Games and Economic Behavior 74 (2012) 651–665

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

The (sub-)optimality of the majority rule

Patrick W. Schmitz [a,b], Thomas Tröger [c,*]

[a] Department of Economics, University of Cologne, Germany
[b] CEPR, United Kingdom
[c] Department of Economics, University of Mannheim, Germany

ARTICLE INFO

Article history:
Received 2 April 2009
Available online 30 August 2011

JEL classification:
D72
D82

Keywords:
Majority rule
Mechanism design
Correlated values
Utilitarianism

ABSTRACT

We consider collective choice from two alternatives. Ex-ante, each agent is uncertain about which alternative she prefers, and may be uncertain about the intensity of her preferences. An environment is given by a probability distribution over utility vectors that is symmetric across agents and neutral across alternatives. In many environments, the majority voting rule maximizes agents' ex-ante expected utilities among all anonymous and dominant-strategy implementable choice rules. But in some environments where the agents' utilities are stochastically correlated, other dominant-strategy choice rules are better for all agents. If utilities are stochastically independent across agents, majority voting is ex-ante optimal among all anonymous and incentive-compatible rules. We also compare rules from an interim-viewpoint.

1. Introduction

A traditional problem in political science is what rule should be used by a society for choosing between two alternatives if each agent is privately informed about her preferences, and if preferences are potentially conflicting. A natural approach to this problem—pioneered by Rae (1969), and subsequently followed by many authors[1]—is to compare rules with respect to the agents' ex-ante expected utilities. We extend existing models in three ways. First, rather than focusing on a pre-fixed set of choice rules, we consider the entire set of dominant-strategy implementable (resp., incentive-compatible) rules, including rules with lottery outcomes. Hence, our contribution may be viewed as an exercise in mechanism design. Second, we allow the agents' preferences to be correlated. This seems appropriate for many applications. Suppose, for example, that there is uncertainty about whether there will be a military threat; if the threat occurs, all agents may be more inclined to prefer an increase of defense expenditure than otherwise. Third, we allow agents to be ex-ante uncertain about the intensities of their future preferences (i.e., preference intensities are state-dependent). Again, this appears natural in many applications. For example, a healthy person may care less about a reduction of the government's health expenditures than a sick person.[2]

Each agent has a privately known von Neumann–Morgenstern utility for each of the two alternatives. Normalizing utilities from one alternative to 0, a *state* is a utility vector that specifies each agent's utility from the other alternative. From an ex-ante viewpoint, the state is distributed according to some commonly known probability distribution that defines the environment. Ex-ante, no agent knows her future utility. At the interim stage, each agent learns her utility and the choice is made. It is important to distinguish an agent's ex-ante preferences (over state-dependent lotteries over alternatives) from

* Corresponding author.
E-mail addresses: patrick.schmitz@uni-koeln.de (P.W. Schmitz), troeger@uni-mannheim.de (T. Tröger).

[1] See the literature review at the end of this section.

[2] The potential importance of preference intensities for institutional design was first emphasized by Buchanan and Tullock (1962). A recent model that allows for uncertainty about preference intensities is Casella (2005).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2011.08.002

her interim preferences (over lotteries over alternatives). At the interim stage, only two preference relations are possible, depending on the *sign* of the agent's utility (we exclude indifference). Ex-ante, however, an agent may also be uncertain about the *absolute value* of her utility, that is, the intensity of her future preferences. In other words, she may face a trade-off between a future self that cares a lot and a different future self that cares little. Formally, this situation fits the state-dependent expected utility model (cf. Fishburn, 1970; Karni, 1985).

We focus on *symmetric* environments, in the sense that the utility probability distribution is symmetric across agents, that is, no agent is special. Given the symmetry across agents, it is natural to focus on *anonymous* choice rules, that is, choice rules that treat all agents in the same manner. Due to symmetry, these choice rules are ex-ante Pareto-ranked, so that no interpersonal utility comparison is needed to find an (ex-ante) optimal rule.[3]

To further simplify matters, we focus on *neutral* environments, in the sense that the utility probability distribution is symmetric around 0. That is, ex-ante both alternatives have the same chance to be preferred with a particular intensity. This implies that there always exists an optimal choice rule that is also *neutral*, in the sense that both alternatives are treated in the same manner.

From the revelation principle, it is without loss of generality to restrict attention to choice rules that are *incentive-compatible* for the agents, that is, no agent must have an incentive to mis-represent her utility if all other agents announce their utilities truthfully. For the most part, we refer to a particularly robust class of incentive-compatible choice rules: *dominant-strategy* choice rules, where announcing her true utility is a dominant strategy for each agent. A dominant-strategy choice rule is robust to the presence of additional private or public signals that may allow agents to revise their beliefs about other agents' utilities.

The class of *weak majority rules* is defined by the property that each alternative is chosen with probability 1/2 unless a sufficiently large fraction of the population prefers one alternative, in which case this alternative is chosen.

Our first result is that in any environment, some weak majority rule is optimal among all dominant-strategy rules; other anonymous and neutral choice rules can only be optimal in non-generic environments. Secondly, we provide a precise characterization of the class of environments where any given weak majority rule is optimal.

An important implication of our characterization is that the (*standard*) *majority rule* is optimal in any environment where the distribution of the utility vector is affiliated. The concept of affiliation (see Milgrom and Weber, 1982, p. 1118) captures a certain type of positive statistical dependence: roughly, the higher an agent's utility, the (weakly) more likely she considers high utilities for the other agents. To paraphrase our earlier example: the stronger an agent's preference for an increase in defense spending, the more likely she considers other agents to have strong preferences in the same direction. A special case of affiliation is stochastic independence. To the extent that affiliation is an appropriate assumption in applications, our result justifies the prevalence of the standard majority rule in real-world institutions.

In some applications, it may be reasonable to assume that members of a minority are somewhat more strongly affected by the collective choice than the members of the majority. In our model, such an assumption corresponds to environments where, conditional on the event that a certain alternative is preferred by few agents, the expected utility of any of these "minority agents" is (in absolute value) larger than the expected utility of the "majority agents". Our characterization result shows that it is in such environments where a weak majority rule can Pareto-dominate the standard majority rule.

The optimality of a weak majority rule in some environments may be a matter of curiosity, but it may also point to a potential improvement of some real-world institutions where the standard majority rule is being used.

Can a higher ex-ante expected utility be achieved if the dominant-strategy requirement is given up, that is, if all incentive-compatible anonymous rules are considered? We show that the answer is "no" if the agents' utilities are stochastically independent—the majority rule is still optimal among all incentive-compatible anonymous rules—, but with stochastically dependent utilities the answer can be "yes". Without independence, we do not have a useful characterization of incentive-compatibility, and a characterization of optimal rules appears difficult. We provide an example which shows that the optimal incentive-compatible rule can lie outside the class of weak majority rules. In the example, there are 3 agents, and both the standard majority rule and the (only other) weak majority rule are outperformed by a rule that chooses each alternative if and only if it is preferred by an *odd* number of agents. Such a rule appears rather fragile, suggesting that a planner's decision for or against the dominant-strategy requirement involves, in general, a trade-off between robustness and expected utility.

We also compare rules from an interim point of view, when each agent knows her utility. An ideal world would be one where, at the interim stage, all agents agree which is the best rule among all anonymous and neutral incentive-compatible rules. We show that the majority rule indeed has this *interim-dominance* property if the agents' utilities are stochastically independent. Finally, dropping the stochastic-independence assumption, we consider a simple class of affiliated environments with three agents and restrict attention to dominant-strategy rules. Here, a rule that is interim-dominant among dominant-strategy rules exists if and only if the positive stochastic dependence between the agents' utilities (as measured by the probability that all agents prefer the same alternative) is not too strong. When no interim-dominant rule exists, many rules are interim-undominated, including, for example, the rule where each agent is a dictator with equal probability.

[3] Our analysis extends to non-anonymous rules if a utilitarian welfare criterion is used, that is, if the sum of the agents' expected utilities is maximized; cf. footnote 13. See Harsanyi (1955) for a discussion of the utilitarian welfare criterion.