Games and Economic Behavior 94 (2015) 15–28

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Symmetry and impartial lotteries

Andrew Mackenzie [1]

Department of Economics, University of Rochester, Rochester, NY, 14627, USA

ARTICLE INFO

Article history:
Received 25 February 2015
Available online 5 September 2015

JEL classification:
D82
D63
D71
D72
M52

Keywords:
Symmetry
Impartiality
Name independence
Voter anonymity
Candidate neutrality

ABSTRACT

A prize is to be awarded, so each candidate designates one of his peers on a ballot. The ballots determine the lottery that selects the winner, and *impartiality* requires that no candidate's choice of designee impacts his own chance of winning, removing incentives for strategic ballot submission. The primary results are (1) a characterization of all impartial rules that treat agents symmetrically as voters, and (2) a characterization of all impartial rules that treat agents symmetrically as candidates. Each rule in either class may be represented as a randomization over a finite set of simple rules. These results have immediate interpretation in a second context: the division of surplus among team members. Corollaries include the constant rule impossibility of Holzman and Moulin (2013), a new dictatorship impossibility, and the first axiomatic characterization of uniform random dictatorship.

1. Introduction

1.1. Overview

A committee wishes to award a prize on the basis of quality to some researcher in a given field, but only the candidates have assessments of one another's quality. How might a system based on peer review treat the researchers symmetrically while remaining immune to manipulation? A team generates a surplus, and each team member has a subjective opinion about the relative contributions. Which symmetric schemes for dividing the surplus using reported evaluations provide no incentives for dishonesty? What does it even mean to treat agents symmetrically in such environments, when each agent is both voter and candidate? For ease of exposition we temporarily focus on the first example, though we return to the second in the discussion of Section 3.

To pursue these questions, we consider the stochastic version of the setting of Holzman and Moulin (2013): we are tasked with selecting an agent to receive a prize, and (as our interest is symmetry) our selection may involve randomization. Unfortunately, we lack the information to reach a decision! So we turn to the agents themselves, each of whom has (1) a "selfish" preference, caring only whether or not he wins, and (2) an informed opinion about his peers. To simplify without losing generality, assume we know the prize is coveted by all.[2] In order to gather the opinions, we ask each agent

E-mail address: andrew.k.mackenzie@gmail.com.

[1] I thank Salvador Barberà, Paulo Barelli, Agustín Bonifacio, Srihari Govindan, Mallesh Pai, Romans Pancs, and William Zwicker. Hervé Moulin and my anonymous reviewers provided especially thorough remarks. Most of all, I thank William Thomson for his colossal feedback and support.

[2] All of our results apply more generally, when any agent might prefer to receive the prize, prefer not to, or be indifferent. For example, if the prize is department chairmanship, some may feel the prestige outweighs the burden while others may feel the reverse. In this case, *efficiency* would require assigning

http://dx.doi.org/10.1016/j.geb.2015.08.007

to submit a ballot designating one of his peers. Our problem is to select a *rule* for making a decision once the ballots are received; formally, a rule is a function which associates each ballot profile with a lottery over agents, where we interpret the lottery's outcome as the prize-winner.

Our concern is that an agent might strategically corrupt the valuable information on his ballot—a scenario that has been portrayed in popular culture[3]—so we focus on rules that satisfy *impartiality*: each agent's chance of winning is determined solely by the reports of his peers, removing incentives for a selfish agent to submit a corrupted ballot. This axiom is in the same spirit as the usual *strategy-proofness*; the difference is that we are not asking for preferences, but rather for opinions that are unrelated to preferences.

We emphasize here that in contrast to Holzman and Moulin (2013), we do not insist that ballots be interpreted as nominations. There may be circumstances where we prefer to ask each agent to designate the *worst* choice among his colleagues, in which case the probability an agent wins should *decrease* when he is designated. We might even want to ask different questions to different agents! We therefore take a general mechanism design approach, investigating all that our axioms allow given each agent's specified message space. Though not as general as the consideration of all *game forms* (Gibbard, 1973), focusing on specified message spaces has been a familiar tactic for mechanism design since the acclaimed *revelation principle* of auction theory (Myerson, 1981). Other message spaces studied for problems similar to ours include

- approval ballots (Alon et al., 2011; Kasher and Rubinstein, 1997; Samet and Schmeidler, 2003),
- rankings of peers (Berga and Gjorgjiev, 2015; Amorós et al., 2002; Ando et al., 2003; Ohseto, 2007),
- numerical ratings of peers (Kurokawa et al., 2015; Ng and Sun, 2003; Ohseto, 2012), and
- relative shares for pairs of peers (de Clippel et al., 2008).

The literature's usual symmetry axiom in settings like ours, proposed by Ng and Sun (2003), is the requirement that when "the names of the agents both as candidates and voters are permuted," the outcomes are permuted accordingly. Interpreting this axiom can be tricky, however, as it somewhat blurs the familiar notions of anonymity and neutrality used in settings where the voters and candidates are disjoint.[4] To avoid any confusion, we propose to call this axiom *name independence*.

But are there other symmetry axioms we might want to consider? Holzman and Moulin (2013) suggest there are by means of an example: an *impartial* rule satisfying *name independence* that sometimes makes distinctions when each agent is designated by one other, simply by favoring a pair of agents who designate each other when there is no other such pair. While we can imagine contexts where such differentiation may be desired—say when the prize is an academic award and the favored pair represents a small subfield—nevertheless we should have at least one symmetry axiom in our arsenal that rules it out.

We propose two such axioms, one in the spirit of anonymity and the other in the spirit of neutrality. To understand these axioms, imagine that each ballot includes both (1) a *signature* recording the name of the designator, and (2) an *entry* recording the name of the designee; when i designates j, the ballot includes i as signature and j as entry. *Voter anonymity* requires that when the signatures are permuted, yielding no self-designations, the outcome be unchanged. *Candidate neutrality* requires that when the entries are permuted, yielding no self-designations, the outcome be permuted accordingly.

Though we feel both proposed axioms are intuitive, they are also both strong. Either rules out the *permutation mechanism* (Fischer and Klimm, 2014), an algorithmic *impartial* rule that processes signed ballots in a random order, at each stage counting only the ballots of agents who have already been eliminated. This rule was designed to ensure the expected number of designations received by the prize-winner is high relative to the maximum number of designations received by an agent, and it performs quite well.

The three symmetry axioms are logically independent. Both the permutation mechanism and the earlier example from Holzman and Moulin satisfy only *name independence*. A *fixed-winner* rule, which always assigns the prize to an exogenously chosen agent (the "fixed winner"), satisfies only *voter anonymity*. A *dictator* rule, which always assigns the prize to the designee of an exogenously chosen agent (the "dictator"), satisfies only *candidate neutrality*.

To state our main findings, we need only introduce a few more rules. The rule that uniformly draws and uses one of the dictator rules is *uniform random dictatorship*. At the opposite extreme of ballot interpretation is the *card-ripping rule*: when there are n agents, (1) each agent's name appears on $n-1$ identical cards, (2) designating an agent means destroying one of his cards, and (3) the winner is drawn uniformly from the remaining cards. Finally, an *eliminator* rule eliminates the designee of an exogenously chosen agent (the "eliminator"), then draws the winner uniformly from the remaining agents.

---

the chairmanship to a faculty member who desires it if there is one, but our objective is to assign the chairmanship to someone who is best-suited. Our approach—the preferences are not used or even asked for—is appropriate if an agent's preference for the prize is private information that reveals nothing about how well-suited he is, but might yield incentives to lie.

[3] In the 1969 comedy film Putney Swope, each board member is to nominate another board member to become chairman, and provided there is no tie the member with the most votes is to win. The result? The underdog wins almost unanimously: "We all voted for him because we thought no one else would vote for him [...] and I will defend that mistake with my life."

[4] This blurring is evident in the literature's various names for the axiom—in their respective settings, Ng and Sun (2003) call it "neutrality," de Clippel et al. (2008) call it "anonymity," and Holzman and Moulin (2013) call it "symmetry."