Universal Pareto Dominance and Welfare for Plausible Utility Functions

Haris Aziz
NICTA and UNSW
Sydney, Australia
haris.aziz@nicta.com.au

Florian Brandl
TU München
Munich, Germany
brandlfl@in.tum.de

Felix Brandt
TU München
Munich, Germany
brandtf@in.tum.de

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.

Keywords: SSB utility, Pareto optimality, social welfare, ordinal efficiency, stochastic dominance
JEL Classifications Codes: D71, D81, C60

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 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,

[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).

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 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.

[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.

- 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.

	existence	convexity	support-dependence	existence of efficient improvements
BD-efficiency	+	+	+	+
ex post efficiency	+	+	+	+
SD-efficiency	+	-	+	+
PC-efficiency	+	-	-	-

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.

The remainder 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 he receives the same object (see, e.g., Aziz et al., 2013a). It can be easily seen that if a lottery maximizes affine welfare for some profile of vNM utility functions it is *SD*-efficient. Bogomolnaia and Moulin (2001) conjectured that the converse is also true within the domain of random assignment, i.e., if a lottery is *SD*-efficient there is some profile of vNM utility functions consistent with the ordinal preferences for which it maximizes affine welfare. This statement, now known as the *ordinal efficiency welfare theorem*, was first proven by McLennan (2002) using a variant of the separating polyhedron hyperplane theorem. Later, constructive proofs were provided by Manea (2008) and Athanassoglou (2011). Dogan and Yildiz (2014) gave a constructive proof for the domain of marriage problems, which is slightly more general than the random assignment domain. All these constructive results follow from our Theorem 6 for the unrestricted social choice domain.

A (non-constructive) generalization of these statements to the unrestricted social choice domain and more general sets of plausible utility functions was proved by Carroll (2010). Carroll's theorem shows that for every universally undominated lottery there exists a utility vector with utility functions from each agent's set of plausible utility functions, such that this lottery maximizes affine welfare. Carroll concludes his paper by stating that "*we have addressed this problem in the context of von Neumann-Morgenstern utility functions over lotteries, but it would be interesting to find non-expected utility models, or more general social choice models, in which analogous results hold.*" (Carroll, 2010, p. 2470). In this paper, we provide a generalization of Carroll's result to skew-symmetric bilinear utility functions as proposed by Fishburn (1982) (see also, Fishburn, 1984b, 1988).

Within the domain of random assignment, a combinatorial characterization of the set of *SD*-efficient lotteries in terms of the acyclicity of a binary relation between objects was given by Bogomolnaia and Moulin (2001) and Katta and Sethuraman (2006). There seems to be no straightforward generalization of this characterization to the general social choice setting.

[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).

3 The Model

Let N be a set of n agents, A a finite set of m alternatives, and $\Delta(A)$ the set of all *lotteries* (or *probability distributions*) over A. We assume that preferences over lotteries are given by *skew-symmetric bilinear (SSB) utility* functions as introduced by Fishburn (1982). An SSB function ϕ_i is a function from $\Delta(A) \times \Delta(A) \to \mathbb{R}$ that is skew-symmetric and bilinear, i.e.,

$$\phi_i(p,q) = -\phi_i(q,p),$$
$$\phi_i(\lambda p + (1-\lambda)q, r) = \lambda\phi_i(p,r) + (1-\lambda)\phi_i(q,r),$$

for all $p, q \in \Delta(A)$ and $\lambda \in [0,1]$. Note that, by skew-symmetry, linearity in the first argument implies linearity in the second argument and that, due to bilinearity, ϕ_i is completely determined by its function values for degenerate lotteries. SSB utility theory is more general than the linear expected utility theory due to von Neumann and Morgenstern (1947) as it does not require independence and transitivity (see Fishburn, 1988, 1984b,d, 1982, for example). Hence, every vNM function u_i can be represented by an SSB function ϕ^{u_i}, where $\phi^{u_i}(p,q) = u_i(p) - u_i(q)$. Whenever we consider ordinal preferences over lotteries, these are given in the form of a binary relation R_i, the strict part of which is denoted by P_i, i.e., $p \; P_i \; q$ iff $p \; R_i \; q$ but not $q \; R_i \; p$ and the symmetric part of which is denoted by I_i, i.e., $p \; I_i \; q$ iff $p \; R_i \; q$ and $q \; R_i \; p$.

Let $\phi = (\phi_1, \dots, \phi_n)$ be a vector of SSB functions. A lottery p (Pareto) dominates another lottery q w.r.t. ϕ if all agents weakly prefer the former to the latter, i.e., $p \; R^\phi \; q$ iff $\phi_i(p,q) \geq 0$ for all $i \in N$. A lottery p is *(Pareto) efficient w.r.t.* ϕ iff

there is no q such that $q \; P^\phi \; p$.

Since P^ϕ is the strict part of R^ϕ, we have that $p \; P^\phi \; q$ iff $\phi_i(p,q) \geq 0$ for all $i \in N$ and $\phi_i(p,q) > 0$ for at least one $i \in N$.

A lottery p *is affine welfare maximizing for* ϕ iff

there is $\lambda \in \mathbb{R}^n_{>0}$ such that $\lambda^T \phi(p,q) \geq 0$ for all $q \in \Delta(A)$.

If agents are endowed with vNM functions, affine welfare maximization is equivalent to the maximization of weighted sums of expected utilities. For fixed ϕ and λ, the existence of a welfare maximizing lottery is guaranteed by the minimax theorem (Fishburn, 1984a).

Now, let $\Phi_1, \dots, \Phi_n$ be non-empty sets of SSB functions and $\Phi = \Phi_1 \times \cdots \times \Phi_n$. A lottery p is *universally (Pareto) undominated w.r.t.* Φ iff

there is no q such that $q \; P^\phi \; p$ for all $\phi \in \Phi$.

Existence of universally undominated lotteries follows from the existence of welfare maximizing lotteries.

4 Efficiency Welfare Theorem

In this section, we will show that, if the sets of plausible SSB functions satisfy certain geometric conditions, universal undominatedness is equivalent to affine welfare maximization for a concrete vector of SSB functions (which in turn is equivalent to efficiency for a concrete vector of SSB functions).

We first state a geometric lemma by Carroll (2010), which will be central to the proof of our main theorem. Let S be a subset of $\mathbb{R}^n$. Then, S is a *polyhedron* if it can be represented as the non-empty intersection of finitely many closed half-spaces. A set $F \subseteq S$ is a *face* of a polyhedron S if there are $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$ such that $F = S \cap \{x \in \mathbb{R}^n \colon a^T x = b\}$. S is a *polytope* if it is the convex hull of finitely many points in $\mathbb{R}^n$. Hence, a polyhedron is a polytope iff it is bounded. The *affine hull* of S is the set of all affine combinations of points in S, i.e.,

$$\mathrm{aff}(S) = \left\{ \sum_{i=1}^{k} a_i s_i \colon\ k \in \mathbb{N},\ s_1, \dots, s_k \in S,\ a_1, \dots, a_k \in \mathbb{R},\ \text{and } \sum_{i=1}^{k} a_i = 1 \right\}.$$

S is *relatively open* if it is open in $\mathrm{aff}(S)$.

Lemma 1 (Carroll, 2010)**.** *Let $U, V \subseteq \mathbb{R}^m$ be non-empty convex sets such that U is relatively open, V is a polyhedron, and $v_0 \in V$. Suppose that, for every $v \in V$, there exists $u \in U$ such that $u^T(v - v_0) \leq 0$. Then there exists $u \in U$ such that $u^T(v - v_0) \leq 0$ for all $v \in V$.*

The main idea for the proof of Theorem 1 is that SSB functions are linear in the first argument. Hence, any SSB function can be identified with a vNM function once the second argument is fixed. In our case, the second argument is the lottery which is assumed to be universally undominated and V is the set of all lotteries. U is the set of functions that map a lottery to its affine welfare level, i.e., the (possibly weighted) sum of all agents' utilities. Now, if v_0 is some universally undominated lottery, we obtain a utility vector u for which this lottery maximizes affine welfare. For the proof we identify SSB functions with matrices in $\mathbb{R}^{m \times m}$ and vNM functions with vectors in $\mathbb{R}^m$.

Theorem 1. *Let $\Phi_1, \dots, \Phi_n$ be non-empty, convex, and relatively open sets of SSB functions and $\Phi = \Phi_1 \times \dots \times \Phi_n$. Then, $p \in \Delta(A)$ is universally undominated w.r.t. Φ iff p is affine welfare maximizing for some $\phi \in \Phi$.*

We note that, as a consequence of the theorem, the case of n agents can be reduced to the case of a single agent by combining the sets of plausible utility functions into one. Let $\Phi = \Phi_1 \times \dots \times \Phi_n$ and $\Phi' = \{\lambda_1 \phi_1 + \dots + \lambda_n \phi_n \colon \phi_i \in \Phi_i \text{ and } \lambda_i > 0 \text{ for all } i \in N\}$. Then, a lottery p is universally undominated w.r.t. Φ if and only if it is universally undominated w.r.t. Φ'. Furthermore, if p maximizes affine welfare for $\phi = (\phi_1, \dots, \phi_n) \in \Phi$ (and weights $\lambda \in \mathbb{R}^n_{>0}$) then p maximizes affine welfare for $\phi' = \lambda_1 \phi_1 + \dots + \lambda_n \phi_n \in \Phi'$. On the other hand, if p maximizes affine welfare for $\phi' \in \Phi'$ where $\phi' = \lambda_1\phi_1 + \cdots + \lambda_n\phi_n$, $\phi_i \in \Phi_i$ and $\lambda_i > 0$ for all $i \in N$, then p maximizes affine welfare for $\phi = (\phi_1, \ldots, \phi_n) \in \Phi$.

By instantiating Φ with a single vector of SSB functions, we obtain the equivalence of efficiency and affine welfare maximization as a corollary.

Corollary 1. *Let ϕ be a vector of SSB functions. Then, p is efficient w.r.t. ϕ iff p is affine welfare maximizing with respect to ϕ.*

The second, less obvious, corollary establishes a geometric property of the set of universally undominated lotteries.

Corollary 2. *Let $\Phi_1, \ldots, \Phi_n$ be non-empty, convex, and relatively open sets of SSB functions. Then, the set of universally undominated lotteries w.r.t. $\Phi_1 \times \cdots \times \Phi_n$ is connected.*

5 Ordinal Efficiency

In the important case when only ordinal preferences between alternatives are known, it is natural to consider plausible SSB functions that are consistent with the ordinal preferences. Three particularly interesting classes of plausible SSB functions are the set of all consistent SSB functions, the set of all consistent vNM functions, and the (unique) canonical consistent SSB function. These sets give rise to three natural extensions of preferences over alternatives to preferences over lotteries and thereby to three notions of ordinal efficiency.

Let $R_i \subseteq A \times A$ be a complete and transitive preference relation of agent i.[4] We will compactly represent a preference relation as a comma-separated list of indifference classes. For example $a \mathrel{P_i} b \mathrel{I_i} c$ is represented by $R_i \colon a, \{b, c\}$.

A *lottery extension* is a function that extends a preference relation $R_i \subseteq A \times A$ to a (possibly incomplete) preference relation over lotteries $R_i \subseteq \Delta A \times \Delta A$ (see, e.g., Cho, 2012; Aziz et al., 2013b). An SSB function ϕ_i is *consistent* with a preference relation R_i if

$$\phi_i(x, y) \geq 0 \quad \text{iff} \quad x \mathrel{R_i} y.$$

A lottery p is *universally preferred* to a lottery q w.r.t. Φ_i iff $\phi_i(p, q) \geq 0$ for all $\phi_i \in \Phi_i$.

Example 1. *Let R_i be a preference relation and Φ_i^{SSB} the set of all SSB functions consistent with R_i. Then, a lottery p is universally preferred to a lottery q w.r.t. Φ_i^{SSB} iff*

$$\text{for all } x, y \in A \text{ with } x \mathrel{P_i} y,\ p(x)q(y) - p(y)q(x) \geq 0. \tag{BD}$$

This equivalence was first shown by Fishburn (1984c). We will refer to this lottery extension as bilinear dominance *(BD). Intuitively, p is BD-preferred to q, denoted $p \mathrel{R_i^{BD}} q$, iff, for every pair of alternatives, the probability that p yields the more preferred alternative and q the less preferred alternative is at least as large as the other way round.*

[4] For BD-efficiency and PC-efficiency transitivity of preferences is not required.

Interestingly, we obtain the same lottery extension even when restricting the set of plausible SSB functions to the set of all consistent weighted linear (WL) *utility functions as introduced by Chew (1983). A WL function is characterized by a vNM function and a non-vanishing weight function. The utility of a lottery is the utility derived by the vNM function weighted according to the weight function. Thus, WL functions are more general than vNM functions, as every vNM function is equivalent to a WL function with constant weight function.*

Example 2. *Let R_i be a preference relation and Φ_i^{vNM} the set of all vNM functions consistent with R_i. As mentioned in Section 3, Φ_i^{vNM} can be conveniently represented by a set of linear utility functions U_i where for each $\phi_i^{vNM} \in \Phi_i^{vNM}$ there is $u_i \in U_i$ such that $\phi_i^{vNM}(p,q) = u_i(p) - u_i(q)$. A lottery p is universally preferred to a lottery q w.r.t. Φ_i iff*

$$\text{for all } y \in A, \sum_{x \in A\colon\, xR_iy} p(x) \geq \sum_{x \in A\colon\, xR_iy} q(x). \tag{SD}$$

This equivalence is well-known and the corresponding lottery extension is referred to as stochastic dominance (SD). *Fishburn (1984c) has shown that the same correspondence also holds for all consistent SSB functions that are monotonic, i.e., increasing in the first argument. Intuitively, p is SD-preferred to q, denoted $p\ R_i^{SD}\ q$, if for each alternative x, the probability that p selects an alternative that is at least as good as x is greater or equal to the probability that q selects such an alternative.*

Example 3. *Let R_i be a preference relation and define $\Phi_i = \{\phi_i^{PC}\}$ by letting*

$$\phi_i^{PC}(x,y) = \begin{cases} 1 & \text{if } x\ P_i\ y, \\ -1 & \text{if } y\ P_i\ x, \\ 0 & \text{otherwise.} \end{cases}$$

Then, a lottery p is universally preferred to a lottery q w.r.t. Φ_i iff

$$\sum_{x,y \in A} p(x)q(y)\phi_i^{PC}(x,y) \geq 0. \tag{PC}$$

We will refer to this lottery extension as pairwise comparison (PC). *Intuitively, p is PC-preferred to q, denoted $p\ R_i^{PC}\ q$, iff the probability that p yields an alternative preferred to the alternative returned by q is at least as large than the other way round. Formally,*

$$p\ R_i^{PC}\ q \quad \text{iff} \quad \sum_{xR_iy} p(x)q(y) \geq \sum_{xR_iy} q(x)p(y).$$

Interestingly, this extension may lead to intransitive preferences over lotteries, even when the preferences over alternatives are transitive (Blyth, 1972; Fishburn, 1988). However, despite the possibility of preference cycles, the minimax theorem implies that every R_i^{PC}

admits at least one most preferred lottery (Fishburn, 1984a). An axiomatic characterization of this extension and empirical evidence supporting the axioms was given by Blavatskyy (2006).

In contrast to R_i^{BD} and R_i^{SD}, R_i^{PC} is complete for every R_i.

In general, the following inclusion relationships can be shown:[5]

$$R_i \subseteq R_i^{BD} \subseteq R_i^{SD} \subseteq R_i^{PC}.$$

The first two inclusions are straightforward whereas the last one follows from the above-mentioned characterization of monotonic SSB functions by Fishburn (1984c).

The following example illustrates the definitions of the extensions above. Consider the preference relation $R_i\colon a, b, c$ and lotteries

$$\begin{array}{rclclclc} \text{p} & = & {}^2\!/\!{}_3 \text{ a} & + & 0 \text{ b} & + & {}^1\!/\!{}_3 \text{ c}, & \\ \text{q} & = & 0 \text{ a} & + & 1 \text{ b} & + & 0 \text{ c} & \text{, and} \\ \text{r} & = & 0 \text{ a} & + & {}^2\!/\!{}_3 \text{ b} & + & {}^1\!/\!{}_3 \text{ c}. & \end{array}$$

Then, $p\ P_i^{PC}\ q$; $p\ P_i^{SD}\ r$; $q\ P_i^{BD}\ r$ (and the relationships implied by the hierarchy of lottery extensions above). No other relationships with respect to P_i^{BD}, P_i^{SD}, and P_i^{PC} hold.

Let $\mathcal{E}$ be a lottery extension and $R = (R_1, \dots, R_n)$ a *preference profile*, i.e., an n-tuple containing a preference relation R_i for every agent $i \in N$. We define (Pareto) dominance w.r.t. $\mathcal{E}$ by letting $p\ R^{\mathcal{E}}\ q$ if $p\ R_i^{\mathcal{E}}\ q$ for all $i \in N$. A lottery p is (Pareto) efficient w.r.t. $\mathcal{E}$ if there is no q such that $q\ P^{\mathcal{E}}\ p$. Hence we obtain three notions of ordinal efficiency.

- p is *BD*-efficient iff it is universally undominated w.r.t. $\Phi_1^{SSB} \times \cdots \times \Phi_n^{SSB}$.
- p is *SD*-efficient iff it is universally undominated w.r.t. $\Phi_1^{vNM} \times \cdots \times \Phi_n^{vNM}$.
- p is *PC*-efficient iff it is universally undominated w.r.t. $\{(\phi_1^{PC}, \dots, \phi_n^{PC})\}$.

Another well-known notion of ordinal efficiency in probabilistic settings is *ex post efficiency*. A lottery p is *ex post* efficient if no alternative in *supp*(p) is Pareto dominated. It is well known that *SD*-efficiency implies *ex post* efficiency. On the other hand, *ex post* efficiency implies *BD*-efficiency. This can be verified with the help of Lemma 2 below.

We thus have a hierarchy of notions of ordinal efficiency where *PC*-efficiency is the strongest notion and *BD*-efficiency is the weakest notion.

$$PC\text{-efficiency} \;\Rightarrow\; SD\text{-efficiency} \;\Rightarrow\; \textit{ex post}\text{ efficiency} \;\Rightarrow\; BD\text{-efficiency}.$$

Applying Theorem 1 to the three classes of plausible SSB functions, we obtain the following corollaries.

[5] In slight abuse of notation, we use R_i to denote both a preference relation over alternatives as well as a preference relation over degenerate lotteries.

Corollary 3. *A lottery p is BD-efficient iff there is $(\phi_1, \dots, \phi_n) \in \Phi_1^{SSB} \times \dots \times \Phi_n^{SSB}$ such that $\sum_{i=1}^n \phi_i(p,q) \geq 0$ for all lotteries q.*

Recall that, for vNM functions, affine welfare maximization is equivalent to the maximization of weighted sums of expected utilities. We thus obtain McLennan's ordinal efficiency welfare theorem as a special case.

Corollary 4 (McLennan, 2002)**.** *A lottery p is SD-efficient iff there is $(u_1, \dots, u_n) \in U_1^{vNM} \times \dots \times U_n^{vNM}$ such that $p \in \arg\max_{q \in \Delta(A)} \sum_{i=1}^n u_i(q)$.*

Corollary 5. *A lottery p is PC-efficient iff there is $\lambda \in \mathbb{R}^n_{>0}$ such that $\sum_{i=1}^n \lambda_i \phi_i^{PC}(p,q) \geq 0$ for all lotteries q.*

6 Geometric properties

In this section, we study basic geometric properties of the sets of efficient outcomes for each of the three notions of ordinal efficiency.

6.1 BD-efficiency

It is straightforward to show that the set of *ex post* efficient lotteries is a face of $\Delta(A)$. It turns out that the set of BD-efficient lotteries is convex and its *closure* is a face of $\Delta(A)$. Furthermore, whether a lottery is BD-efficient only depends on its support.

We show this by proving that BD-efficiency can be characterized using a well-known preference extension from alternatives to *sets* of alternatives that Gärdenfors (1979) attributes to Fishburn (1972).

Definition 1. *Let $S, T \subseteq A$ be two sets of alternatives. Then, S* Fishburn dominates *T, denoted by $S \; R_i^F \; T$, if*

$$\begin{aligned} &\text{for all } x \in S \setminus T,\ y \in T \colon x \; R_i \; y,\\ &\text{for all } x \in T \setminus S,\ y \in S \colon y \; R_i \; x. \end{aligned}$$

we say that S Fishburn dominates T w.r.t. a preference profile R, written $S \; R^F \; T$, if $S \; R_i^F \; T$ for all $i \in N$. Lastly, S is Fishburn efficient *if there is no T such that $T \; P^F \; S$.*

Fishburn's preference extension was studied extensively in the context of strategyproof set-valued social choice functions (Gärdenfors, 1979; Ching and Zhou, 2002; Brandt and Brill, 2011; Sanver and Zwicker, 2012; Brandt and Geist, 2014). The following lemma will be useful for illuminating the geometric and computational properties of the set of BD-efficient lotteries.

Lemma 2. *Let R be a preference profile. A lottery p is BD-efficient w.r.t. R iff $supp(p)$ is Fishburn efficient w.r.t. R.*

R_1: a, b, c Set of BD-efficient lotteries:

R_2: c, a, b $\Delta(\{a, b, c\}) \setminus \{\lambda a + (1 - \lambda) b \colon \lambda \in [0, 1)\}$

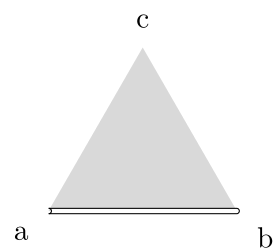

Figure 1: The triangle on the right represents the set of all lotteries on three alternatives. The area shaded in gray denotes the set of BD-efficient lotteries for the preference profile given on the left.

Now we are in a position to characterize the geometric properties of the set of BD-efficient lotteries.

Theorem 2. *The set of BD-efficient lotteries is convex and its closure is a face of* $\Delta(A)$*.*

The set of BD-efficient lotteries itself may not be closed. To see this, consider the preference profile in Figure 1. The lottery $p^\epsilon = (1 - \epsilon) b + \epsilon c$ is BD-efficient for every $\epsilon \in (0, 1)$ since the set $\{b, c\}$ is Fishburn efficient. However, p^0 is not BD-efficient because b is Pareto dominated by a.

Hence, even though BD-efficiency only depends on the support of a lottery, it is not the case that all lotteries whose support is contained in the support of some BD-efficient lottery is BD-efficient as well.

6.2 Plausible vNM functions

If the sets of plausible utility functions only contain vNM functions satisfying the technical conditions from Theorem 1, the set of universally undominated lotteries is closed. Moreover, this set is a union of faces of $\Delta(A)$. This implies that, in this case, universal undominatedness only depends on the support of a lottery. Moreover, if some lottery is universally undominated, any lottery with smaller support is also universally undominated since the set of undominated lotteries is closed.

Theorem 3. *Let* $U_1, \ldots, U_n$ *be non-empty, convex, and relatively open sets of vNM functions. Then, the set of universally undominated lotteries w.r.t.* $U_1 \times \cdots \times U_n$ *is a union of faces of* $\Delta(A)$*.*

Corollary 6. *Whether a lottery is SD-efficient only depends on its support.*

However, in contrast to the set of BD-efficient lotteries, the set of SD-efficient lotteries may be non-convex. To this end, consider the profile in Figure 2.[6] Every alternative

[6]This example is minimal with respect to the number of alternatives, i.e., for every profile with at most three alternatives, the set of SD-efficient lotteries is convex. However, there are profiles with three agents and a non-convex set of SD-efficient lotteries.

$R_1\colon a, c, b, d$
$R_2\colon a, d, b, c$
$R_3\colon b, c, a, d$
$R_4\colon b, d, a, c$

Set of *SD*-efficient lotteries:
$\Delta(\{a, b, c\}) \cup \Delta(\{a, b, d\})$

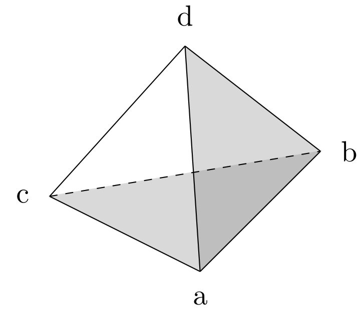

Figure 2: The tetrahedron on the right represents the set of all lotteries on four alternatives. The area shaded in gray denotes the set of *SD*-efficient lotteries for the preference profile given on the left.

is Pareto optimal. Hence $1/2\, c + 1/2\, d$ is *ex post* efficient. However, this lottery is *SD*-dominated by $1/2\, a + 1/2\, b$. Consider the degenerate lotteries c and d. Both are *SD*-efficient. For example, for lottery c, shifting some weight to b is not beneficial for agent 1. If some weight is moved from c to a, then it is not beneficial for agent 3. Although, lotteries c and d are *SD*-efficient, their convex combination $1/2\, c + 1/2\, d$ is not.[7]

We remark that a lottery is *ex post* efficient iff it is in the convex hull of the set of *SD*-efficient lotteries. The argument for the statement is as follows. Consider a lottery p that is *ex post* efficient. It is, by definition, a lottery of Pareto optimal alternatives. Each degenerate lottery corresponding to a Pareto optimal alternative is *SD*-efficient. For the other direction, observe that if a lottery is in the convex hull of *SD*-efficient lotteries, then it is a convex combination of *ex post* efficient degenerate lotteries.

6.3 PC-efficiency

The set of *PC*-efficient lotteries appears to be much more difficult to characterize than the sets of *BD*-efficient and *SD*-efficient lotteries. For example, *PC*-efficiency does not only depend on the support of a lottery. Consider the preference profile in Figure 3. The lottery $p = 1/2\, b + 1/2\, c$ is not *PC*-efficient, since it is *PC*-dominated by $p' = 1/2\, a + 1/2\, d$. It can be checked that $p' \; R_i^{PC} \; p$ for $i = 1, 2, 4$ and $p' \; P_3^{PC} \; p$. But $q = 1/4\, b + 3/4\, c$ has the same support as p and is *PC*-efficient. We show this by solving a system of linear inequalities. Let $x_1, x_2, x_3, x_4 \in \mathbb{R}_{\geq 0}$ such that $x_1 + x_2 + x_3 + x_4 = 1$ and $x_1 a + x_2 b + x_3 c + x_4 d$ is a

[7] The proof of non-convexity of *SD*-efficient lotteries uses a preference profile with linear preferences. For the more general case of preferences that admit indifferences, the fact that random serial dictatorship (which is a convex combination of serial dictators) is not *SD*-efficient (Bogomolnaia and Moulin, 2001; Bogomolnaia et al., 2005; Aziz et al., 2013b) and the observation that outcomes of serial dictatorships are *SD*-efficient already imply that the set of *SD*-efficient lotteries is not convex.

$R_1\colon a, b, c, d$
$R_2\colon d, b, c, a$
$R_3\colon a, b, d, c$
$R_4\colon c, a, d, b$

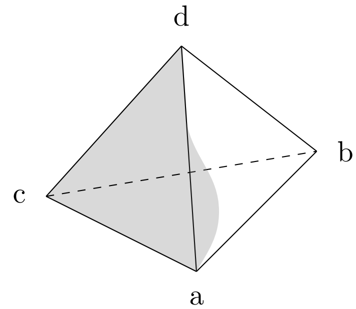

Figure 3: The tetrahedron on the right represents the set of all lotteries on four alternatives. The area shaded in gray was plotted with the help of a computer and denotes the set of PC-efficient lotteries for the preference profile given on the left.

lottery that PC-dominates q. Then we get one equation and one inequality for each agent.

$$x_1, x_2, x_3, x_4 \geq 0 \tag{1}$$
$$x_1 + x_2 + x_3 + x_4 = 1 \tag{2}$$
$$x_1 + {}^3\!/\!_4\, x_2 - {}^1\!/\!_4\, x_3 - x_4 \geq 0 \tag{3}$$
$$-x_1 + {}^3\!/\!_4\, x_2 - {}^1\!/\!_4\, x_3 + x_4 \geq 0 \tag{4}$$
$$x_1 + {}^3\!/\!_4\, x_2 - {}^1\!/\!_4\, x_3 + {}^1\!/\!_2\, x_4 \geq 0 \tag{5}$$
$$-{}^1\!/\!_2\, x_1 - {}^3\!/\!_4\, x_2 + {}^1\!/\!_4\, x_3 - {}^1\!/\!_2\, x_4 \geq 0 \tag{6}$$

Adding up (3) and (4) gives $3x_2 \geq x_3$. Then we plug $3x_2$ instead of x_3 in (6) and get $-{}^1\!/\!_2\, x_1 - {}^1\!/\!_2\, x_4 \geq 0$, which implies that $x_1 = x_4 = 0$. Using $x_1 = x_4 = 0$ in (6) yields $x_3 \geq 3x_2$. Hence, $3x_2 = x_3$, since $3x_2 \geq x_3$ from before. Thus, putting $3x_2 = x_3$ in (2) gives $4x_2 = 1$. So finally $x_2 = {}^1\!/\!_4$ and $x_3 = {}^3\!/\!_4$. Hence, if a lottery PC-dominates q, it is the same lottery.

Interestingly, the example in Figure 3 also shows that the set of PC-efficient lotteries need not be the union of a finite number of polytopes.

Furthermore, the set of PC-efficient lotteries may not be convex. Consider the preference profile in Figure 4. Every alternative is ranked first once. Hence, a and c are PC-efficient. However, ${}^1\!/\!_2\, a + {}^1\!/\!_2\, c$ is PC-dominated by b. Agents 1 and 3 are indifferent between both lotteries and agent 2 strictly prefers b to ${}^1\!/\!_2\, a + {}^1\!/\!_2\, c$.

7 Computational properties

We consider the following three computational problems for various notions of universal domination:

- finding an efficient lottery,

$R_1 \colon a, b, c$
$R_2 \colon b, c, a$
$R_3 \colon c, b, a$

Set of PC-efficient lotteries:
$\Delta(\{a, b\}) \cup \Delta(\{b, c\})$

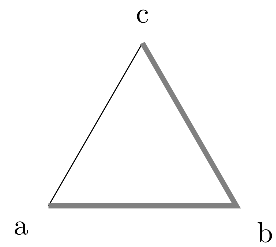

Figure 4: The triangle on the right represents the set of all lotteries on three alternatives. The gray lines denote the set of PC-efficient lotteries for the preference profile given on the left.

- verifying whether a given lottery is efficient, and
- finding an efficient lottery that dominates a given lottery.

Note that these problems are independent from each other. For example, verifying whether a given lottery is efficient can be computationally intractable while an efficient lottery can be found in polynomial time. The third problem (finding outcomes that are not only efficient but also an improvement over a given *status quo*) has received considerable attention in other special domains such as matching and coalition formation (Abdulkadiroğlu and Sönmez, 1999; Chen, 2012; Aziz et al., 2013c).

7.1 Finding universally undominated lotteries

It follows from the "if" statement of Theorem 1 that universally undominated lotteries can easily be found by instantiating concrete SSB functions ϕ_i (and weights λ_i) and computing an affine welfare maximizing lottery using linear programming. Feasibility of the linear program follows from the minimax theorem. This holds without placing any restrictions on the sets of plausible SSB functions.

If the sets of plausible utility functions only contain vNM functions (as in the case of *SD*-efficiency), universally undominated lotteries can be found in linear time by simply fixing a profile of plausible utility functions and computing the degenerate lottery with the highest affine welfare for this profile.

It was shown earlier that *BD*-efficiency of a lottery is equivalent to Fishburn efficiency of its support (Lemma 2). This fact can be used to find a *BD*-efficient lottery in linear time by identifying Fishburn undominated sets of vertices in the Pareto digraph of the given preference profile.[8]

The Pareto relation R_N for a preference profile R is defined as follows. For $x, y \in A$, $x \; R_N \; y$ if $x \; R_i \; y$ for all $i \in N$. The *Pareto digraph* of R is the directed graph (A, R_N). It can be seen from Definition 1 that Fishburn dominance only depends on the Pareto

[8] Note that in the case of intransitive preferences, the Pareto digraph may have cycles.

digraph. Hence, it suffices to find a Fishburn undominated set in the Pareto digraph. This can for example be achieved by identifying a minimal set S such that for all $x \in S$ and $y \in A \setminus S, x \; R_N \; y$. S can be computed in linear time using well-known efficient algorithms for finding strongly connected components (e.g., Tarjan, 1972). Every lottery with support S is then BD-efficient by Lemma 2.

7.2 Verifying whether a given lottery is universally undominated

Using our main theorem (Theorem 1), we present a linear program which, for very general sets of plausible SSB functions, checks whether a given lottery is universally undominated.

Theorem 4. *It can be checked in polynomial time whether a lottery is universally undominated w.r.t. Φ if every Φ_i is given as the non-empty intersection of finitely many hyperplanes and open half spaces in $\mathbb{R}^{m \times m}$. Furthermore, if a lottery is universally undominated the utility functions for which it maximizes affine welfare can be computed in polynomial time.*

Theorem 4 has consequences for checking whether a given lottery is BD-efficient, SD-efficient, or PC-efficient. For example, the set of all consistent vNM functions can be written as the intersection of $\binom{m-1}{2}$ hyperplanes and $m-1$ open half spaces which implies that checking SD-efficiency of a lottery is polynomial-time computable. Every singleton set is the intersection of finitely many hyperplanes and, thus, checking PC-efficiency is polynomial-time computable. This general result can also be used to check BD-efficiency of a lottery. As we will see, there is even a linear-time algorithm for this special case (cf. Theorem 5).

7.3 Finding a universally undominated lottery that universally dominates a given lottery

For some notions of ordinal efficiency we can solve the problem of finding a universally undominated lottery that universally dominates a given lottery in polynomial time. In general, however, it surprisingly turns out that this problem may not even have a solution, i.e., for some universally dominated lottery p there is no universally undominated lottery that universally dominates p.

For the special case of BD-efficiency, we can again leverage the graph-theoretical characterization to find a universally undominated lottery in linear time.[9]

Theorem 5. *For a given lottery p, a BD-efficient lottery that BD-dominates p can be computed in linear time.*

A strengthening of this computational problem is to not only find a universally undominated lottery dominating a given lottery, but to also compute the SSB functions for which

[9] When preferences are intransitive, a BD-efficient lottery that BD-dominates a given lottery may not exist.

the undominated lottery maximizes affine welfare. In general it is unclear how to solve this problem efficiently. For the special cases of BD-efficiency and SD-efficiency, however, we propose a linear program which solves this problem. We present the corresponding theorem for SD-efficiency.

Theorem 6. *For a given lottery p, an SD-efficient lottery q that SD-dominates p can be computed in polynomial time. Furthermore, the utility functions for which q maximizes affine welfare can be computed in polynomial time.*

If the sets of plausible SSB functions are subsets of vNM functions, the problem of finding a universally undominated lottery dominating a given lottery always has a solution due to the transitivity of the induced preferences over lotteries. This, however, does not hold in general. The following example shows that, for a PC-dominated lottery p, there may be no PC-efficient lottery that PC-dominates p.

Consider the following preference profile.

$$\begin{aligned} R_1&\colon \{d,e\}, c, \{a,b\} \\ R_2&\colon b, d, c, a, e \\ R_3&\colon a, e, c, b, d \end{aligned}$$

Let $p = 1/2\, a + 1/2\, b$. We will show that there is no PC-efficient lottery that PC-dominates p. To this end, let q be a lottery that PC-dominates p, i.e., every agent weakly prefers q to p according to the PC preference extension. From these conditions we get the following inequalities.

$$1/2\,(q(b)+q(d)+q(c)) - 1/2\,(q(d)+q(c)+q(a)) - q(e) \geq 0, \tag{7}$$
$$1/2\,(q(a)+q(e)+q(c)) - 1/2\,(q(e)+q(c)+q(b)) - q(d) \geq 0, \tag{8}$$
$$q \geq 0. \tag{9}$$

The inequalities (7) and (8) ensure that agents 2 and 3 weakly prefer q to p. Note that agent 1 weakly prefers every lottery to p. The last inequality is due to the fact that q has to be a lottery. Simplifying (7) and (8), we get

$$\begin{aligned} 1/2\,(q(b)-q(a)) - q(e) &\geq 0, \\ 1/2\,(q(a)-q(b)) - q(d) &\geq 0, \\ q &\geq 0. \end{aligned}$$

Hence, $q(a) = q(b)$ and $q(d) = q(e) = 0$. The resulting lotteries can be parameterized by $\lambda \in [0,1]$ by letting $q_\lambda = \lambda a + \lambda b + (1-2\lambda)c$. Out of these lotteries, the only candidate for an efficient lottery is $q_0 = c$, since q_0 dominates q_λ for every $\lambda > 0$ (agents 2 and 3 are indifferent and agent 1 strictly prefers q_0 to q_λ). But q_0 is dominated by $r = 1/2\, d + 1/2\, e$. Again, agents 2 and 3 are indifferent between both lotteries and agent 1 strictly prefers r to q_0.

It is an interesting open problem whether there always is a path of Pareto improvements from every universally dominated lottery to some universally undominated lottery. We leave this problem as future work.

Acknowledgments

This material is based upon work supported by Deutsche Forschungsgemeinschaft under grants BR 2312/7-2 and BR 2312/10-1 and by NICTA which is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. Parts of this paper were presented at the 15th ACM Conference on Economics and Computation (Palo Alto, June 2014) and the 12th Meeting of the Society for Social Choice and Welfare (Boston, June 2014). The authors thank Paul Stursberg for helpful discussions.

References

A. Abdulkadiroğlu and T. Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999.

A. Abdulkadiroğlu and T. Sönmez. Ordinal efficiency and dominated sets of assignments. *Journal of Economic Theory*, 112(1):157–172, 2003.

M. Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. *Econometrica*, 21(4):503–546, 1953.

S. Athanassoglou. Efficiency under a combination of ordinal and cardinal information on preferences. *Journal of Mathematical Economics*, 47(2):180–185, 2011.

H. Aziz, F. Brandt, and M. Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013a.

H. Aziz, F. Brandt, and M. Brill. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 455–462. IFAAMAS, 2013b.

H. Aziz, F. Brandt, and P. Harrenstein. Pareto optimality in coalition formation. *Games and Economic Behavior*, 82:562–581, 2013c.

P. R. Blavatskyy. Axiomatization of a preference for most probable winner. *Theory and Decision*, 60(1):17–33, 2006.

C. R. Blyth. Some probability paradoxes in choice from among random alternatives. *Journal of the American Statistical Association*, 67(338):366–373, 1972.

A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.

F. Brandt and M. Brill. Necessary and sufficient conditions for the strategyproofness of irresolute social choice functions. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 136–142. ACM Press, 2011.

F. Brandt and C. Geist. Finding strategyproof social choice functions via SAT solving. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1193–1200. IFAAMAS, 2014.

G. Carroll. An efficiency theorem for incompletely known preferences. *Journal of Economic Theory*, 145(6):2463–2470, 2010.

N. Chen. On computing Pareto stable assignments. In *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 384–395, 2012.

S. H. Chew. A generalization of the quasilinear mean with applications to the measurement of income and decision theory resolving the Allais paradox. *Econometrica*, 51(4):1065–1092, 1983.

S. Ching and L. Zhou. Multi-valued strategy-proof social choice rules. *Social Choice and Welfare*, 19(3):569–580, 2002.

W. J. Cho. Probabilistic assignment: A two-fold axiomatic approach. Mimeo, 2012.

B. Dogan and K. Yildiz. Efficiency and stability in probabilistic assignments in marriage problems. Technical report, SSRN, 2014.

P. C. Fishburn. Even-chance lotteries in social choice theory. *Theory and Decision*, 3(1): 18–40, 1972.

P. C. Fishburn. Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26 (1):31–67, 1982.

P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Review of Economic Studies*, 51(4):683–692, 1984a.

P. C. Fishburn. SSB utility theory: An economic perspective. *Mathematical Social Sciences*, 8(1):63–94, 1984b.

P. C. Fishburn. Dominance in SSB utility theory. *Journal of Economic Theory*, 34(1): 130–148, 1984c.

P. C. Fishburn. SSB utility theory and decision-making under uncertainty. *Mathematical Social Sciences*, 8(3):253–285, 1984d.

P. C. Fishburn. *Nonlinear preference and utility theory.* The Johns Hopkins University Press, 1988.

P. Gärdenfors. On definitions of manipulation of social choice functions. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. North-Holland, 1979.

D. M. Grether and C. R. Plott. Economic theory of choice and the preference reversal phenomenon. *American Economic Review*, 69(4):623–638, 1979.

A-K. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131(1):231–250, 2006.

M. J. Machina. Generalized expected utility analysis and the nature of observed violations of the independence axiom. In B. Stigum and F. Wenstop, editors, *Foundations of Utility and Risk Theory with Applications*, chapter 5. Springer, 1983.

M. Manea. A constructive proof of the ordinal efficiency welfare theorem. *Journal of Economic Theory*, 141(1):276–281, 2008.

M. Manea. Asymptotic ordinal inefficiency of random serial dictatorship. *Theoretical Economics*, 4(2):165–197, 2009.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

A. McLennan. Ordinal efficiency and the polyhedral separating hyperplane theorem. *Journal of Economic Theory*, 105(2):435–449, 2002.

M. R. Sanver and W. S. Zwicker. Monotonicity properties and their adaption to irresolute social choice rules. *Social Choice and Welfare*, 39(2–3):371–398, 2012.

R. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972.

W. Thomson. Strategyproof allocation rules. Mimeo, December 2013.

J. von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100:295–320, 1928.

J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947.

Appendix: Proofs

Efficiency Welfare Theorem

Theorem 1. *Let $\Phi_1, \dots, \Phi_n$ be non-empty, convex, and relatively open sets of SSB functions and $\Phi = \Phi_1 \times \cdots \times \Phi_n$. Then, p is universally undominated w.r.t. Φ iff p is affine welfare maximizing for some $\phi \in \Phi$.*

Proof. First, we show that if $p \in \Delta(A)$ is affine welfare maximizing for some $\phi' \in \Phi$ then it is universally undominated w.r.t. Φ. Assume for contradiction that this is not the case, i.e., there is $q \in \Delta(A)$ such that $q \ P^{\phi} \ p$ for all $\phi \in \Phi$. In particular, $q \ P^{\phi'} \ p$. Then by definition of $q \ P^{\phi'} \ p$, we get $\lambda^T \phi'(q,p) > 0$ for every $\lambda \in \mathbb{R}^n_{>0}$ and hence a contradiction.

Now we prove that universal undominatedness of a lottery is sufficient for affine welfare maximization. The set of lotteries $\Delta(A)$ is a polytope in $\mathbb{R}^m$ and, as such, a polyhedron. Let $p \in \Delta(A)$ be universally undominated w.r.t. Φ. For $i \in N$, define $U_i = \Phi_i p = \{\phi p \colon \phi \in \Phi_i\}$. Every $u_i \in U_i$ can be associated with a vNM function where $u_i(q) = q^T u_i$. Since Φ_i is non-empty, convex, and relatively open for all $i \in N$ by assumption, the same holds for U_i for all $i \in N$ by linearity and continuity of matrix multiplication. Let $U \subseteq \mathbb{R}^m$ be the set of weighted sums of plausible utility functions, i.e.,

$$U = \left\{ \sum_{i=1}^{n} \lambda_i u_i \colon \lambda_i > 0 \text{ and } u_i \in U_i \text{ for all } i \in N \right\}.$$

It is easily seen that U is non-empty, convex, and relatively open.
For $u \in U$ and $q \in \Delta(A)$, $u(q) = q^T u$ is the affine welfare of q under u. First note that $p^T u_i = p^T \phi_i p = 0$ for all $u_i \in U_i$ and $i \in N$ by skew-symmetry of every ϕ_i since $p^T \phi_i p = p^T(-\phi_i^T)p = -(p^T \phi_i^T)p = -p^T(p^T \phi_i^T)^T = -p^T \phi_i p$. Hence, $p^T u = 0$ for all $u \in U$. Since p is universally undominated w.r.t. Φ, for every $q \in \Delta(A)$, there are $u_i \in U_i$ such that either $q^T u_i < 0$ for some $i \in N$ or $q^T u_i = 0$ for all $i \in N$. In the former case, choose $u_j \in U_j$ and $\lambda_j > 0$ arbitrarily for $j \in N \setminus \{i\}$ and

$$\lambda_i \geq \frac{1}{q^T u_i} \sum_{j \in N \setminus \{i\}} \lambda_j q^T u_j.$$

Then, for $u = \sum_{i=1}^n \lambda_i u_i$, we have $q^T u \leq 0$. In the latter case, let $u = \sum_{i=1}^n u_i$, which implies $q^T u = 0$. In either case, we have $u \in U$ such that $q^T u \leq 0 = p^T u$.

Since all the premises are met we can apply Lemma 1 to U and $\Delta(A)$ and obtain $u \in U$ such that $q^T u \leq 0 = p^T u$ for all $q \in \Delta A$. Translating things back to SSB functions, we get $u = \sum_{i=1}^n \lambda_i u_i$ such that, for all $q \in \Delta(A)$,

$$0 \geq \sum_{i=1}^{n} \lambda_i q^T u_i = \sum_{i=1}^{n} \lambda_i q^T \phi_i p = \sum_{i=1}^{n} \lambda_i \phi_i(q,p)$$

for some $\phi_i \in \Phi_i$ such that $u_i = \phi_i p$ for all $i \in N$. Thus, p maximizes affine welfare if agents have SSB functions according to $\phi_1, \dots, \phi_n$ as obtained before. $\square$

Geometric Properties

Corollary 2. *Let $\Phi_1, \dots, \Phi_n$ be non-empty, convex, and relatively open sets of SSB functions. Then, the set of universally undominated lotteries w.r.t. $\Phi = \Phi_1 \times \cdots \times \Phi_n$ is connected.*

Proof. We have to show that for two arbitrary universally undominated lotteries p and p', there is a path from p to p' along universally undominated lotteries. By Theorem 1, there are $\phi, \phi' \in \Phi$ such that $\phi = \sum_{i\in N} \lambda_i \phi_i$ and $\phi' = \sum_{i \in N} \lambda_i' \phi_i'$ for $\lambda_i, \lambda_i' > 0$ and $\phi_i, \phi_i' \in \Phi_i$ for all $i \in N$ and

$$\phi(p,q) = p^T \phi q \geq 0 \quad \text{and} \quad \phi'(p',q) = (p')^T \phi' q \geq 0 \quad \text{for all } q \in \Delta(A). \tag{10}$$

By linearity of ϕ and ϕ' in the second argument, (10) is equivalent to

$$\phi(p,a) = p^T \phi a \geq 0 \quad \text{and} \quad \phi'(p',a) = (p')^T \phi' a \geq 0 \quad \text{for all } a \in A.$$

Now, let $\phi^\lambda = \lambda\phi + (1-\lambda)\phi'$ and $u_a^\lambda = \phi^\lambda a$ for all $\lambda \in [0,1]$ and $a \in A$. By the minimax theorem, the set

$$E^\lambda = \{x \in \Delta(A) \colon x^T \phi^\lambda q \geq 0 \text{ for all } q \in \Delta(A)\}$$

is non-empty for all $\lambda \in [0,1]$. With the notation introduced above, we can write

$$E^\lambda = \{x \in \Delta(A) \colon x^T u_a^\lambda \geq 0 \text{ for all } a \in A\}.$$

So E^λ is a non-empty polytope for all $\lambda \in [0,1]$.

First, we show that the closure of $E = \bigcup_{\lambda\in[0,1]} E^\lambda$ is connected. For $X, Y \subseteq \mathbb{R}^m$, let $\text{dist}(X,Y) = \inf\{||x-y|| \colon x \in X, y \in Y\}$. It suffices to show that for all $\epsilon > 0$ there is $\delta_0 > 0$ such that $\text{dist}(E^\lambda, E^{\lambda+\delta}) < \epsilon$ for all $\delta \in (0, \delta_0)$. Intuitively this means that varying the vectors u_a^λ slightly does not lead to a polytope which is "far away" from E^λ. To prove this, assume for contradiction that there is $\epsilon > 0$ such that for all $\delta_0 > 0$ there is $\delta \in (0,\delta_0)$ with $\text{dist}(E^\lambda, E^{\lambda+\delta}) \geq \epsilon$. Hence, there is a sequence $(\delta_k)_{k\in\mathbb{N}} \subseteq \mathbb{R}_{>0}$ such that $\text{dist}(E^\lambda, E^{\lambda+\delta_k}) \geq \epsilon$ for all $k \in \mathbb{N}$. From this, define $(y^k)_{k\in\mathbb{N}} \subseteq \Delta(A)$ such that $y^k \in E^{\lambda+\delta_k}$ for all $k \in \mathbb{N}$. Since $(y^k)_{k\in\mathbb{N}}$ is bounded, it has a cumulation point $y \in \Delta(A)$. Thus, for every $\delta_0 > 0$ there is $k \in \mathbb{N}$ such that $\delta_k \in (0,\delta_0)$, $\text{dist}(E^\lambda, \{y_k\}) \geq \epsilon$, and $||y_k - y||$ is arbitrarily small. In particular, $\text{dist}(E^\lambda, \{y\}) \geq \epsilon$ and, hence, $y \notin E^\lambda$. Thus, there is $a \in A$ such that $y^T u_a^\lambda < 0$. Then, we get

$$\begin{aligned}
(y^k)^T u_a^{\lambda+\delta^k} &= (y^k - y + y)^T u_a^{\lambda+\delta_k} \\
&= (y^k - y)^T u_a^{\lambda+\delta_k} + y^T u_a^{\lambda+\delta_k} \\
&= (y^k - y)^T u_a^{\lambda+\delta_k} + y^T (u_a^{\lambda+\delta_k} - u_a^\lambda + u_a^\lambda) \\
&= (y^k - y)^T u_a^{\lambda+\delta_k} + y^T (u_a^{\lambda+\delta_k} - u_a^\lambda) + y^T u_a^\lambda \\
&\leq ||y^k - y|| \cdot ||u_a^{\lambda+\delta_k}|| + ||y|| \cdot ||u_a^{\lambda+\delta_k} - u_a^\lambda|| + y^T u_a^\lambda < 0
\end{aligned}$$

for k large enough, which contradicts $y^k \in E^{\lambda+\delta_k}$ and establishes connectedness of the closure of E.

Now we prove that E is closed. To this end, let $(x^k)_{k\in\mathbb{N}} \subseteq E$ such that $x^k \to x$ for some $x \in \Delta(A)$. For every $k \in \mathbb{N}$ we can find $\lambda_k \in [0,1]$ such that x^k maximizes affine welfare for ϕ^{λ_k}, i.e., $(x^k)^T\phi^{\lambda_k} q \geq 0$ for all $q \in \Delta(A)$. Since $(\lambda_k)_{k\in\mathbb{N}}$ is bounded, it contains a convergent subsequence. Hence, we may assume without loss of generality that $\lambda_k \to \lambda \in [0,1]$. It follows that, for every $q \in \Delta(A)$,

$$x^T\phi^\lambda q = \lim_{k\to\infty} x^k\phi^{\lambda_k} q \geq 0,$$

since $x^k \in E^{\lambda_k}$ for all $k \in \mathbb{N}$ and, hence, $x \in E^\lambda$. Thus, E is closed and, from before, it follows that E is connected. □

Lemma 2. *Let R be a preference profile. A lottery p is BD-efficient w.r.t. R iff $supp(p)$ is Fishburn efficient w.r.t. R.*

Proof. For the direction from left to right, assume for contradiction that p is BD-efficient and $supp(p)$ is not Fishburn efficient, i.e., there is $S \subseteq A$ which Fishburn dominates $supp(p)$. We distinguish two cases.

Case 1: $S \setminus supp(p) \neq \emptyset$. Hence there is $x \in S$ which Pareto dominates every $y \in supp(p)$. Therefore, the degenerate lottery x BD-dominates p, i.e., $x\ R_i^{BD}\ p$ for all $i \in N$ and $x\ P_i^{BD}\ p$ for some $i \in N$.

Case 2: $S \subsetneq supp(p)$. Let $\lambda = \sum_{x\in S} p(x)$. Define q such that $q(x) = {}^1\!/\!_\lambda\, p(x)$ for $x \in S$ and $q(x) = 0$ for $x \in A\setminus S$. By definition of λ, q is a lottery. We show that q BD-dominates p. Let $i \in N$ and $x, y \in A$ such that $x\ P_i\ y$. Since $S\ R_i^F\ supp(p)$, one of three cases occurs: either $x, y \in S$ or $x, y \in supp(p) \setminus S$ or $x \in S$ and $y \in supp(p) \setminus S$. In the first two cases, we have $q(x)p(y) - q(y)p(x) = 0$. In the third case, $q(x)p(y) - q(y)p(x) = q(x)p(y) > 0$ and, hence, $q\ P_i^{BD}\ p$. So, for every $i \in N$, either there are $x, y \in supp(p)$ such that $x\ P_i\ y$ and we get $q\ P_i^{BD}\ p$ or i is indifferent between all alternatives in $supp(p)$, which implies $q\ R_i^{BD}\ p$. Since $S \subset supp(p)$ Fishburn dominates $supp(p)$, there is at least one $i \in N$ that is not indifferent between all alternatives in $supp(p)$. Thus, p is not BD-efficient which is a contradiction.

For the direction from right to left, assume for contradiction that $supp(p)$ is Fishburn efficient and p is not BD-efficient. By transitivity of R^{BD}, there is some BD-efficient lottery q which BD-dominates p. First, we show that $supp(q) \neq supp(p)$. So assume for contradiction this is not the case and let $P^+ = \{x \in supp(p)\colon p(x) > q(x)\}$. From $p \neq q$, it follows that P^+ is non-empty. Since q BD-dominates p, $x\ R_i\ y$ for all $x \in supp(p) \setminus P^+$ and $y \in P^+$. The intuition is that the alternatives in P^+ are the lowest ranked alternatives among the alternatives in $supp(p)$ for every agent. Let $\lambda = \sum_{x\in supp(p)\setminus P^+} q_x$ and define $q'(x) = {}^1\!/\!_\lambda\, q(x)$ for all $x \in supp(p) \setminus P^+$ and $q'(x) = 0$ otherwise. With a similar argument as in the second case above, this implies that $q'\ R_i^{BD}\ q$ for all $i \in N$ and $q'\ P_i^{BD}\ q$ for some $i \in N$, which contradicts BD-efficiency of q. Hence $supp(p) \neq supp(q)$.

Now we show that $supp(q)$ Fishburn dominates $supp(p)$. If this is not true, one of the following two cases applies. Either there are $x \in supp(q) \setminus supp(p)$, $y \in supp(p)$, and $i \in N$ such that $y \; P_i \; x$. Then, $q_y p_x - q_x p_y = -q_x p_y < 0$, which is a contradiction to the fact that $q \; P_i^{BD} \; p$. Or there are $x \in supp(p) \setminus supp(q)$, $y \in supp(q)$, and $i \in N$ such that $x \; P_i \; y$. Then, $q_x p_y - q_y p_x = -q_y p_x < 0$, which is again a contradiction. Thus, we have $supp(q) \; R_i^F \; supp(p)$ for all $i \in N$. An agent can only be indifferent between $supp(q)$ and $supp(p)$ if he is indifferent between all alternatives in $supp(q) \cup supp(p)$. If every agent is indifferent between all alternatives in $supp(q)$ and $supp(p)$, q cannot *BD*-dominate p. Thus, there is $i \in N$ such that $supp(q) \; P_i^F \; supp(p)$. This contradicts Fishburn efficiency of $supp(p)$. □

Theorem 2. *The set of BD-efficient lotteries is convex and its closure is a face of* $\Delta(A)$.

Proof. First, we show that the set of *BD*-efficient lotteries is convex. By Lemma 2, it suffices to show that the union of two Fishburn efficient sets X and Y is also Fishburn efficient. To this end, let $X, Y \subseteq A$ be Fishburn efficient and assume for contradiction that $X \cup Y$ is not Fishburn efficient. Thus, there exists $Z \subseteq A$ which Fishburn dominates $X \cup Y$. Clearly, $Z \neq X \cup Y$. We distinguish two cases:
Case 1: $Z \setminus (X \cup Y) \neq \emptyset$. Let $z \in Z$. Then $z \; R_i \; x$ for all $x \in X \cup Y$ and all $i \in N$ and $z \; P_i \; x$ for some $x \in X \cup Y$ and $i \in N$. Hence, either X or Y is Fishburn dominated by $\{z\}$, which is a contradiction.

Case 2: $Z \subsetneq (X \cup Y)$. In this case, either $Z \cap X$ is non-empty and Fishburn dominates X or $Z \cap Y$ is non-empty and Fishburn dominates Y. In both cases we get a contradiction.

Now we prove that the closure of set of *BD*-efficient lotteries is a face of $\Delta(A)$. By convexity, there is $p \in \Delta(A)$ such that, for all *BD*-efficient $q \in \Delta(A)$, $supp(q) \subseteq supp(p)$. The claim is that $F = \{q \in \Delta(A) \colon supp(q) \subseteq supp(p)\}$ is the closure of the set of *BD*-efficient lotteries. For every q in the relative interior of F, it holds that $q(a) > 0$ for all $a \in supp(p)$, i.e., $supp(q) = supp(p)$. Again, by Lemma 2, we know that q is *BD*-efficient. By the definition of p, it is clear that no lottery that is not contained in F is *BD*-efficient. Since the closure of the relative interior of F is F, the claim follows. □

Theorem 3. *Let* $U_1, \dots, U_n$ *be non-empty, convex, and relatively open sets of vNM functions. Then, the set of universally undominated lotteries w.r.t.* $U_1 \times \dots \times U_n$ *is a union of faces of* $\Delta(A)$.

Proof. Let $U_1, \dots, U_n$ be non-empty, convex, and relatively open sets of vNM functions and p a lottery that is universally undominated w.r.t. $U_1 \times \dots \times U_n$. Then,

$$F = \{q \in \Delta(A) \colon supp(q) \subseteq supp(p)\}$$

is the smallest face of $\Delta(A)$ in which p is contained. The theorem is proven if we can show that every lottery in F is universally undominated. By Theorem 1, there is $(u_1, \dots, u_n) \in$

$U_1 \times \cdots \times U_n$ for which p maximizes affine welfare. Hence, there is $\lambda \in \mathbb{R}^n_{>0}$ such that

$$u(p) := \sum_{i=1}^{n} \lambda_i u_i(p) \geq \sum_{i=1}^{n} \lambda_i u_i(q) =: u(q)$$

for all $q \in \Delta(A)$. In particular, $u(p) \geq u(a)$ for all $a \in A$. As a weighted sum of linear functions, u is also linear. This implies that $u(p) = u(a)$ for all $a \in supp(p)$. Hence, for every $q \in F$, we have $u(q) = u(p)$, i.e., q maximizes affine welfare for $(u_1, \ldots, u_n)$. This implies that q is universally undominated by Theorem 1. □

Corollary 6. *Whether a lottery is SD-efficient only depends on its support.*

Proof. *SD*-efficiency is equivalent to universal undominatedness w.r.t. $\Phi_i^{vNM} \times \cdots \times \Phi_i^{vNM}$. Hence, by Theorem 3, the set of *SD*-efficient lotteries is the union of faces of $\Delta(A)$. Since two lotteries lie on the same faces of $\Delta(A)$ if they have the same support, the statement follows. □

Computational Properties

Theorem 4. *It can be checked in polynomial time whether a lottery is universally undominated w.r.t. Φ if each Φ_i is given as the non-empty intersection of finitely many hyperplanes and open half spaces in $\mathbb{R}^{m\times m}$. Furthermore, if a lottery is universally undominated, the utility functions for which it maximizes affine welfare can be computed in polynomial time.*

Proof. Let p be the lottery for which we want to check if it is universally undominated. We assume without loss of generality that each Φ_i is conic, i.e., $\phi_i \in \Phi_i$ implies $\lambda_i\phi_i \in \Phi_i$ for all $\lambda_i > 0$ and explain later why this is no restriction. If each Φ_i can be written as the non-empty intersection of finitely many hyperplanes and open half spaces, so can $\Phi = \Phi_1 \times \cdots \times \Phi_n$. Furthermore, Φ is conic. Let $c^1, \ldots, c^k, d^1, \ldots, d^l \in \mathbb{R}^{m\times m\times n}$, $a_1, \ldots, a_k$, $b_1, \ldots, b_l \in \mathbb{R}$, such that

$$\Phi = \bigcap_{j=1}^{k} \{\phi \in \mathbb{R}^{m\times m\times n} \colon \phi^T c^j < a_j\} \cap \bigcap_{j=1}^{l} \{\phi \in \mathbb{R}^{m\times m\times n} \colon \phi^T d^j = b_j\}.$$

As the intersection of finitely many relatively open sets, Φ is relatively open. Note that the distance of a point $y \in \mathbb{R}^{m\times m\times n}$ to the hyperplane $H = \{\phi \in \mathbb{R}^{m\times m\times n} \colon \phi^T c = a\}$ is the absolute value of $\frac{a - y^T c}{||c||^2}$. Consider the following linear program.

$$\begin{aligned}
&\max_{\Phi,\epsilon} && \epsilon && \\
&\text{subject to} && \sum_i \phi_i(p,a) \geq 0 && \text{for all } a \in A, && (11)\\
& && \phi^T c^j + \epsilon \leq a_j && \text{for all } j \in \{1,\ldots,k\}, && (12)\\
& && \phi^T d^j = b_j && \text{for all } j \in \{1,\ldots,l\}, && (13)\\
& && \epsilon \geq 0. &&
\end{aligned}$$

The claim is that p is universally undominated iff $\epsilon > 0$ in every optimal solution of the LP. As a result, if (ϵ^*, ϕ^*) is an optimal solution of the linear program, then p maximizes affine welfare for ϕ^*. By assumption, every Φ_i is non-empty, convex, and relatively open. Hence, we can apply Theorem 1, i.e., p is universally undominated iff it maximizes affine welfare for some profile of plausible SSB functions. Since Φ is assumed to be conic, we do not need weights λ_i for agents' utilities. The constraints (11) are linear in ϕ, since $\phi_i(p, a) = p^T \phi a$. If a lottery does not maximize affine welfare for some profile $(\phi_1, \ldots, \phi_n)$ of SSB functions, then, by linearity of the ϕ_i in the second argument, there is $a \in A$ such that $\sum_i \phi_i(p, a) < 0$. Hence, if (11) is satisfied, p maximizes affine welfare for some profile of SSB functions. The constraints (12) and (13) guarantee that every feasible ϕ is in the closure of Φ.

Now let (ϕ^*, ϵ^*) be an optimal solution of the linear program. If $\epsilon^* > 0$ the inequalities in (12) are satisfied with strict inequality, which implies $\phi^* \in \Phi$. In this case p maximizes affine welfare for ϕ^*. If $\epsilon^* = 0$, there is no $\phi \in \Phi$ for which p maximizes affine welfare. Hence, by Theorem 1, p is universally undominated iff $\epsilon^* > 0$.

Finally, we have to address conicity of the Φ_i. If p maximizes affine welfare for some profile of SSB functions $(\phi_1, \ldots, \phi_n) \in \Phi$ with corresponding weights $\lambda_1, \ldots, \lambda_n$, it maximizes affine welfare for $(\lambda_1\phi_1, \ldots, \lambda_n\phi_n) \in \text{pos}(\Phi_1) \times \cdots \times \text{pos}(\Phi_n)$ and vice versa, where $\text{pos}(\Phi_i) = \{\lambda_i\phi_i \colon \phi_i \in \Phi_i, \lambda_i > 0\}$. Hence, we can assume without loss of generality that all Φ_i are conic. □

Theorem 5. *For a given lottery p, a BD-efficient lottery that BD-dominates p can be computed in linear time.*

Proof. Let p be a lottery and $S = supp(p)$. If there is $a \in A \backslash S$ such that $\{a\}$ Fishburn dominates S in (A, R_N), we compute the set of all such alternatives. Among these alternatives there has to be a Pareto optimal alternative a which is therefore Fishburn undominated in (A, R_N). Then, the degenerate lottery a weakly BD-dominates p and is itself BD-efficient. In the other case, we compute a minimal set T such that for all $x \in T, y \in S \setminus T, x \; R_N \; y$. This can be done in linear time as noted in Section 7.1. We claim that T is itself Fishburn undominated in (A, R_N). By minimality, T is Fishburn undominated in (S, R_N) and, if there is $T' \not\subseteq S$ that Fishburn dominates T, the first case applies, i.e., there is $a \in A \setminus S$ that Fishburn dominates S. Thus, any lottery with support T is BD-efficient by Lemma 2. So we are done if we can find a lottery with support T which BD-dominates p. Note that $T \cap S \neq \emptyset$ as otherwise again the first case applies. To this end, let $\lambda = \sum_{a \in T} p(a)$ and q such that $q(a) = {}^1\!/\!_\lambda \, p(a)$ for all $a \in T$ and $q(a) = 0$ otherwise. As in the proof of Lemma 2, it can be shown that q weakly BD-dominates p. Linear time computability follows from the fact that all steps can be carried out in linear time. □

Theorem 6. *For a given lottery p, an SD-efficient lottery q that SD-dominates p can be computed in polynomial time. Furthermore, the utility functions for which q maximizes affine welfare can be computed in polynomial time.*

Proof. We first show the statement for the case of transitive and anti-symmetric preference relations and then generalize it to transitive preference relations. The proof technique is similar to the one used for the proof of Theorem 1 by Athanassoglou (2011). For this proof, we identify the set of alternatives A with the set $\{1, \dots, m\}$. A lottery p if SD-efficient iff the optimal objective value of following primal linear program is zero.

$$
\begin{aligned}
\min_{q,r,s} \quad & \sum_{i=1}^{n} \sum_{k=1}^{m} -r_{ik} & & \\
\text{subject to} \quad & \sum_{j R_i j_i(k)} q_j - r_{ik} = \sum_{j R_i j_i(k)} p_j & & \text{for all } k \in \{1, \dots, m\},\ i \in N, \qquad (14)\\
& q_j + s_j = 1 & & \text{for all } j \in \{1, \dots, m\},\\
& \sum_{j=1}^{m} q_j = 1, & & \qquad (15)\\
& q \geq 0,\ r \geq 0,\ s \geq 0. & &
\end{aligned}
$$

Let (q^*, r^*, s^*) be an optimal solution of the primal linear program. Then, q^* weakly SD-dominates p and is itself SD-efficient. So we can assume without loss of generality that p is SD-efficient, i.e., the optimal target value of the primal linear program is 0. To construct utility functions for which p maximizes affine welfare, we consider the dual of the linear program above.

$$
\begin{aligned}
\max_{x,y,x} \quad & \sum_{i=1}^{n} \sum_{k=1}^{m} x_{ik} \sum_{j R_i j_i(k)} p_j + y + \sum_{j=1}^{m} z_j & & \\
\text{subject to} \quad & \sum_{i=1}^{n} \sum_{j=\text{rank}_i(k)}^{m} x_{ij} + y + z_k \leq 0 & & \text{for all } k \in \{1, \dots, m\}, \qquad (16)\\
& x \geq 1, & & \\
& y \text{ free variable},\ z \leq 0, & & \qquad (17)
\end{aligned}
$$

where $\text{rank}_i(k)$ is the rank of alternative k in i's preference relation, so $j_i(\text{rank}_i(k)) = k$.

By strong duality, the primal problem has an optimal target value of 0 iff the optimal solution of the dual linear program (x^*, y^*, z^*) satisfies

$$
\sum_{i=1}^{n} \sum_{k=1}^{m} x^*_{ik} \sum_{j R_i j_i(k)} p_j + y^* + \sum_{j=1}^{m} z^*_j = 0. \qquad (18)
$$

Let u denote a profile of vNM functions such that

$$
u_i(j_i(k)) = \sum_{j=k}^{m} x^*_{ij} \text{ for all } k \in \{1, \dots, m\},\ i \in N. \qquad (19)
$$

Since (x^*, y^*, z^*) is dual-feasible, u is consistent with the preference relation R_i for all $i \in N$. Rearranging terms in (18) yields

$$\sum_{i=1}^{n}\sum_{k=1}^{m} u_i(j_i(k))p_{j_i(k)} = -y^* - \sum_{j=1}^{m} z_j^* \Longrightarrow \sum_{i=1}^{n}\sum_{j=1}^{m} u_i(j)p_j = -y^* - \sum_{j=1}^{m} z_j^*. \tag{20}$$

For every $q \in \Delta(A)$, we have

$$\begin{aligned}
\sum_{i=1}^{n}\sum_{j=1}^{m} u_i(j)q_j &= \sum_{i=1}^{n}\sum_{k=1}^{m} u_i(j_i(k))q_{j_i(k)} \overset{(19)}{=} \sum_{i=1}^{n}\sum_{k=1}^{m} q_{j_i(k)} \sum_{j=k}^{m} x_{ij}^* \\
&= \sum_{i=1}^{n}\sum_{k=1}^{m} q_k \sum_{j=\mathrm{rank}_i(k)}^{m} x_{ij}^* = \sum_{k=1}^{m} q_k \sum_{i=1}^{n} \sum_{j=\mathrm{rank}_i(k)}^{m} x_{ij}^* \\
&\overset{(16)}{\leq} \sum_{k=1}^{m} -(y^* + z_k^*)q_k \overset{(15)+(17)}{\leq} -y^* - \sum_{k=1}^{m} z_k^* \overset{(20)}{=} \sum_{i=1}^{n}\sum_{j=1}^{m} u_i(j)p_j.
\end{aligned}$$

The case of anti-symmetric preference relations can be easily modified to allow for indifferences between alternatives. To this end, for every $i \in N$, the set of alternatives A is partitioned into $m_i \in \{1, \ldots, m\}$ indifference classes. For $x \in A$ and $k \in \{1, \ldots, m_i\}$, we say that x is in i's kth indifference class $I_i(k)$, if

$$y \ P_i \ x \text{ for all } y \in \bigcup_{j=1}^{k-1} I_i(k) \quad \text{and} \quad x \ R_i \ y \text{ for all } y \in A \setminus \bigcup_{j=1}^{k-1} I_i(k).$$

We introduce variables r_{ik} where $k \in \{1, \ldots, m_i\}$ and adapt the constraints (14) in the primal problem as follows:

$$\sum_{jR_iI_i(k)} q_j - r_{ik} = \sum_{jR_iI_i(k)} p_j \qquad \text{for all } k \in \{1, \ldots, m_i\},\ i \in N.$$

The corresponding constraints in the dual are modified to

$$\sum_{i=1}^{n} \sum_{j=\mathrm{rank}_i(k)}^{m} x_{ij} + y + z_j \leq 0 \quad \text{for all } j \in I_i(k),\ k \in \{1, \ldots, m_i\}.$$

Consequently the profile of vNM functions u is defined as

$$u_i(j) = \sum_{j=k}^{m_i} x_{ij}^*, \quad \text{for all } j \in I_i(k),\ k \in \{1, \ldots, m_i\},\ i \in N.$$

$\square$