Journal of Mathematical Economics 69 (2017) 84–90

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

An extreme point characterization of strategy-proof and unanimous probabilistic rules over binary restricted domains ☆

Hans Peters [a,*], Souvik Roy [b,a], Soumyarup Sadhukhan [b], Ton Storcken [a]

[a] Department of Quantitative Economics, Maastricht University, Netherlands
[b] Economic Research Unit, Indian Statistical Institute, Kolkata, India

ARTICLE INFO

Article history:
Received 29 February 2016
Received in revised form
13 January 2017
Accepted 16 January 2017
Available online 24 January 2017

Keywords:
Strategy-proofness
Probabilistic rules
Binary restricted domains
Single-dipped domains

ABSTRACT

We show that every strategy-proof and unanimous probabilistic rule on a binary restricted domain has binary support, and is a probabilistic mixture of strategy-proof and unanimous deterministic rules. Examples of binary restricted domains are single-dipped domains, which are of interest when considering the location of public bads. We also provide an extension to infinitely many alternatives.

1. Introduction

Suppose that in choosing between red and white wine, half of the dinner party is in favor of red wine while the other half prefers white wine. In this situation a deterministic (social choice) rule has to choose one of the two alternatives, where a fifty–fifty lottery seems to be more fair. In general, for every preference profile a probabilistic rule selects a lottery over the set of alternatives. Gibbard (1977) provides a characterization of all strategy-proof probabilistic rules over the complete domain of preferences (see also Sen, 2011). In particular, if in addition a rule is unanimous, then it is a probabilistic mixture of deterministic rules. This result implies that in order to analyze probabilistic rules it is sufficient to study deterministic rules only.

In Peters et al. (2014) it is shown that if preferences are single-peaked over a finite set of alternatives then every strategy-proof and unanimous probabilistic rule is a mixture of strategy-proof and unanimous deterministic rules.[1] The same is true for the multi-dimensional domain with lexicographic preferences (Chatterji et al., 2012). But it is not necessarily true for all dictatorial domains (Chatterji et al., 2014), in particular, there are domains where all strategy-proof and unanimous deterministic rules are dictatorial but not all strategy-proof and unanimous probabilistic rules are random dictatorships.

A binary restricted domain over two alternatives x and y is a domain of preferences where the top alternative(s) of each preference belong(s) to the set $\{x, y\}$ (we allow for indifferences); and moreover, for every preference with top x there is a preference with top y such that the only alternatives weakly preferred to y in the former and x in the latter preference, are x and y.

Outstanding examples of binary restricted domains are domains of single-dipped preferences with respect to a given ordering of the alternatives. Single-dipped preferences are of central interest in situations where the location of an obnoxious facility (public bad) has to be determined by voting: think of deciding on the location of a garbage dump along a road, such that every inhabitant has a single dip (his house, or the school of his children, etc.) and prefers a location for the garbage dump as far away as possible from this dip. Peremans and Storcken (1999) have shown the equivalence between individual and group strategy-proofness in subdomains of single-dipped preferences. They characterize a general class of strategy-proof deterministic rules. In Manjunath (2014) the problem of locating a single public bad along a line segment when agents' preferences are single-dipped, is studied. In particular, all strategy-proof and unanimous deterministic rules are characterized. In Barberà et al. (2012) it is shown that, when all single-dipped preferences are admissible, the range

☆ The authors would like to thank Arunava Sen, Gopakumar Achuthankutty, as well as the editor and two reviewers for their invaluable suggestions. Support from the European Union (EU project 660341) is gratefully acknowledged.

* Corresponding author.
E-mail addresses: h.peters@maastrichtuniversity.nl (H. Peters), souvik2004@gmail.com (S. Roy), t.storcken@maastrichtuniversity.nl (T. Storcken).

[1] Ehlers et al. (2002) characterize such probabilistic rules for single-peaked preferences where the set of alternatives is the real line.

http://dx.doi.org/10.1016/j.jmateco.2017.01.003
0304-4068/

of a strategy-proof and unanimous deterministic rule contains at most two alternatives. The paper also provides examples of subdomains admitting strategy-proof rules with larger ranges. Ayllón and Caramuta (2016) consider group strategy-proofness under single-dipped preferences when agents become satiated: above a certain distance from their dips they become indifferent, and thus they go beyond the binary restricted domain. Further works on strategy-proofness under single-dipped preferences include Öztürk et al. (2013, 2014), Lahiri et al. (forthcoming), and Chatterjee et al. (2016). For strong Nash implementation under single-dipped preferences see Yamamura (2016). There is also a literature on this topic when side payments are allowed, e.g., Lescop (2007) or Sakai (2012).

In the present paper we show that every strategy-proof and unanimous probabilistic rule over a binary restricted domain with top alternatives x and y has binary support, i.e., for every preference profile probability 1 is assigned to $\{x, y\}$. We also show that if a strategy-proof and unanimous probabilistic rule has binary support then it can be written as a convex combination of deterministic rules. Moreover, we present a complete characterization of such rules, by using so-called admissible collections of committees.

Closely related papers are Larsson and Svensson (2006) and Picot and Sen (2012). Larsson and Svensson (2006) include a characterization of all strategy-proof surjective deterministic rules for the case of two alternatives with indifferences allowed. Their Theorem 3 is close to our Theorem 3.9—our theorem is slightly more general since we allow for more than two alternatives. Picot and Sen (2012) show that every probabilistic rule is a convex combination of deterministic rules if there are only two alternatives and no indifferences are allowed.

The paper is organized as follows. Section 2 introduces the model and definitions. Section 3 contains the main results, Section 4 contains an application to single-dipped preference domains, and Section 5 presents an extension to the case where the number of alternatives may be infinite.

2. Preliminaries

Let A be a finite set of at least two alternatives and let $N = \{1, \ldots, n\}$ be a finite set of at least two agents. Subsets of N are called *coalitions*. Let $\mathbb{W}(A)$ be the set of (*weak*) *preferences* over A.[2] By P and I we denote the asymmetric and symmetric parts of $R \in \mathbb{W}(A)$. For $R \in \mathbb{W}(A)$ by $\tau(R)$ we denote set of the first ranked alternative(s) in R, i.e., $\tau(R) = \{x \in A : xRy \text{ for all } y \in A\}$. In general, the notation $\mathcal{D}$ will be used for a set of admissible preferences for an agent $i \in N$. As is clear from the notation, we assume the same set of admissible preferences for every agent. A *preference profile*, denoted by $R_N = (R_1, \ldots, R_n)$, is an element of $\mathcal{D}^n$, the Cartesian product of n copies of $\mathcal{D}$. For a coalition S, R_S denotes the restriction of R_N to S. For notational convenience we often denote a singleton set $\{z\}$ by z.

Definition 2.1. A *deterministic rule* (DR) is a function $f : \mathcal{D}^n \to A$.

Definition 2.2. A DR f is *unanimous* if $f(R_N) \in \cap_{i=1}^n \tau(R_i)$ for all $R_N \in \mathcal{D}^n$ such that $\cap_{i=1}^n \tau(R_i) \neq \emptyset$.

Agent $i \in N$ *manipulates* DR f *at* $R_N \in \mathcal{D}^n$ *via* R_i' if $f(R_i', R_{N\setminus i}) P_i f(R_N)$.

Definition 2.3. A DR f is *strategy-proof* if for all $i \in N$, $R_N \in \mathcal{D}^n$, and $R_i' \in \mathcal{D}$, i does not manipulate f at R_N via R_i'.

[2] I.e., for all $R \in \mathbb{W}(A)$ and $x, y, z \in A$, we have xRy or yRx (completeness), and xRy and yRz imply xRz (transitivity). Note that reflexivity (xRx for all $x \in A$) is implied.

Definition 2.4. A *probabilistic rule* (PR) is a function $\Phi : \mathcal{D}^n \to \triangle A$ where $\triangle A$ is the set of probability distributions over A. A *strict* PR is a PR that is not a DR.

Observe that a deterministic rule can be identified with a probabilistic rule by assigning probability 1 to the chosen alternative.

For $a \in A$ and $R_N \in \mathcal{D}^n$, $\Phi_a(R_N)$ denotes the probability assigned to a by $\Phi(R_N)$. For $B \subseteq A$, we denote $\Phi_B(R_N) = \sum_{a\in B} \Phi_a(R_N)$.

Definition 2.5. A PR Φ is *unanimous* if $\Phi_{\cap_{i=1}^n \tau(R_i)}(R_N) = 1$ for all $R_N \in \mathcal{D}^n$ such that $\cap_{i=1}^n \tau(R_i) \neq \emptyset$.

Definition 2.6. For $R \in \mathcal{D}$ and $x \in A$, the *upper contour set* of x at R is the set $U(x, R) = \{y \in X : yRx\}$. In particular, $x \in U(x, R)$.

Agent $i \in N$ *manipulates* PR Φ *at* $R_N \in \mathcal{D}^n$ *via* R_i' if $\Phi_{U(x,R_i)}(R_i', R_{N\setminus i}) > \Phi_{U(x,R_i)}(R_i, R_{N\setminus i})$ for some $x \in A$.

Definition 2.7. A PR Φ is *strategy-proof* if for all $i \in N$, $R_N \in \mathcal{D}^n$, and $R_i' \in \mathcal{D}$, i does not manipulate Φ at R_N via R_i'.

In other words, strategy-proofness of a PR means that a deviation results in a (first order) stochastically dominated lottery for the deviating agent.

For PRs Φ^j, $j = 1, \ldots, k$ and nonnegative numbers λ^j, $j = 1, \ldots, k$, summing to 1, we define the PR $\Phi = \sum_{j=1}^k \Phi^j$ by $\Phi_x(R_N) = \sum_{j=1}^k \lambda^j \Phi_x^j(R_N)$ for all $R_N \in \mathcal{D}^n$ and $x \in A$. We call Φ a *convex combination* of the PRs Φ^j.

Definition 2.8. A domain $\mathcal{D}$ is said to be a *deterministic extreme point* domain if every strategy-proof and unanimous PR on $\mathcal{D}^n$ can be written as a convex combination of strategy-proof and unanimous DRs on $\mathcal{D}^n$.

For $a \in A$, let $\mathcal{D}_a = \{R \in \mathcal{D} : \tau(R) = a\}$.

Definition 2.9. Let $x, y \in A$, $x \neq y$. A domain $\mathcal{D}$ is a *binary restricted domain* over $\{x, y\}$ if

(i) for all $R \in \mathcal{D}$, $\tau(R) \in \{\{x\}, \{y\}, \{x, y\}\}$,
(ii) for all $a, b \in \{x, y\}$ with $a \neq b$, and for each $R \in \mathcal{D}_a$, there exists $R' \in \mathcal{D}_b$ such that $U(b, R) \cap U(a, R') = \{a, b\}$.

Condition (ii) in the definition of a binary restricted domain is used in the proof of Proposition 3.5. There, we also provide an example (see Remark 3.6) to show that this condition cannot be dispensed with.

We conclude this section with the following definition.

Definition 2.10. Let $x, y \in A$, $x \neq y$. A domain $\mathcal{D}$ is a *binary support domain* over $\{x, y\}$ if $\Phi_{\{x,y\}}(R_N) = 1$ for every $R_N \in \mathcal{D}^n$ and every strategy-proof and unanimous PR Φ on $\mathcal{D}^n$.

3. Results

In this section we present the main results of this paper. First we show that every binary support domain is a deterministic extreme point domain (Corollary 3.3). Next we show that every binary restricted domain is a binary support domain (Theorem 3.4). Consequently, every binary restricted domain is a deterministic extreme point domain (Corollary 3.8). Next, we characterize the set of all strategy-proof and unanimous rules on such binary restricted domains.