

A. o. on . n . A

David Dillenberger[†] Uzi Segal[‡]

March 6, 2024

Abstract

- n o s o n

Ten thousand children need to be allocated into ten schools, each accommodating one thousand of them. The schools are not the same, and parents may rank them in different ways. However, if all children are considered equal, then a social lottery seems to be the best solution, where each student has an equal chance to attend each of the ten schools.¹ This procedure is egalitarian — everyone gets the same lottery — and feasible. But is it efficient? Specifically, is there no other procedure such that ex-ante, before people know their allocated school, they will get a better lottery?

If individual preferences over the schools are not the same, then this procedure may be inefficient — for example, if each school is ranked best by exactly 1000 parents. It is true that if all individuals are expected utility maximizers and have the same preferences over lotteries (and in particular, over the schools), then this procedure leads to an efficient allocation. This is also the case if all have the same quasi-concave preferences over lotteries. But if preferences are quasi-convex, and a mixture of two indifferent lotteries is inferior to the mixed lotteries, then we show that this procedure is never efficient, regardless of whether individual preferences are the same or not. Such preferences are implied by some well known alternatives to expected utility theory (for example, Tversky and Kahneman's [35] Cumulative Prospect Theory, where risk aversion implies quasi-convexity. See discussion below).

mixed. Most of the experimental literature that documents violations of expected utility (e.g., Coombs and Huang [8]) found either preference for randomization or aversion to it. Camerer and Ho [6] find support for quasi-convexity over gains and quasi-concavity over losses. An example of behavior that distinguishes between the two attitudes to mixture is the probabilistic insurance problem of Kahneman and Tversky [17]. They showed that in contrast with experimental evidence, any risk averse expected utility maximizer must prefer probabilistic insurance to regular insurance. Sarver [30]

like in standard expected utility, this inefficiency does not rely on cardinal information which can be used to assess the intensity of preferences over the basic goods, but on the ordinal property of the preferences over lotteries

That is, the average lottery faced by the participants is a uniform distribution over the N goods. Obviously, this distribution is feasible. The sum of its components must be 1, as the original lottery satisfies eq. (1). And if the average lottery is not uniform, then the original allocation is not feasible as it must violate eq. (2).

Any solution q specifies the probability distribution over final outcomes for each individual. The Birkhoff-von Neumann Theorem ([4],[39]) guarantees that for any q there is always a (social) lottery over all possible permutations of the allocations of the final outcomes that induces the marginal probabilities of q .³

3. An efficient allocation

We first characterize solutions that are feasible, that is, satisfy equations (1) and (2), and are ex-ante Pareto efficient, in the sense that there is no other solution in which some individuals are strictly better off and no one is worse off.⁴ As preferences are continuous over a compact domain, feasible efficient allocations exist. We show that in such allocations, and without any further assumptions on individuals' preferences, all but 'not too many' individuals obtain either a degenerate lottery or a lottery with positive probabilities on two goods only.

³ In the case of a uniform distribution, this is simply the original allocation. ⁴ This definition is used in the literature to refer to allocations that are efficient in the sense of Pareto.

x_r x_s x_t and that they both receive lotteries with positive probabilit-

This suggests a broader point. There are known results that imply the equivalence of different randomized mechanisms and random serial dictatorship (Abdulkadiroğlu and Sönmez [3]; see also Pathak and Sethuraman [25]), in the sense that they induce the same ex-ante probability distribution over the final goods. But then those seemingly identical mechanisms are also typically ex-ante inefficient. If social planners know the individuals' preferences over lotteries, and in particular that they are strictly quasi-convex, they can improve the agents' welfare ex-ante.⁷ Importantly, this argument only relies on simple, observable information: strict quasi-convexity of preferences and the size of the supports of the lotteries that are used, rather than on the intensity of preferences over the goods or the weights given to each of them in the corresponding lotteries.⁸

While for exposition purposes we confine our attention to the case of strict quasi-convex preferences, Theorem 1 generically also holds under expected utility, which is linear (and hence also weakly quasi-convex) in probabilities. Suppose all individuals are expected utility maximizers. Hylland and Zeckhauser [16] use competitive equilibrium with equal incomes to show the existence of a solution in which almost all agents receive a binary lottery.⁹ Our result holds without relying on any market mechanism.

Also assuming expected utility, Bogomolnaia and Moulin [5] show how random serial dictatorship, which uses uniform distribution to rank agents, is not necessarily even ordinally efficient; it may induce for each agent a distribution over the goods that is stochastically dominated, with respect to that agent's ordinal preferences, by another feasible distribution. Their suggested

⁷Note that we ignore here the question of strategy-proofness, that is, how to guarantee that agents truly reveal their preferences. We are instead focusing only on the properties of the induced allocation (of lotteries) for any given set of preferences.

⁸Abdulkadiroğlu, Che, and Yasuda [2] point out that cardinal information allows the

probabilistic serial mechanism (which is ordinally efficient) is typically not ex-ante efficient. It is also worth noting that their solution implies that agents with the same ordinal preferences must receive the same lottery over goods. In our case, even if all agents have the same cardinal preferences (and are strictly quasi-convex), *necessarily* not all of them receive the same lottery, as otherwise, the same binary lottery to all will not allocate all available goods.

$$x_i = \frac{1}{n} \sum_{j=1}^n x_j$$

When all individuals have the same preferences, it is natural to require that

outcomes that can simultaneously be used. Note that many individuals may hold the same binary lottery, but only one individual can hold any non-binary

characterizing allocations that are efficient and satisfy the following No-Envy criterion.

Definition For all a and b , $q^a \succ_a q^b$.

No-Envy postulates that in the allocation of lotteries, no individual would prefer to replace their lottery with that of any other agent.¹¹ Clearly, if $q_1 = \dots = q_N = q$, then No-Envy implies equality, in the sense that for all $a, b \in \mathcal{A}$, $f(a) \sim f(b)$.

No-Envy is appealing on a normative ground.

JR-457.076(t)-0.925088(h)-1.28246(e)-0.126613(i)-0.7



they do not satisfy monotonicity with respect to first order stochastic dominance, and equilibrium does not exist. We show in the proof of Theorem 3 that this stronger version of monotonicity eliminates the existence problem.

Let T be the number of lotteries used in the proposed solution. Then for $h = 1, 2, \dots, T$ there is a continuum of agents who receive the same binary lottery, say $(x^h, \rho^h; y^h, 1 - \rho^h)$ for some outcomes x^h, y^h and $\rho^h \in [0, 1]$. The implementation of this, so that the fraction of the people in this group that receives x^h is ρ^h , can be guaranteed by using the appropriate law of large numbers for a continuum of independent random variables. Such approach appears, for example, in Sun [34], and we adopt here his measure theoretic framework.¹²

may know which of the two vases is Ming and which is a modern counterfeit,

Conclusion

The use of binary lotteries is pervasive in economics. Many experimental works are conducted with choices among such lotteries (or between them



Since all probabilities are between 0 and 1, it follows by standard arguments that there is a subsequence of q^h , without loss of generality the sequence itself, such that for all $n = 1, \dots, Nk$, $q^{n,h} \rightarrow q^{n,*}$. The vector $q^* = (q^{1,*}, \dots, q^{Nk,*})$ satisfies eqs. (1) and (2), hence it is a solution. Since V is continuous it satisfies equality, and as by the continuity of V , $V(q^{n,*}) = v^n$, it follows by the definition of v that q^* is optimal.

We next establish the bound M on the number of possible binary lotteries. Relabel the basic goods so that all agents agree that $x_1 \quad x_2 \quad \dots \quad x_N$. Consider the set $B := \{(x_i, x_j) : i \leq \frac{N}{2}, j > \frac{N}{2}\}$. Note that for any pair $(x_k, x_l) \notin B$, there is $(x_i, x_j) \in B$ with either (i) $j \leq l, k$ with at least one

