Journal of Mathematical Economics 60 (2015) 123–133

Contents lists available at ScienceDirect

Journal of Mathematical Economics

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

Universal Pareto dominance and welfare for plausible utility functions

Haris Aziz[a,b], Florian Brandl[c,*], Felix Brandt[c]

[a] NICTA, Sydney, Australia
[b] UNSW, Sydney, Australia
[c] TU München, Munich, Germany

ARTICLE INFO

Article history:
Received 15 September 2014
Received in revised form
5 June 2015
Accepted 24 June 2015
Available online 6 July 2015

Keywords:
SSB utility
Pareto optimality
Social welfare
Ordinal efficiency
Stochastic dominance

ABSTRACT

We study Pareto efficiency in a setting that involves two kinds of uncertainty: Uncertainty over the possible outcomes is modeled using lotteries whereas uncertainty over the agents' preferences over lotteries is modeled using sets of plausible utility functions. A lottery is *universally Pareto undominated* if there is no other lottery that Pareto dominates it for *all* plausible utility functions. We show that, under fairly general conditions, a lottery is universally Pareto undominated iff it is Pareto efficient for *some* vector of plausible utility functions, which in turn is equivalent to affine welfare maximization for this vector. In contrast to previous work on linear utility functions, we use the significantly more general framework of skew-symmetric bilinear (SSB) utility functions as introduced by Fishburn (1982). Our main theorem generalizes a theorem by Carroll (2010) and implies the *ordinal efficiency welfare theorem*. We discuss three natural classes of plausible utility functions, which lead to three notions of ordinal efficiency, including stochastic dominance efficiency, and conclude with a detailed investigation of the geometric and computational properties of these notions.

1. Introduction

Consider two agents, Alice and Bob, and an unpleasant task that may be assigned to Alice (a), to Bob (b), or to neither of them (c). All we know about their pairwise preferences over the possible assignments is that they both strongly prefer not being assigned the task and that their preference between letting the other agent perform the task or not having the task assigned at all is less intense. In other words, Alice prefers b and c with equal intensity to a. Her preference between b and c is unknown, but known to be less intense than her preference between b and a (or, equivalently, c and a). Bob's preferences are defined analogously. All preferences that match the above description will be called *plausible*. Clearly, outcome c, in which the task is not assigned, is Pareto efficient for every plausible preference configuration. In general, however, such outcomes need not exist and a reasonable extension of the notion of Pareto efficiency in the face of uncertainty is to consider an outcome efficient if there is no other outcome that is preferred by all agents for all plausible preferences. In the example, all three outcomes are efficient according to this definition. However, not every *lottery* over these outcomes is efficient. In fact, it turns out that the only efficient lotteries are those that do not put positive probability on both a and b. The set of efficient lotteries thus exhibits two phenomena that we will observe frequently in this paper: It fails to be convex and whether a lottery is efficient only depends on its support.

More generally, following McLennan (2002), Manea (2008), Carroll (2010), and others, this paper investigates Pareto efficiency in a setting that involves two kinds of uncertainty: Uncertainty over the possible outcomes is modeled using probability distributions (lotteries) whereas uncertainty over the agents' preferences over lotteries is modeled using sets of plausible preferences relations over lotteries. A lottery is *potentially efficient* if it is Pareto efficient for *some* vector of plausible preference relations while it is *universally undominated* if there is no other lottery that Pareto dominates it for *all* plausible preference relations. It is easily seen that every potentially efficient lottery is universally undominated. Our main theorem shows that, under fairly general conditions, the converse holds as well, i.e., the set of universally undominated and the set of potentially efficient lotteries coincide. We prove the statement for the unrestricted social choice domain, which implies the same statement for many subdomains of interest such as roommate markets, marriage markets, or house allocation. As we will see, the set of universally undominated lotteries may not even be a geometric object with flat sides, i.e., it may fail to be the union of

* Corresponding author.
E-mail addresses: haris.aziz@nicta.com.au (H. Aziz), brandlfl@in.tum.de (F. Brandl), brandtf@in.tum.de (F. Brandt).

http://dx.doi.org/10.1016/j.jmateco.2015.06.014
0304-4068/

Table 1
Properties of varying notions of ordinal efficiency. An efficiency notion satisfies *existence* if every preference profile admits an efficient lottery. An efficiency notion satisfies *convexity* if the convex combination of two efficient lotteries is efficient. An efficiency notion is *support-dependent* if a lottery is efficient iff every lottery with the same support is efficient. *Efficient improvements exist* for an efficiency notion if, for any given lottery, there is an efficient lottery that dominates the original lottery.

	Existence	Convexity	Support-dependence	Existence of efficient improvements
BD-efficiency	+	+	+	+
Ex post efficiency	+	+	+	+
SD-efficiency	+	−	+	+
PC-efficiency	+	−	−	−

finitely many polytopes. One corollary of our main theorem is that the set of universally undominated lotteries is always connected.

In contrast to previous work, which is based on von Neumann–Morgenstern (vNM) utility functions, we assume that preferences over lotteries are given by sets of *skew-symmetric bilinear (SSB) utility functions*. Classic vNM utility theory postulates the independence axiom[1] and the transitivity axiom. However, there is experimental evidence that both of these axioms are violated systematically in real-world decisions. The Allais Paradox (Allais, 1953) is perhaps the most famous example pointing out violations of independence. A detailed review of such violations is provided by Machina (1983). Mas-Colell et al. (1995, p. 181) conclude that "because of the phenomena illustrated [...] the search for a useful theory of choice under uncertainty that does not rely on the independence axiom has been an active area of research".

Even the widely accepted transitivity axiom seems too demanding in some situations. For example, the preference reversal phenomenon[2] (see, e.g., Grether and Plott, 1979) shows failures of transitivity. SSB utility theory assumes neither independence nor transitivity and can accommodate both effects, the Allais Paradox and the preference reversal phenomenon. Still, the existence of maximal elements, arguably the main appeal of transitivity, is guaranteed for SSB utility functions by the minimax theorem (von Neumann, 1928). For a more thorough discussion of SSB utility theory we refer the reader to Fishburn (1988).

Sets of plausible utility functions are typically interpreted as incomplete information on behalf of the social planner. Indeed, it seems quite natural to assume that the social planner's information about the agents' utility functions is restricted to ordinal preferences, top choices, or subsets of pairwise comparisons with further conditions implied by domain restrictions. Three particularly interesting classes of plausible utility functions arise when contemplating that only ordinal preferences over pure outcomes are known. For a given binary preference relation R_i, we consider

- the set of all SSB functions consistent with R_i,
- the set of all vNM functions consistent with R_i, and
- the unique canonical SSB function consistent with R_i (where canonical means that all pairwise comparisons have the same intensity).

These sets give rise to three natural extensions of preferences over alternatives to preferences over lotteries and thereby to three notions of ordinal efficiency. While the second notion is equivalent to the well-studied notion of *stochastic dominance (SD) efficiency*, the other two notions, one weaker and one stronger than *SD*-efficiency, have not been studied before. We call the weaker notion *bilinear dominance (BD) efficiency* and the stronger one *pairwise comparison (PC) efficiency*. The preference extension underlying *PC*-efficiency seems particularly natural because it prescribes that an agent prefers lottery p to lottery q iff it is more likely that p yields a better alternative than q. In contrast to the other preference extensions, the *PC* extension always yields a complete preference relation. Yet, *PC* preferences cannot be modeled using vNM utility functions.

In the second part of the paper, we investigate geometric as well as computational properties of efficiency notions obtained via universal undominatedness. Our findings include the following observations (see also Table 1).

- Whether a lottery is *BD*-efficient or whether it is *SD*-efficient only depends on its support.
- The set of *SD*-efficient lotteries and the set of *PC*-efficient lotteries may fail to be convex. As a consequence, the convex combination of two mechanisms that return *SD*-efficient lotteries may violate *SD*-efficiency.
- Universally undominated lotteries can generally be found in polynomial time. When imposing only very mild conditions on the set of plausible SSB utility functions, it can also be verified in polynomial time whether a given lottery is universally undominated. These conditions capture all notions of ordinal efficiency mentioned in the paper.
- An *SD*-efficient lottery that *SD*-dominates a given lottery can be found in polynomial time.
- For *BD*-efficiency, all considered computational problems can be solved in linear time due to a combinatorial characterization of *BD*-efficiency in terms of undominated sets of vertices in the corresponding Pareto digraph.
- It is possible that there is no *PC*-efficient lottery that *PC*-dominates a given lottery.

The remaining part of the paper is structured as follows. An overview of the literature related to our work is given in Section 2. The formal model is introduced in Section 3 and the main theorem is presented in Section 4. In Section 5, we introduce three notions of ordinal efficiency and discuss their geometric properties in Section 6. Finally, in Section 7, we examine three basic computational problems for varying notions of efficiency. All proofs are deferred to the Appendix.

2. Related work

The notion of *SD*-efficiency was popularized by Bogomolnaia and Moulin (2001) and has received considerable attention in the domain of *random assignment* where agents express preferences over objects and the outcome is a randomized allocation of objects to agents (e.g., Abdulkadiroğlu and Sönmez, 2003; Manea, 2009).[3] The random assignment setting constitutes a subdomain of the more general randomized social choice setting considered in this paper. Each discrete assignment can be seen as an alternative such that each agent is indifferent between all assignments in which

[1] The independence axiom requires that if a lottery p is preferred to a lottery q, then a coin toss between p and a third lottery r is preferred to a coin toss between q and r (with the same coin used in both cases).

[2] The preference reversal phenomenon prescribes that a lottery p is preferred to a lottery q, but the certainty equivalent of p is lower then the certainty equivalent of q.

[3] Bogomolnaia and Moulin use the term *ordinal efficiency* for *SD*-efficiency. In order to distinguish *SD*-efficiency from the other notions of ordinal efficiency considered in this paper, we use *SD-efficiency* as advocated by Thomson (2013) (see also, Cho, 2012; Aziz et al., 2013b).