Games and Economic Behavior 90 (2015) 119–127

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Size versus fairness in the assignment problem ☆

Anna Bogomolnaia, Herve Moulin *

University of Glasgow, United Kingdom

ARTICLE INFO

Article history:
Received 17 June 2014
Available online 11 February 2015

JEL classification:
D63
D70

Keywords:
Envy-freeness
Random assignment
Probabilistic serial

ABSTRACT

When not all objects are acceptable to all agents, maximizing the number of objects actually assigned is an important design concern. We compute the guaranteed size ratio of the Probabilistic Serial mechanism, i.e., the worst ratio of the actual expected size to the maximal feasible size. It converges decreasingly to $1 - \frac{1}{e} \simeq 63.2\%$ as the maximal size increases. It is the best ratio of any Envy-Free assignment mechanism.

© 2015 Elsevier Inc. All rights reserved.

1. The problem and the punchline

Lotteries are commonly used to allocate indivisible resources (*objects*), especially so when monetary transfers are ruled out. Examples include the assignment of jobs to time-slots, of workers to tasks or offices, the allocation of seats in overdemanded public schools (Abdulkadiroğlu and Sonmez, 2003; Kojima and Unver, 2014), of students to dormitory rooms or courses, etc. An excellent survey is the work of Sonmez and Unver (2011). Using cash transfers and prices in such problems skews the distribution toward the wealthier agents, which is arguably inefficient (Che et al., 2013); they are also ruled out by moral objections to commoditizing certain objects like human organs (Roth, 2007). Randomization is then a practical way to restore fairness at the ex ante stage.

A great deal of recent economic research applies the methodology of mechanism design to the random allocation of objects. The earliest results (briefly reviewed below) bear on the benchmark *random assignment* problem where each agent wants at most one object, reports an *ordinal* preference ranking of those objects, and receives a random object, or no object. The compelling test of ex ante fairness is the *No Envy* property: when Ann compares the probability distribution of the object she will receive to the distribution of the object Bob will receive, she finds that her distribution stochastically dominates Bob's. We focus here on the tension between No Envy and the potential wastefulness of the mechanism when agents have outside options.

Outside options are pervasive in many practical instances of assignment: in the school choice problem they are offered by private schools; college students can live off campus; jobs have deadlines so a time slot beyond that date is worse than

☆ We are especially grateful to our colleagues David Manlove and Baharak Rastegari for introducing us to their research project EPSRC Grant # EP/K010042/1, and sharing the results already obtained with their colleagues at the University of Liverpool. Special thanks also to Jay Sethuraman for guiding us through the literature on online matching, and to Bettina Klaus, Tadashi Hashimoto, and seminar participants in Warwick, Manchester and Toulouse, for stimulating conversations. The comments of three anonymous referees have been very helpful.

* Corresponding author.
E-mail address: herve.moulin@glasgow.ac.uk (H. Moulin).

http://dx.doi.org/10.1016/j.geb.2014.11.006
0899-8256/© 2015 Elsevier Inc. All rights reserved.

dropping the job at the outset, and so on. An agent will not accept an object worse than his/her outside option, and this affects the size of the realized assignment (number of agents who receive an object). This size is a measure of utilization of the resources, therefore maximizing it is important in its own right: filling the largest possible number of seats/rooms/jobs, is a component of the system performance, to which public school administrators, the housing office on campus, the job manager, etc., are paying attention.

Note that the largest feasible size of an assignment only depends upon the bipartite graph of acceptability, and ignores the finer information in the profile of individual preferences. So it is not surprising that size maximization often conflicts with fairness and incentive compatibility. This is obvious in the following elementary example with two objects a, b and two agents Ann, Bob, who both prefer a to b. If both objects are acceptable (better than his outside option) to Bob but Ann only accepts a, then assigning a to Ann and b to Bob is the only assignment of maximal size. It is obviously unfair to Bob who *envies* Ann's allocation. Moreover selecting this assignment also gives Bob the incentive to report that only a is acceptable, if he prefers a 50% chance of getting a to a 100% chance of b.

We give a precise lower bound on the trade-off size versus fairness in the random assignment problem with outside options. We define the *size ratio* of an assignment at a given profile of preferences as the ratio of its size to the maximal feasible size when we must only ensure that everyone gets an acceptable object. The *guaranteed m-size ratio* of a random assignment mechanism is its worst size ratio over all assignment problems such that the maximal size of a feasible assignment is m.

We discuss first the *Probabilistic Serial* mechanism (hereafter PS; see next section and Section 5), the only known mechanism to date combining Envy-Freeness with Efficiency (Pareto optimality). We compute exactly the guaranteed m-size ratio of PS: it decreases with m from $\frac{3}{4}$ for $m = 2$ (achieved in the example above) and converges to $1 - \frac{1}{e} \simeq 0.632$ as m grows arbitrarily large. Then we show that this is the greatest guaranteed m-size ratio among *all* envy-free assignment mechanisms.

2. Related literature

1) The first random assignment mechanism in Hylland and Zeckhauser (1979) is a competitive equilibrium with fiat money to buy lotteries, and relies on cardinal (von Neuman Morgenstern) individual utilities over objects. Such individual reports are too complex in practice, so attention turned to the more realistic *ordinal* mechanisms where a report is simply a ranking of the acceptable objects. The most natural ordinal mechanism is the time honored *Random Priority* (RP), a.k.a. serial dictatorship, discussed first in Abdulkadiroğlu and Sonmez (1998) who offer a market-like interpretation of RP. Next Bogomolnaia and Moulin (2001) proposed the alternative *Probabilistic Serial* mechanism that fares better than RP in terms of efficiency and fairness, but has worse incentive properties: PS is Envy-Free but RP is not, while RP is strategyproof but PS is not. Subsequent work considerably refined the comparison of RP and PS; for instance Erdil (2014) discusses a different wasteful aspect of RP that PS does not share, while Bogomolnaia and Heo (2012), and Hashimoto et al. (2014) characterize PS axiomatically. Particularly relevant here is the asymptotic equivalence of PS and RP along certain expansion paths of the economy with a fixed, finite number of types of objects, while the number of copies of each object grows at roughly the same rate as the number of agents. First established in Che and Kojima (2010), this result was recently generalized in Liu and Pycia (2013) to a broad class of random assignment mechanisms. However for any fixed finite number of agents, the expected sizes achieved by PS and RP are not comparable at all preference profiles.

2) The goal of maximizing the assignment size appears first in the algorithmic mechanism design literature. An early instance is the work of Procaccia and Tennenholtz (2009), discussing the tradeoff between Strategy-Proofness and the utilitarian minimization of aggregate cost. Another seminal example, closer to home, is in the bilateral matching problem. When preferences have ties and remaining single is preferred to some potential partners, not all stable matchings are of the same size (the "rural hospital theorem" does not apply), so it is natural to look for a stable matching of maximal size (Irving and Manlove, 2009), or for a maximal cardinality matching with the smallest number of blocking pairs (Biró et al., 2010): both questions turn out to be NP-hard.

3) It results from our Theorem and earlier results in Cres and Moulin (2001) that the guaranteed m-size ratio of RP is always bounded above by r_m, the guaranteed ratio of PS. On the other hand, the results in Bhalgat et al. (2011) and Krysta et al. (2014) provide the lower bound $1 - (1 - \frac{1}{m+1})^m - \frac{1}{m}$, a sequence converging **in**creasingly to $1 - \frac{1}{e}$, for the guaranteed m-size ratio of RP. Thus, for problems with a large feasible assignment, RP and PS have approximately the same guaranteed ratio $1 - \frac{1}{e}$. The interesting fact is that the proof techniques in Bhalgat et al. (2011), Krysta et al. (2014) are radically different than ours. They are closely related to the problem of designing an online bilateral matching algorithm maximizing the match size relative to the maximal size feasible offline. The Ranking algorithm of Karp et al. (1990) selects randomly and uniformly an ordering of the objects, then assigns to the incoming agent the highest acceptable object in that ordering; its m-guaranteed size is no less than $1 - (1 - \frac{1}{m+1})^m$ (see also Birnbaum and Mathieu, 2008 for a simpler proof and Kalyanasundaram and Pruhs (2000) for a generalization to multiple objects).

These results suggest that, for any m, RP *may have* the best guaranteed m-size ratio among all strategyproof mechanisms, despite the fact that it is dominated by some less wasteful strategy-proof mechanisms (Erdil, 2014). In fact Theorem 6.2 in Krysta et al. (2014) establishes this conjecture for the case $m \leq 3$. Thus our results confirm the intuition that PS and RP are similar for large problems, but no part of our Theorem can be deduced from existing results, even in an asymptotic sense.

Download English Version:

https://daneshyari.com/en/article/5071635

Download Persian Version:

https://daneshyari.com/article/5071635

Daneshyari.com