Mathematical Social Sciences 73 (2015) 1–5

Contents lists available at ScienceDirect

Mathematical Social Sciences

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

Fair sharing under dichotomous preferences

Conal Duddy

Discipline of Economics, National University of Ireland Galway, Ireland

ARTICLE INFO

Article history:
Received 13 October 2014
Accepted 27 October 2014
Available online 4 November 2014

ABSTRACT

In an important article on collective choice with dichotomous preferences, Bogomolnaia et al. (2005) propose a fairness criterion called *fair welfare share*. We argue that this criterion permits mechanisms that are not fair and we propose an alternative fairness concept called *proportional sharing*. It guarantees to each subgroup, such as a social class, gender or ethnic group, a fair probability that at least one of its members will like the final outcome. It is compatible with *anonymity*, *neutrality*, *ex ante efficiency* and *strategy-proofness* if and only if there are at most four agents or at most three possible outcomes.

1. Introduction

An agent has a dichotomous preference over a set of outcomes if he or she simply partitions the outcomes into "good" and "bad", and makes no further distinction within either category. It may be that the agent has poor information about the outcomes or little time to consider them, or that each outcome impacts on his or her welfare in one of just two possible ways. For example, suppose the agent is considering different possible dates for a meeting. Some dates would entail making a special journey, while the others would not. If there is no other significant distinguishing feature then the agent may have a dichotomous preference over the possible dates.

If a group of agents has dichotomous preferences over a set of possible outcomes, what mechanism of collective choice should be used to choose one of those outcomes?

Approval Voting can be viewed as the utilitarian solution to this problem. Under this mechanism we choose the outcome that is liked (considered "good") by the greatest number of agents. If more than one outcome is maximally popular then we choose randomly from those most popular outcomes. This solution has a number of desirable properties. It is *neutral* and *anonymous*, meaning that neither the names of the outcomes nor of the agents matter to the collective choice. It is *strategy-proof* since no agent ever has any incentive to misrepresent his or her preference. And it is *efficient* in both the *ex post* and *ex ante* sense, which, loosely speaking, means that no welfare is wasted. We will define these properties formally in Section 2.[1]

Yet, while the utilitarian solution has much to recommend it, there may be circumstances in which it is not attractive. Suppose that within the group of agents there is a majority with somewhat homogeneous preferences and a minority with preferences quite different from those of the majority. Under utilitarianism, the minority agents have no chance of obtaining an outcome that they like. The cohesion of the group may be threatened if there is a minority with little or no prospect of happiness. It may be desirable to share welfare, or at least expected welfare, more fairly.

In response to this, Bogomolnaia et al. (2005) propose a criterion called *fair welfare share*. This guarantees to each one of the n agents a probability of at least $1/n$ that the final outcome will be one that he or she likes. They write that "*fair welfare share* uses the random dictator mechanism as the disagreement option that each participant is entitled to enforce. In other words, we give a fair share of control over the final outcome to each participant" (p. 167).

Under the random dictator mechanism an agent is selected at random and given the right to choose the final outcome. The term "disagreement option" comes from bargaining theory. In this context, the disagreement option is the mechanism that the group resorts to if no consensus can be reached on any other mechanism. If the random dictator mechanism is the disagreement option then any other mechanism will be feasible only if it respects the uniform distribution of power across the set of agents under that option. This rules out Approval Voting since that mechanism would be vetoed by minority agents who would fare better under the random dictator mechanism.

Each agent, under the random dictator mechanism, is selected to be dictator with a probability of $1/n$. Therefore, any other proposed mechanism must guarantee that the sum of probabilities over the outcomes that an agent likes is at least $1/n$, which is exactly what is required by the fair welfare share criterion.

E-mail address: conal.duddy@gmail.com.

[1] For extensive discussion of topics relating to Approval Voting see Brams and Fishburn (2007) and Laslier and Sanver (2010).

http://dx.doi.org/10.1016/j.mathsocsci.2014.10.005
0165-4896/

However, let us consider the following example. There are just two possible outcomes a and b, and ninety-nine agents. Fifty agents like a only, while the other forty-nine like b only. Suppose that we assign a probability of just $1/99$ to b and $98/99$ to a. The minority of forty-nine may well feel aggrieved by this lottery. It is only a small step removed from the "tyranny of the majority" they would face under utilitarianism. They would have much better prospects under the random dictator mechanism. Yet, this lottery meets the fair welfare share criterion.

This example reveals that, though *necessary*, it may not be *sufficient* for a mechanism to meet the fair welfare share criterion in order to be feasible when the random dictator mechanism is the disagreement option. This is contrary to the impression given in the earlier quote.

A crucial feature of the random dictator mechanism is that it shares out power fairly not just between individual agents but also between groups of agents. We have noted that an agent's chance of liking the final outcome is at least $1/n$ under the random dictator mechanism. But, critically, it is also true that if we take any two agents then the probability that at least one of them will like the final outcome is at least $2/n$, and is at least $3/n$ for any set of three agents, and so on. By ignoring this, the fair welfare share criterion permits mechanisms that are not feasible when the random dictator mechanism is the disagreement option.

We therefore propose the following fairness criterion, which we call *proportional sharing*. For any subset G of the set of agents, there must be a probability of at least $|G|/n$ that the final outcome will be one that is liked by an agent in G.

This criterion guarantees that every subgroup receives a fair share of welfare, not just each individual. The requirement applies to every subset of the set of agents. This means that no matter how a group is partitioned, whether by social class, gender, ethnicity, age, etc., it can be said that each subgroup has been assigned a fair share of welfare. In the example above, for instance, we are required to assign a probability of $50/99$ to a and $49/99$ to b, the probabilities being in proportion to the sizes of the two groups.

Proportional sharing is an attractive criterion. Is it compatible with other desirable properties of collective choice mechanisms, such as strategy-proofness and efficiency? A simple example of a mechanism satisfying proportional sharing is a variation of random dictatorship whereby we first select an agent at random and then choose randomly from among the outcomes liked by that agent. This mechanism is clearly strategy-proof. However, it is not efficient. For instance, suppose that there are two possible outcomes a and b, and two agents. Both agents like a, but only one of them likes b. Under this simple mechanism, outcome b is chosen with a probability of $1/4$ even though it is Pareto-inferior to outcome a.

Can a mechanism more sophisticated than this satisfy proportional sharing while also being both efficient and strategy-proof?

1.1. Impossibility

Unfortunately, Bogomolnaia et al. (2005) find that when there are at least seventeen outcomes and at least five agents their fair welfare share criterion is not jointly compatible with the standard criteria of social choice theory, (i) ex ante efficiency, (ii) strategy-proofness, (iii) neutrality and (iv) anonymity. Interestingly, it is not known whether those five criteria are compatible when there are, say, sixteen outcomes.

Alas, this result immediately implies that the proportional sharing criterion is also incompatible with criteria (i)–(iv) if there are at least seventeen outcomes and at least five agents, since the proportional sharing criterion is logically stronger than their fair welfare share criterion. However, in many collective choice scenarios there will be fewer than seventeen outcomes to choose between. What upper limit is imposed on the number of possible outcomes by the conjunction of these five properties? Of course, this is likely to depend on the number of agents in the group. If so, then what is the relationship between the number of agents in the group and the upper limit on the number of possible outcomes they may consider? We answer this question by showing that the upper limit is three for any number of agents greater than four, and that there is no upper limit if there are four or fewer agents.

In summary, then, we propose a new fairness concept called proportional sharing and we prove it is consistent with other desirable social choice properties if and only if there are at most four agents or at most three possible outcomes. From a formal point of view, there is a kind of trade-off between our result and that of Bogomolnaia et al. (2005). They use a fairness criterion that is logically weaker than ours, but (unlike theirs) our result is "tight" in regard to the numbers of outcomes and agents. More importantly, we have argued that proportional sharing is the correct fairness condition to require when an appeal is made to the random dictator mechanism.

2. The model

We borrow our notation from Bogomolnaia et al. (2005). Let N be a finite set of n agents and A a finite set of outcomes. A profile U is an $N \times A$ matrix. If agent i likes outcome a then entry u_i^a is one, otherwise u_i^a is zero. The i-row U_i is the dichotomous preferences of agent i, and the a-column U^a indicates the agents who like outcome a. By convention, if agent i is indifferent between all of the outcomes then U_i is a vector of ones, rather than a vector of zeros.

A lottery p is a column vector, indexed by A, of non-negative numbers that sum to one. We write p_a for the probability of outcome a at lottery p. The utility of agent i at lottery p is the probability that an outcome that i likes will be chosen. This is equal to the dot product of U_i and p, which we denote by $U_i \cdot p$.

A mechanism π is a mapping from the set of all profiles to the set of all lotteries. So $\pi(U)$ is a lottery and $\pi(U)_a$ is the probability of outcome a at that lottery.[2]

The following are four properties that a mechanism π may have.

Anonymity. If profile U can be obtained by interchanging rows of profile U' then $\pi(U) = \pi(U')$.

Neutrality. If profile U is obtained by interchanging columns a and b of profile U' then $\pi(U)$ is obtained by interchanging the a and b components of $\pi(U')$.

Strategy-proofness. For any agent i and profiles U and U', if $U_j = U'_j$ for all $j \neq i$ then $U_i \cdot \pi(U) \geq U_i \cdot \pi(U')$.

Proportional sharing. For any profile U and subset G of N, the sum of probabilities $\pi(U)_a$ over all outcomes a liked by at least one member of G is at least $|G|/n$.

An outcome a is Pareto-superior to another outcome b if no agent prefers b to a and at least one agent prefers a to b. We say that an outcome is *efficient* if there is no outcome that is Pareto-superior to it.

In a very similar way, we say that lottery p is Pareto-superior to lottery q if no agent prefers q to p and at least one agent prefers p to q. While there is just one concept of efficiency for outcomes, we distinguish between two concepts of efficiency for lotteries. These are as follows.

2 There is a substantial literature on probabilistic collective choice. Gibbard (1977) is a seminal paper on the aggregation of linear orderings into a lottery. In Gibbard's model each agent has a latent von Neumann–Morgenstern utility scale that induces an ordering over lotteries. The case in which the utility scales are actually elicited from the agents is considered by, for example, Hylland (1980) and Dutta et al. (2007).