to search for other strategy-proof decision schemes. Unfortunately, Hylland proved that random dictatorships constitute the only class of unanimous and strategy-proof cardinal decision schemes. In this section, we provide a relatively simple proof of Hylland's theorem, which is stated below.

Theorem 1 *A CDS satisfies strategy-proofness and unanimity if and only if it is a random dictatorship.*

Proof It is clear that a random dictatorship satisfies unanimity and is also strategy-proof. We prove the converse.

Step 1 We first show that for $|N| = 2$, a unanimous and strategy-proof CDS ϕ is a random dictatorship.

In the proof of this step, for $k, j \in \{1, \dots, M\}$ with $k \neq j$ and a positive

Proof Suppose that the claim is false. Assume w.l.o.g. that $\lim_{\eta \rightarrow 0} \phi_k(u_1, u_{kj}^\eta) = \lambda_k > \lambda'_k = \lim_{\eta \rightarrow 0} \phi_k(v_1, u_{kj}^\eta)$. Observe that Claim 1 implies that $\lim_{\eta \rightarrow 0} \sum_t u_1(a_t) \phi_t(u_1, u_{kj}^\eta) = 1 - \lambda_k + u_1(a_k) \lambda_k$ and $\lim_{\eta \rightarrow 0} \sum_t u_1(a_t) \phi_t(v_1, u_{kj}^\eta) = 1 - \lambda'_k + u_1(a_k) \lambda'_k$. Therefore $\lim_{\eta \rightarrow 0} \sum_t u_1(a_t) (\phi_t(v_1, u_{kj}^\eta) - \phi_t(u_1, u_{kj}^\eta)) = (1 - u_1(a_k)) (\lambda_k - \lambda'_k)$. But the RHS of this expression is strictly positive by assumption. Therefore there exists η small enough such that $\sum_t u_1(a_t) \phi_t(v_1, u_{kj}^\eta) > \sum_t u_1(a_t) \phi_t(u_1, u_{kj}^\eta)$. This implies that voter 1 can manipulate ϕ at (u_1, u_{kj}^η) via v_1 which contradicts strategy-proofness of ϕ . \square

Let $u_{jk}^\eta, u_{kj}^\eta \in \mathcal{U}$ and let $u_2, u_1 \in \mathcal{U}$ with $\tau(u_1) = a_j$ and $\tau(u_2) = a_k$.

Claim 3 $\lim_{\eta \rightarrow 0} \phi_j(u_{jk}^\eta, u_2) = \lim_{\eta \rightarrow 0} \phi_j(u_1, u_{kj}^\eta)$.

Proof Let $\lim_{\eta_1 \rightarrow 0} \phi_j(u_{jk}^{\eta_1}, u_2) = \lambda_j$ and let $\lim_{\eta_2 \rightarrow 0} \phi_j(u_1, u_{kj}^{\eta_2}) = \lambda'_j$. According to Claim 2, $\lim_{\eta_2 \rightarrow 0} \phi_j(u_{jk}^{\eta_1}, u_{kj}^{\eta_2}) = \lambda'_j$ for all η_1 . Therefore

$$\lim_{\eta_1, \eta_2 \rightarrow 0} \phi_j(u_{jk}^{\eta_1}, u_{kj}^{\eta_2}) = \lim_{\eta_1 \rightarrow 0} \lambda'_j = \lambda'_j.$$

But Claim 2 also implies that $\lim_{\eta_1 \rightarrow 0} \phi_j(u_{jk}^{\eta_1}, u_{kj}^{\eta_2}) = \lambda_j$ for all η_2 . Therefore

$$\lambda'_j = \lim_{\eta_1, \eta_2 \rightarrow 0} \phi_j(u_{jk}^{\eta_1}, u_{kj}^{\eta_2}) = \lim_{\eta_2 \rightarrow 0} \lambda_j = \lambda_j$$

which is what we have to prove. \square

Let $a_j, a_k, a_s, a_t \in A$ with $a_j \neq a_k$ and $a_s \neq a_t$. Let u_1 and v_1 be admissible utility functions such that $\tau(u_1) = a_j$ and $\tau(v_1) = a_s$.

Claim 4 $\lim_{\eta \rightarrow 0} \phi_j(u_1, u_{kj}^\eta) = \lim_{\eta \rightarrow 0} \phi_s(v_1, u_{ts}^\eta)$.

Proof We know from Claim 2 that $\lim_{\eta \rightarrow 0} \phi(u_1, u_{kj}^\eta)$ does not depend on u_1 as long as the first-ranked alternative in u_1 is a_j . We can therefore denote this limit w.l.o.g. as $\lambda_j(j, k)$. So we have to prove that $\lambda_j(j, k) = \lambda_s(s, t)$. We consider two cases.

Case I $s \neq k$.

We will first prove that $\lambda_j(j, k) = \lambda_s(s, k)$.

Let δ, ϵ and γ be positive numbers and let v_1^ϵ be an admissible utility function with $\tau(v_1^\epsilon) = a_s$, $v_1^\epsilon(a_j) = 1 - \epsilon$ and $v_1^\epsilon(a_l) \leq \epsilon$ for all $a_l \neq a_s, a_j$.

Now consider voter 1 in the profile $(v_1^\epsilon, u_{ks}^\gamma)$. Her maximal expected utility from truth-telling is $\phi_s(v_1^\epsilon, u_{ks}^\gamma) + (1 - \epsilon) \phi_j(v_1^\epsilon, u_{ks}^\gamma) + \epsilon (1 - \phi_s(v_1^\epsilon, u_{ks}^\gamma) - \phi_j(v_1^\epsilon, u_{ks}^\gamma))$.

If she announces u_{jk}^δ instead her minimal expected utility is $\phi_s(u_{jk}^\delta, u_{ks}^\gamma) + (1 - \epsilon)\phi_j(u_{jk}^\delta, u_{ks}^\gamma)$. Since ϕ is strategy-proof, we have

$$\begin{aligned} & \phi_s(v_1^\epsilon, u_{ks}^\gamma) + (1 - \epsilon)\phi_j(v_1^\epsilon, u_{ks}^\gamma) + \epsilon(1 - \phi_s(v_1^\epsilon, u_{ks}^\gamma) - \phi_j(v_1^\epsilon, u_{ks}^\gamma)) \\ & \geq \phi_s(u_{jk}^\delta, u_{ks}^\gamma) + (1 - \epsilon)\phi_j(u_{jk}^\delta, u_{ks}^\gamma). \end{aligned}$$

Since the inequality above is true for all δ, ϵ and γ , we can take limits to obtain

$$\begin{aligned} & \lim_{\epsilon, \gamma, \delta \rightarrow 0} (\phi_s(v_1^\epsilon, u_{ks}^\gamma) + (1 - \epsilon)\phi_j(v_1^\epsilon, u_{ks}^\gamma) + \epsilon(1 - \phi_s(v_1^\epsilon, u_{ks}^\gamma) - \phi_j(v_1^\epsilon, u_{ks}^\gamma))) \\ & \geq \lim_{\epsilon, \gamma, \delta \rightarrow 0} (\phi_s(u_{jk}^\delta, u_{ks}^\gamma) + (1 - \epsilon)\phi_j(u_{jk}^\delta, u_{ks}^\gamma)). \end{aligned}$$

Observe that Claims 2 and 3 imply that $\lim_{\delta \rightarrow 0} \phi_j(u_{jk}^\delta, u_{ks}^\gamma) = \lambda_j(j, k)$ and $\lim_{\gamma \rightarrow 0} \phi_s(v_1^\epsilon, u_{ks}^\gamma) = \lambda_s(s, k)$. Also, Claim 1 implies that $\lim_{\delta \rightarrow 0} \phi_s(u_{jk}^\delta, u_{ks}^\gamma) = 0$ and $\lim_{\gamma \rightarrow 0} \phi_j(v_1^\epsilon, u_{ks}^\gamma) = 0$. Therefore the inequality above reduces to

$$\lim_{\epsilon \rightarrow 0} (\lambda_s(s, k) + \epsilon(1 - \lambda_s(s, k))) \geq \lim_{\epsilon \rightarrow 0} (1 - \epsilon)\lambda_j(j, k)$$

So $\lambda_s(s, k) \geq \lambda_j(j, k)$. By reversing the roles of a_s and a_j we also have the reverse inequality, and thus $\lambda_s(s, k) = \lambda_j(j, k)$.

Define $\lambda_k(s, k) := \lim_{\eta \rightarrow 0} \phi_k(v_1, u_{ks}^\eta)$, then by Claim 3 we have $\lambda_k(s, k) := \lim_{\eta \rightarrow 0} \phi_k(u_{sk}^\eta, v_2)$, where $v_2 \in \mathcal{U}$ has $\tau(v_2) = a_k$. By an argument symmetric to the one in the first part of the proof Claim 4, we obtain $\lambda_k(s, k) = \lambda_t(s, t)$. So altogether we have

$$\lambda_j(j, k) = \lambda_s(s, k) = 1 - \lambda_k(s, k) = 1 - \lambda_t(s, t) = \lambda_s(s, t),$$

where the second and last equalities follow from Claim 1.

Case II $s = k$.

Since there are at least three alternatives, we can find a_r distinct from a_j and a_k . Applying Case I repeatedly, we have $\lambda_j(j, k) = \lambda_r(r, j) = \lambda_k(k, j) = \lambda_k(k, t)$.

Cases I and II exhaust all possibilities and establish the claim. \square

We now summarize the implication of Claims 1 through 4. There exists a real number λ lying between 0 and 1 with the following properties. Let a_j and a_k be two arbitrary but distinct alternatives. Consider a utility profile where a_j and a_k are first-ranked for voters 1 and 2, respectively. Now consider a sequence of utility profiles where the utility function of voter 1 is fixed but the utility function of voter 2 is changed in a way such that a_k remains first-ranked and the utility of a_j is increased to 1. Then the sequence of probabilities associated with alternative a_j converges to λ while that of a_k converges to $1 - \lambda$. Similarly, if we fix voter 2's utility function and consider a sequence of utility functions for voter 1 where a_k increases to 1, then the sequence of probabilities associated with a_j and a_k converges once again to λ and $1 - \lambda$, respectively.

Claim 5 For all admissible utility profiles u and all $j \in \{1, \dots, M\}$, if $\phi_j(u) > 0$, then $a_j \in \{\tau(u_1), \tau(u_2)\}$.

Proof Suppose that the Claim is false. Assume w.l.o.g. that there exist distinct alternatives a_j, a_k and a_s and an admissible utility profile u where a_j and a_k are first-ranked by voters 1 and 2 respectively and $\phi_s(u) > 0$. Let η and δ be positive numbers and let $u_{j_s}^\eta$ and $u_{k_s}^\delta$ be admissible utility functions, and denote $\lim_{\eta, \delta \rightarrow 0} \phi_l(u_{j_s}^\eta, u_{k_s}^\delta)$ by λ'_l for all $a_l \in A$. We first prove that $\lambda'_s > 0$.

In order to establish this, we start with a general observation. Let w be a profile and a_t be an alternative which is not first-ranked in w_1 . Let v_1 be an admissible utility function with $v_1(a_t) > w_1(a_t)$ and $v_1(a_l) = w_1(a_l)$ for all $a_l \neq a_t$. Then $\phi_t(v_1, w_2) \geq \phi_t(w)$. In order to see this, observe that since ϕ is strategy-proof, we must have $\sum_r w_1(a_r) \phi_r(w) \geq \sum_r w_1(a_r) \phi_r(v_1, w_2)$ and $\sum_r v_1(a_r) \phi_r(v_1, w_2) \geq \sum_r v_1(a_r) \phi_r(w)$. Combining these two inequalities we have $\sum_r (v_1(a_r) - w_1(a_r)) (\phi_r(v_1, w_2) - \phi_r(w)) \geq 0$, which implies $\phi_t(v_1, w_2) \geq \phi_t(w)$. Thus if we increase the utility of an alternative for a voter in a profile, the probability associated with that alternative cannot decline. Notice that this observation together with our assumption that $\phi_s(u) > 0$ implies that for η, δ small enough, ϕ_s

Combining inequalities (1) and (2) we obtain

$$1 \geq 1 + \sum_{l \neq j, k} \lambda'_l.$$

This implies that $\lambda'_l = 0$ for each $l \neq j, k$. This contradicts $\lambda'_s > 0$, and hence completes the proof of Claim 5. \square

Combining Claims 1–5, we see that for any profile with unequal top alternatives all probability is assigned to the top alternatives (Claim 5), and that agents 1 and 2 can guarantee probabilities as close to λ and $1 - \lambda$ as desired on their respective top alternatives (Claims 1–4). Hence, ϕ is a random dictatorship with weights λ and $1 - \lambda$. This completes the proof of Step 1. \square

Step 2 We now show that a unanimous and strategy-proof CDS is a random dictatorship for arbitrary N

Step 2.1 Suppose $\beta = 0$.

We want to show that for all (u_1, u_2) , $\phi(u_1, u_2) = \phi(u_1, u_1)$.
Suppose not. Then, there are u_1, u_2 and a_k such that

$$\phi_k(u_1, u_2) > \phi_k(u_1, u_1). \quad (3)$$

Now, for $\epsilon > 0$, choose u^ϵ such that

$$\tau(u^\epsilon) = a_k, u^\epsilon(a_j) = \epsilon \text{ for all } a_j \neq a_k.$$

Note that

$$\phi(u^\epsilon, u^\epsilon) = \phi(u_1, u_1) \quad (4)$$

since the coalesced individual has zero weight in the random dictatorship g . From Eqs. (3) and (4) and the specification of u^ϵ , it follows that

$$\lim_{\epsilon \rightarrow 0} \sum_{j=1}^M u^\epsilon(a_j) \phi_j(u_1, u_2) > \lim_{\epsilon \rightarrow 0} \sum_{j=1}^M u^\epsilon(a_j) \phi_j(u^\epsilon, u^\epsilon). \quad (5)$$

In order to prevent individual 1 from manipulating ϕ at $(u^\epsilon, u^\epsilon, u_3, \dots, u_N)$, we need

$$\sum_{j=1}^M u^\epsilon(a_j) \phi_j(u^\epsilon, u^\epsilon) \geq \sum_{j=1}^M u^\epsilon(a_j) \phi_j(u_1, u^\epsilon).$$

In order to prevent individual 2 from manipulating ϕ at $(u_1, u^\epsilon, u_3, \dots, u_N)$, we need

$$\sum_{j=1}^M u^\epsilon(a_j) \phi_j(u_1, u^\epsilon) \geq \sum_{j=1}^M u^\epsilon(a_j) \phi_j(u_1, u_2).$$

Putting these inequalities together, we need

$$\sum_{j=1}^M u^\epsilon(a_j) \phi_j(u^\epsilon, u^\epsilon) \geq \sum_{j=1}^M u^\epsilon(a_j) \phi_j(u_1, u_2). \quad (6)$$

But, Eq. (5) shows that this cannot be satisfied for all values of ϵ , a contradiction. Hence, in this case, ϕ is a random dictatorship with weights $(0, 0, \beta_3, \dots, \beta_N)$.

Step 2.2 Suppose $\beta > 0$.

Let $I' = \{3, \dots, N\}$. Define a function $h: \mathcal{U}^2 \rightarrow \mathcal{L}$ as follows:

$$\text{for all } u_1, u_2, a_j: h_j(u_1, u_2) = \frac{1}{\beta} \left[\phi_j(u_1, u_2) - \sum_{\{i \in I' | \tau(u_i) = a_j\}} \beta_i \right].$$

We want to show that h is a 2-person CDS satisfying strategy-proofness and unanimity.

First, we show that h is a CDS. That is, $h_j(u_1, u_2) \geq 0$ for all $a_j \in A$, and $\sum_j h_j(u_1, u_2) = 1$.

Note that $\sum_j h_j(u_1, u_2) = 1$ follows from the definition of h itself. So, we only need to show that each $h_j(u_1, u_2)$ is non-negative.

Consider u_1, u_2 such that $\tau(u_1) = a_j \neq a_k = \tau(u_2)$.

Claim 6 $\phi_l(u_1, u_2) \geq \phi_l(u_1, u_1)$ for all $a_l \neq a_j$.

Proof Suppose there is $a_l \neq a_j$ such that $\phi_l(u_1, u_2) < \phi_l(u_1, u_1)$. Choose u^ϵ such that $\tau(u^\epsilon) = a_j, u^\epsilon(a_i) \geq 1 - \epsilon$ for all $a_i \neq a_j, a_l$, and $u^\epsilon(a_l) = 0$. Then, since $\phi(u_1, u_1) = \phi(u^\epsilon, u^\epsilon)$,

$$\lim_{\epsilon \rightarrow 0} \sum_{i=1}^M u^\epsilon(a_i) \phi_i(u_1, u_2) > \lim_{\epsilon \rightarrow 0} \sum_{i=1}^M u^\epsilon(a_i) \phi_i(u^\epsilon, u^\epsilon). \quad (7)$$

But, this shows that Eq. (6) is not satisfied for some value of ϵ , and hence contradicts the assumption that ϕ is strategy-proof. \square

Claim 6 establishes that for all $l \neq j, h_l(u_1, u_2) \geq 0$. We still need to show that $h_j(u_1, u_2) \geq 0$. But, notice that we could have “started” from u_2 , and proved that $\phi_l(u_1, u_2) \geq \phi_l(u_2, u_2)$ for all $l \neq k$. This shows that $h_j(u_1, u_2) \geq 0$.

We now want to show that h satisfies unanimity. Choose any u_1, u_2 such that $\tau(u_1) = \tau(u_2) = a_j$ for some $a_j \in A$. Take any $a_k \in A$, and let the *upper contour set* of u_1 for a_k be

$$B(k, u_1) = \{l \in \{1, \dots, M\} \mid u_1(a_l) > u_1(a_k)\}.$$

Claim 7 $\phi_j(u_1, u_2) = \phi_j(u_1, u_1)$.

Proof Suppose there is some a_k such that

$$\sum_{l \in B(k, u_1)} [\phi_l(u_1, u_1) - \phi_l(u_1, u_2)] < 0. \quad (8)$$

For small $\epsilon > 0$ choose u^ϵ such that

- (i) u_1 and u^ϵ are ordinally equivalent.
- (ii) $u^\epsilon(a_l) \geq 1 - \epsilon$ for all $l \in B(k, u_1)$.
- (iii) $u^\epsilon(a_l) \leq \epsilon$ for all $l \notin B(k, u_1)$.

Now, strategy-proofness of ϕ implies that Eq. (6) also holds for the new specification of u^ϵ .

Noting that $\phi(u^\epsilon, u^\epsilon) = \phi(u_1, u_1)$, Eqs. (8) and (6) cannot hold simultaneously as $\epsilon \rightarrow 0$.

Suppose $\phi(u_1, u_1)$ stochastically dominates $\phi(u_1, u_2)$, i.e., all sums in the LHS of (8) are non-negative and at least one sum is positive. But, note that if $\phi(u_1, u_1)$ stochastically dominates $\phi(u_1, u_2)$, then it is well known⁷ that

$$\sum_{a_l \in A} u_1(a_l) \phi_l(u_1, u_1) > \sum_{a_l \in A} u_1(a_l) \phi_l(u_1, u_2). \quad (9)$$

⁷ See, for instance Quirk and Saposnik (1962).

Noting that $\phi(u_1, u_1) = \phi(u_2, u_2)$, Eq. (9) shows that 1 manipulates ϕ at (u_1, u_2) via u_2 . Therefore we must have that for all k

$$\sum_{l \in B(k, u_1)} \phi_l(u_1, u_1) = \sum_{l \in B(k, u_1)} \phi_l(u_1, u_2). \tag{10}$$

Since a_j is the maximal element in the ordering represented by u_1 , the conclusion follows immediately. \square

Claim 7 immediately establishes that h satisfies unanimity.

We now show that h is strategy-proof. Pick any utility functions u_1, u_2, u'_1 . Then

$$\begin{aligned} \sum_{j=1}^M u_1(a_j) h_j(u_1, u_2) &= \sum_{j=1}^M u_1(a_j) \frac{1}{\beta} \left[\phi_j(u_1, u_2, u_3, \dots, u_N) - \sum_{\{i \in I' | \tau(u_i) = a_j\}} \beta_i \right] \\ &\geq \sum_j u_1(a_j) \frac{1}{\beta} \left[\phi_j(u'_1, u_2, u_3, \dots, u_N) - \sum_{\{i \in I' | \tau(u_i) = a_j\}} \beta_i \right] \\ &= \sum_{j=1}^M u_1(a_j) h_j(u'_1, u_2). \end{aligned}$$

Therefore voter 1 cannot manipulate in h . An identical argument establishes that 2 cannot manipulate h either.

Hence, h must be a random dictatorship with weights α_1 and α_2 .

Let $\beta_1 = \alpha_1 \beta$ and $\beta_2 = \alpha_2 \beta$. We want to show that ϕ is a random dictatorship with weights β_1, \dots, β_N . Notice that we would have proved this if we can show that the weights of the 2-agent h constructed earlier do not depend on the choice of (u_3, \dots, u_N) used in the construction of h . In fact, it is sufficient to show that the weights do not change when (say) u_3 changes to u'_3 , because we can change the profile from (u_3, \dots, u_N) to (u'_3, \dots, u'_N) by changing utility functions one at a time.

Suppose that the weights change to α'_1 and α'_2 with $\alpha'_1 > \alpha_1$ when u_3 changes to u'_3 . We show that this violates strategy-proofness of ϕ .

First, suppose $\tau(u_3) = \tau(u'_3) = a_j$. Consider u_1, u_2 such that $\tau(u_1) = a_j$ and $\tau(u_2) = a_l$ where $u_3(a_l) = 0$, that is, a_l is the worst element in terms of u_3 . Then, it is easy to check that 3 manipulates ϕ at $(u_1, u_2, u_3, \dots, u_N)$ since there is a probability transfer of $(\beta \alpha'_1 - \beta \alpha_1)$ from a_l to a_j (with probabilities on all other elements remaining the same) when 3 states u'_3 rather than u_3 . Hence, the weights cannot change if the top elements of u_3 and u'_3 are the same.

Now, suppose $\tau(u_3) = a_j$ and $\tau(u'_3) = a_k \neq a_j$. Using arguments of the previous paragraph, we can assume that $u_3(a_k) = 1 - \epsilon$ and $u_3(a_l) = 0$. Again, assume that $\tau(u_1) = a_j$ and $\tau(u_2) = a_l$. Then,

$$\sum_{i=1}^M u_3(a_i) [\phi_i(u_1, u_2, u_3, \dots, u_N) - \phi_i(u_1, u_2, u'_3, \dots, u_N)] = -\beta (\alpha'_1 - \alpha_1) + \epsilon \beta_3.$$

This difference can be made negative by choosing ϵ small enough. So, ϕ violates strategy-proofness.

This concludes the proof of the induction step, and thus of Theorem 1. \square

Remark 1 Theorem 1 is a slightly weaker version of the result proved in Hylland (1980) and Nandeibam (2004). This is because we assume that the maximal elements of admissible utility functions are unique while Hylland and Nandeibam allow for the possibility of multiple maximal elements. The result in that case is that strategy-proofness and unanimity imply that the CDS is a *weak random dictatorship*. As in a random dictatorship, a weak random dictatorship is characterized by probability weights β_1, \dots, β_N . At any profile u , each voter i can allocate the weight β_i amongst the alternatives which comprises her maximal set $\tau(u_i)$. The probability associated with any alternative is then obtained by aggregating across voters. An important observation is that not all weak random dictatorships are strategy-proof. As we discuss below, both Hylland and Nandeibam first prove their result on the unique maximal element domain and then extend it to the more general domain. Their extension technique can therefore be “added on” to our arguments to prove the more general result.

Relationship to the literature We comment briefly on the related literature. In particular, we point out the differences in our proof and that of Hylland (1980) and Nandeibam (2004). The proofs of Nandeibam and Hylland are actually quite similar. Both papers first show that if the domain of decision schemes is restricted to utility profiles such that each individual has a *unique* maximal element, then strategy-proofness and unanimity⁸ imply the existence of an *additive coalitional power structure*. That is, for each set $I' \subseteq N$, one can define a number $a(I')$ such that if all members of I' unanimously prefer say x to all other elements while all members of the complementary coalition unanimously prefer say y to all other elements and if x, y are the only Pareto-optimal elements, then $a(I')$ is the probability attached to x while $1 - a(I')$ is the probability attached to y . Moreover, $a(\cdot)$ is additive in the sense that $a(I') = a(I_1) + a(I_2)$ if $\{I_1, I_2\}$ is a partition of I' , and $a(\cdot)$ does not depend upon the alternatives x, y . This result is then used to obtain a random dictatorship result on this domain. Use of the Duality Theorem of linear programming (Hylland) or Farkas Lemma (Nandeibam) allows the extension of the result to the domain where the multiple maximal elements are admissible.

We follow a very different approach. We assume that the *minimum* difference in utility levels of alternatives which have different utilities is some positive number $\eta > 0$ where η is the grid size. We prove the theorem in two steps. First, we consider the case of *two* individuals and a domain where each individual has a *unique* top alternative. Using proof techniques developed by Sen (2001) to prove the (deterministic) Gibbard-Satterthwaite result, we now consider the implications of strategy-proofness and unanimity as the grid size η approaches zero. We show that the two-person decision scheme must be a random dictatorship result. The second step uses induction on the number of individuals to establish the random dictatorship result for arbitrary numbers of individuals.

We should also mention related work of Barberà et al. (1998). This paper also uses the same model as we do i.e. they assume that individual utility functions are cardinal valued and that decision schemes are sensitive to cardinalizations. However, instead of assuming unanimity, it explores the consequences of imposing

⁸ Hylland actually does not use unanimity, but a version of Citizen’s Sovereignty.

strategy-proofness and various *smoothness* conditions. Their main result is that if decision schemes are strategy-proof and belong to the class \mathcal{C}^2 , then it must be a unilateral scheme.

4 Strateg -proofness with utilit grids

Our proof technique suggests an interesting extension of the basic framework. In particular, our proof relies heavily on the fact that we can specify utility profiles where the utility of some alternative is arbitrarily close to 1 although it is not maximal. How essential is this in generating the random dictatorship result? In order to answer this question, assume that an admissible utility function has the property that the minimal difference in utility levels of alternatives which have different utilities is at least some $\eta > 0$. The random dictatorship result no longer holds in this framework. The following counter-example demonstrates that non-maximal elements can obtain positive probability for some utility profiles.

Example 2 Let $I = \{1, 2\}$, $|A| = 3$. As before, the best alternative has utility 1, the worst has utility 0, while the maximum utility that the middle alternative can get is $1 - \eta$.

Consider the following rule ϕ^* .

- (i) If $\tau(u_1) = \tau(u_2)$, then ϕ^* assigns probability 1 to the unanimous top alternative.
- (ii) If there are only two Pareto optimal alternatives at a profile, then ϕ^* assigns 0.5 to each of these.
- (iii) If there are three Pareto optimal alternatives at the profile u , but $u_i(a_k) < 0.5$ for some i where a_k is the alternative not ranked first by either voter, then ϕ^* assigns probability 0.5 to each top alternative.
- (iv) Otherwise, ϕ^* assigns $0.5 - d$ to each top alternative and $2d$ to the middle alternative, where d is independent of the profile and $d \leq \eta/2(1 + \eta)$.

Clearly, ϕ^* is unanimous. To see that ϕ^* is strategy-proof, suppose the true profile u is such that either cases (ii) or (iii) apply. Without loss of generality, let $u_1(a_1) > u_1(a_2) \geq u_1(a_3)$. Clearly, 1 cannot increase the probability weight on

In order to prevent this lottery from giving 1 an expected utility greater than 0.5, we need the upper bound on d . \square

The example suggests the following related lines of inquiry. First, notice that there is an upper bound on the probability on the middle alternative. Moreover, this upper bound is an increasing function of the grid size. So, is it generally true that if a CDS is strategy-proof and unanimous, then the maximum probability on non-maximal elements is an increasing function of grid size? The question is interesting because the maximum possible probability on non-maximal alternatives is a crude measure of the distance from some random dictatorship since the latter assigns zero probability to such alternatives.

Second, the CDS constructed in the example approaches a random dictatorship in the limit as the grid size approaches zero. Again, it is of considerable interest to see whether such a ‘limit’ random dictatorship result is true. We turn to these questions in the following subsection.

4.1 A limit result

In this section, we first prove a ‘limit’ random dictatorship result, thus answering the second question at the end of the preceding subsection. We then turn briefly to the first question concerning the maximal probability on non-maximal elements.

For $0 < \eta < 1$ let U^η be the set of utility functions with the property that the minimal difference in utility levels of alternatives which have different utilities is at least η . Consider the following situation. Let $\eta^1, \eta^2, \dots, \eta^k, \dots$ be a decreasing sequence of real numbers in $(0, 1)$ converging to 0. For each k let ϕ^{η^k} be a strategy-proof and unanimous CDS defined on $(U^{\eta^k})^N$. Note that for any $u \in U^N$ there is a minimal number k_u such that $u \in (U^{\eta^k})^N$ for all $k \geq k_u$. With some abuse of notation we can therefore define

$$\lim_{k \rightarrow \infty} \phi^{\eta^k}(u) = \lim_{k \rightarrow \infty, k \geq k_u} \phi^{\eta^k}(u).$$

Obviously, this limit does not have to exist for every u . For instance, take different recurring random dictatorships in the sequence of CDS’s. We will show, however, that *there is a random dictatorship that assigns this limit to any utility profile for which it exists.*

Theorem 2 *There is a random dictatorship $\bar{\phi}$ such that*

$$\bar{\phi}(u) = \lim_{k \rightarrow \infty} \phi^{\eta^k}(u)$$

for all $u \in U^N$ for which the limit exists.

Proof For each $p = 1, 2, 3, \dots$ let $U^p \subset U^N$ be the set of profiles of those utility functions that take values in the set $\left\{0, \frac{1}{p}, \frac{2}{p}, \dots, \frac{p-1}{p}, 1\right\}$.

(i) We first construct a subsequence of $\{\phi^{\eta^k}\}$ which converges on every utility profile in $\bigcup_{p=1}^\infty U^p$, as follows.

Since U^1 is finite, we can construct a subsequence of the given sequence of CDSs which converges on every $u \in U^1$. So we have a subsequence $\phi^{1,k}$ such that $\phi^1(u) \equiv \lim_{k \rightarrow \infty} \phi^{1,k}(u)$ exists for every $u \in U^1$.

We now show that ϕ^1 is a strategy-proof and unanimous CDS on U^1 . To check unanimity, pick any $u \in U^1$ such that for some $a_j \in A$, $\tau(u_i) = a_j$ for all $i \in I$. Then, for all k , $\phi_j^{1,k}(u) = 1$. Hence, $\phi_j^1(u) = \lim_{k \rightarrow \infty} \phi_j^{1,k}(u) = 1$.

We now check that ϕ^1 is strategy-proof. Suppose to the contrary that ϕ^1 is not strategy-proof. Then, there are $u \in U^1$, $i \in I$ and $u'_i \in U^1$ such that

$$\sum_{j=1}^M u_i(a_j) \phi_j^1(u'_i, u_{-i}) > \sum_{j=1}^M u_i(a_j) \phi_j^1(u).$$

But, this contradicts the fact that for each k ,

$$\sum_{j=1}^M u_i(a_j) \phi_j^{1,k}(u'_i, u_{-i}) \leq \sum_{j=1}^M u_i(a_j) \phi_j^{1,k}(u).$$

Next, since U^2 is finite, we may construct a subsequence of the sequence $\phi^{1,k}$ which converges on every $u \in U^2$. So we have a subsequence $\phi^{2,k}$ such that $\phi^2(u) \equiv \lim_k \phi^{2,k}(u)$ exists for every $u \in U^2$. Then, it follows from previous arguments that ϕ^2 is a strategy-proof and unanimous CDS on U^2 . Also, by construction, ϕ^1 and ϕ^2 coincide on $U^1 \subset U^2$.

Continuing in this way, we construct an infinite sequence ϕ^1, ϕ^2, \dots of CDS's such that each ϕ^k is a strategy-proof and unanimous CDS on U^k , and coincides with ϕ^ℓ on U^ℓ for each $\ell < k$.

(ii) Let $u \in \bigcup_k U^k$. Then $u \in \bigcap_{k \geq k_u} U^k$, and therefore $\lim_{k \geq k_u} \phi^k(u)$ exists, and is in fact equal to $\phi^{k_u}(u)$. Denote this limit by $\bar{\phi}(u)$. Then it follows that $\bar{\phi}$ is a strategy-proof and unanimous CDS on $\bigcup_k U^k \subset \mathcal{U}^N$. It is not hard to verify from our proof of Theorem 1 that this theorem still applies (the set $\bigcup_k U^k$ is sufficiently rich), and therefore $\bar{\phi}$ is a random dictatorship. Obviously, $\bar{\phi}$ is defined on all of \mathcal{U}^N .

(iii) Let now $u \in \mathcal{U}^N$ be a utility profile for which $\lim_{k \rightarrow \infty} \phi^{n^k}(u)$ exists. The latter means that this limit is equal to the limit of any subsequence; therefore, we can add the utility profile u to the set $\bigcup_k U^k$ without changing the subsequences $\{\phi^{1,k}\}$, $\{\phi^{2,k}\}$, etc. In particular, it follows that $\lim_{k \rightarrow \infty} \phi^{n^k}(u)$ is equal to $\bar{\phi}(u)$. \square

Theorem 2 implies that, when applied to a fixed utility profile, the probabilities put on non-maximal elements by a converging sequence of unanimous and strategy-proof CDS's must converge to zero as the grid size converges to zero. So this provides a partial answer to the first question raised at the end of the preceding subsection.

5 Conclusion

We have investigated the structure of strategy-proof, cardinal-valued decision schemes satisfying unanimity. One of our contributions is to provide a new and

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.