Convex preferences: A new definition

MICHAEL RICHTER
Department of Economics, Royal Holloway, University of London

ARIEL RUBINSTEIN
School of Economics, Tel Aviv University and Department of Economics, New York University

We suggest a concept of convexity of preferences that does not rely on any alge-braic structure. A decision maker has in mind a set of orderings interpreted as evaluation criteria. A preference relation is defined to be convex when it satis-fies the following condition: If, for each criterion, there is an element that is both inferior to b by the criterion and superior to a by the preference relation, then b is preferred to a. This definition generalizes the standard Euclidean definition of convex preferences. It is shown that under general conditions, any strict convex preference relation is represented by a maxmin of utility representations of the criteria. Some economic examples are provided.

KEYWORDS. Convex preferences, abstract convexity, maxmin utility.

JEL CLASSIFICATION. C60, D01.

1. INTRODUCTION

The canonical definition of convex preferences requires that if a is preferred to b, then any convex combination of a and b is also preferred to b. This definition relies on the existence of an algebraic structure attached to the space of alternatives. In this paper, we present a new definition of convex preferences. It has an attractive verbal and intuitive meaning, it generalizes the standard Euclidean notion of convex preferences and it also can be applied to spaces without algebraic structure.

In our approach, the agent has in mind a set of *primitive orderings* $\Lambda = \{\geq_k\}$, where each $\geq_k$ is a complete and transitive binary relation (which may have indifferences) over a set of alternatives X. Each ordering represents a criterion for evaluating the alterna-tives. The agent employs these criteria when forming his preferences. To be Λ-*convex,* a preference is required to satisfy the following consistency requirement: Given any two alternatives a and b, if, for each criterion, there is an element that is (i) inferior by that criterion to b and (ii) preferred to a, then b must be preferred to a.

According to this definition, convexity can be perceived as a scheme of argumenta-tion used by either the agent himself or someone trying to persuade him. The argument is as follows. You should prefer b to a, since for each of your relevant evaluation criteria,

there is an alternative inferior to b by that criterion that you prefer to a. To illustrate, assume that job candidates are evaluated according to research, teaching, and charm. To persuade a colleague that b should be hired rather than a, one needs to demonstrate that there is a candidate c, who is a worse researcher than b and preferred by the colleague to a, that there is a candidate (who may or may not be c) who is a worse teacher than b and is preferred by the colleague to a, and that there is a less charming candidate than b whom the colleague ranks above a.

The concept we introduce depends crucially on the set Λ. Obviously, the same set X endowed with different sets of primitive orderings may have different sets of convex preferences. We think about the orderings in Λ as the building blocks in the agent's formation of preferences (for the related concept of "definable preferences," see Rubinstein (1978, 1998)). The orderings in Λ may describe objective features of the alternatives (such as height, weight, or geographical position), but they may also reflect subjective criteria that the agent employs when ranking the alternatives (such as attractiveness and charisma). In the analysis, we take these orderings to be primitives and explore the preferences that are convex with respect to them. This is in contrast to Richter and Rubinstein (2015) who study general-equilibrium-like environments with the notion of *abstract convexity* (Edelman and Jamison, 1985) and under some conditions induced the primitive relations from the notion of abstract convexity.

There are several reasons why Λ-*convexity* is an attractive concept.

(a) We find the consistency requirement compelling for its own sake. The analysis will clarify its role in shaping preference relations.

(b) For Euclidean spaces, choosing the set Λ to contain all algebraic linear orderings induces the standard convex preferences notion. The new definition allows an extension of the standard definition in two directions: First, it applies to setups where no algebraic structure is specified; second, for spaces *with* algebraic structure, alternative specifications of Λ induce alternative convexity-like requirements.

(c) Our main analytical result states that any Λ-*convex* strict preference relation can be represented by a utility function of the form $\min_k(u_k(x))$, where each u_k is a utility representation of some ordering in Λ. This is a meta-representation theorem. For any set of primitive orderings Λ, it delivers a representation result for the Λ-convex preferences, making it possible to derive both new and known maxmin-like representation results.

2. A NEW DEFINITION OF CONVEX PREFERENCES

DEFINITION 1. Let X be a set and let Λ be a set of primitive orderings on X. We say that a preference relation $\succsim$ (complete and transitive) on X is Λ-*convex* if for every $a, b \in X$, the following condition holds: If for every $\geq_k \in \Lambda$, there is a y_k, such that $b \geq_k y_k$ and $y_k \succsim a$, then $b \succsim a$.

We say that a preference relation $\succsim$ on X is Λ-*strictly-convex* if for every $a, b \in X$, the following stronger condition holds: If for every $\geq_k \in \Lambda$, there is a $y_k \neq b$ such that $b \geq_k y_k$ and $y_k \succsim a$, then $b \succ a$.

We now present some *traits* of Λ-convex preferences.

(I) Every primitive ordering $\geq_l \in \Lambda$ is Λ-convex: If, for every $\geq_k$, there is a y_k such that $b \geq_k y_k$ and $y_k \geq_l a$, then, in particular, for l, there is a y_l such that $b \geq_l y_l$ and $y_l \geq_l a$, and, thus, $b \geq_l a$.

(II) A Λ-convex preference relation must satisfy the weak "Pareto" property: If $b \geq_k a$ for every $\geq_k \in \Lambda$, then $b \succsim a$ (apply the definition with $y_k = a$ for every primitive ordering in Λ). If $\succsim$ on X is Λ-strictly-convex, then it satisfies a stronger version of Pareto: For any two distinct elements a and b, if $b \geq_k a$ for every $\geq_k \in \Lambda$, then $b \succ a$.

(III) If Λ is finite and $\succsim$ is Λ-convex, then for each alternative a, there is a direction $\geq_l$ for which a weak decline cannot be strictly improving (for all $y \neq a$, $a \geq_l y \Rightarrow a \succsim y$): It is impossible that for all $\geq_k$ there is $y_k \neq a$ such that $a \geq_k y_k$ and $y_k \succ a$, since letting y_l be $\succsim$-minimal from among those $\{y_k\}$, then, by Λ-convexity, $a \succsim y_l \succ a$.

Furthermore, for any Λ (even infinite), if $\succsim$ is Λ-strictly-convex, then for each alternative a, there is a direction $\geq_l$ for which a weak decline is strictly disimproving (for all $y \neq a$, $a \geq_l y \Rightarrow a \succ y$): It is impossible that for all $\geq_k$ there is $y_k \neq a$ such that $a \geq_k y_k$ and $y_k \succsim a$, since then it would follow that $a \succ a$.

(IV) For a finite set X, when the primitive orderings are strict, a Λ-convex preference relation can be built inductively as follows: Take an alternative x_1 that is at the bottom of one of the primitive relations and place it at the bottom of $\succsim$. Then take an alternative x_2 that is at the bottom of $X - \{x_1\}$ with respect to one of the primitive orderings and place it strongly or weakly above x_1. Continue this procedure until you exhaust all alternatives. The constructed preference is Λ-convex, since if, for every $\geq_k \in \Lambda$, there is a $y_k \neq b$ such that $b \geq_k y_k$ and $y_k \succsim a$, then at least one such y_k must be picked by the procedure before b, and, thus, $b \succsim y_k$ and $y_k \succsim a$, which implies $b \succsim a$.

(V) For any finite set X, where the primitive orderings are strict, every Λ-convex preference relation $\succsim$ must be consistent with the procedure described in (IV). For each $\geq_k$, let b_k be a $\geq_k$-minimal alternative and let b be $\succsim$-minimal from among $\{b_k\}$. For every other alternative, $a \in X \backslash \{b\}$; for every $\geq_k$, it is the case that $a \geq_k b_k$ and $b_k \succsim b$. Thus, Λ-convexity implies that $a \succsim b$. Therefore, b is $\succsim$-minimal. Let $x_1 = b$ and continue with the set $X - \{x_1\}$ to identify the sequence $x_2, \ldots, x_{|X|}$ such that x_i is $\geq_k$-minimal for some $\geq_k$ and $\succsim$-minimal from among $X - \{x_1, \ldots, x_{i-1}\}$.

(VI) If $\Lambda \supset \Gamma$ and $\succsim$ is Γ-convex, then $\succsim$ is Λ-convex. To see it, suppose that $\succsim$ is Γ-convex. If for every $\geq_k \in \Lambda$, there is a y_k such that $b \geq_k y_k$ and $y_k \succsim a$, then for every $\geq_k \in \Gamma$, there is a y_k, such that $b \geq_k y_k$ and $y_k \succsim a$, and by Γ-convexity, $b \succsim a$. Hence $\succsim$ is Λ-convex.

(VII) If $\succsim$ is Λ-convex, then the sets of $\Lambda \cup \{\succsim\}$-convex and Λ-convex preferences are identical. One side of the statement follows from (VI). To see the converse, suppose that $\succsim'$ is $\Lambda \cup \{\succsim\}$-convex. Assume that for every $\geq_k \in \Lambda$, there is a $y_k \neq b$, such that $b \geq_k y_k$ and $y_k \succsim' a$. Let $y = \min(y_k, \succsim)$. Therefore, $y \succsim' a$ and by the Λ-convexity of $\succsim$, $b \succsim y$. Then $\Lambda \cup \{\succsim\}$-convexity of $\succsim'$ implies that $b \succsim' a$. Thus, $\succsim'$ is Λ-convex.

(VIII) For strict preferences, there is no difference between Λ-convexity and Λ-strict-convexity.

Three examples of Λ-convex orderings follow.

EXAMPLE 1. Let X be a (finite or not) subset of $\mathbb{R}$ and let Λ contain exactly two orderings: the increasing ordering $\geq_I$ and the decreasing ordering $\geq_D$.   ◇

OBSERVATION. *A preference is Λ-strictly convex if and only if it is singled-peaked on X (that is, there are no three alternatives $x < y < z$ such that $x, z \succsim y$).*

PROOF. Suppose that $\succsim$ is singled-peaked. Assume that there are $y_I \neq b$ and $y_D \neq b$ such that $b \geq_I y_I$ and $b \geq_D y_D$, and $y_I, y_D \succsim a$. Then, by single-peakness of $\succsim$, we must have $b \succ y_I$ or $b \succ y_D$ and, thus, $b \succ a$. Thus, $\succsim$ is Λ-strictly-convex.

Alternatively, if $\succsim$ is Λ-strictly-convex, then by trait (III), there are no three alternatives $x < y < z$ such that $x, z \succsim y$.   □

EXAMPLE 2. Let X be a convex and closed subset of $\mathbb{R}^N$. Each nonzero vector v defines an *algebraic linear ordering* by $x \geq_v y$ if $v \cdot x \geq v \cdot y$. Denote the set of all algebraic linear orderings by Ψ. We will show that for continuous preference relations on X, the standard notion of convexity is equivalent to Ψ-convexity.   ◇

OBSERVATION. *The following two statements about a continuous preference relation $\succsim$ are equivalent:*

  (i) *The relation $\succsim$ is convex by the standard definition.*

  (ii) *The relation $\succsim$ is Ψ-convex.*

PROOF. Assume (i). Take two different points $a, b \in X$ such that for every $\geq_k \in \Psi$, there is a $y_k \neq b$ such that $b \geq_k y_k$ and $y_k \succsim a$. We show $b \succsim a$ by contradiction. Suppose $a \succ b$. Since $\succsim$ is continuous and convex, the set $U_\succsim(a) = \{z : z \succsim a\}$ is closed and convex. Thus, by the separating hyperplane theorem, there is some algebraic ordering $\geq_l$ such that b lies strictly below $U_\succsim(a)$. Since $b \geq_l y_l$, it follows that $y_l \notin U_\succsim(a)$ and, therefore, $a \succ y_l$, which is contradiction.

Assume (ii). Take two points a and b such that $b \succsim a$. Then, for any point c between a and b and any algebraic linear ordering $\geq_k$, it is the case that $c \geq_k a$ or $c \geq_k b$ and both a and b are preferred to a. Thus, by the definition of Ψ-convexity, $c \succsim a$.   □

The observation demonstrates that the notion of Λ-convexity generalizes the standard convexity notion for continuous preferences. Notice, however, that other familiar properties of preference relations can also be expressed in the language of Λ-convexity. For example, for the case that X is a convex closed subset of $\mathbb{R}^N$, let Ψ_+ be the set of algebraic linear orderings with nonnegative coefficients. Then one can show that a continuous preference relation $\succsim$ is Ψ_+-convex if and only if $\succsim$ is weakly increasing and convex in the standard sense. Thus, by properly choosing the set of primitive orderings, Λ-convexity can express both convexity and monotonicity.

EXAMPLE 3. Let $X = \mathbb{R}_+^2$ (or $X = \mathbb{R}^2$) and let Λ consist of the two primitive orderings $\geq_R$ ("right") and $\geq_U$ ("up"). The following observation implies that a preference relation that is continuous, Λ-convex, and monotonic (if $y_1 > x_1$ and $y_2 > x_2$, then $(y_1, y_2) \succ (x_1, x_2)$) must have indifference curves that are vertical, horizontal, or right-angled only.   ◇

OBSERVATION. *Any continuous Λ-convex and monotonic preference relation has a utility representation of the form $U(x_1, x_2) = \min(f(x_1), g(x_2))$, where f and g are strictly increasing functions.*

PROOF. By monotonicity, the function $U(x, x) = \frac{e^x}{1+e^x}$ represents $\succsim$ along the main diagonal onto $(1/2, 1)$. This representation can be extended by attaching to each alternative the unique alternative on the main diagonal to which it is indifferent (its existence is guaranteed by monotonicity and continuity).

Define $f(z_1) = \sup(U(z_1, z_2) : z_2 \in \mathbb{R})$ and $g(z_2) = \sup(U(z_1, z_2) : z_1 \in \mathbb{R})$. We first show that $U(y_1, y_2) \equiv \min(f(y_1), g(y_2))$. By definition, $U(y_1, y_2) \leq \min(f(y_1), g(y_2))$. Suppose the inequality is strict for some (y_1, y_2). Then there exists $z_2 > y_2$, $z_1 > y_1$ such that $(y_1, z_2) \succ (y_1, y_2)$ and $(z_1, y_2) \succ (y_1, y_2)$, violating trait (III).

The functions f and g are weakly increasing (because $\succsim$ is monotonic). If f and g are strictly increasing everywhere, then we are done.

If not, without loss of generality (WLOG), suppose that $f(y_1) = f(x_1)$, where $y_1 > x_1$. Then, for every z_2, $U(y_1, z_2) = U(x_1, z_2)$ and, thus, $(y_1, z_2) \sim (x_1, z_2)$. Consequently, by monotonicity, for any $y_2 > x_2$, $(y_1, y_2) \succ (x_1, x_2) \sim (y_1, x_2)$. Therefore, g is strictly increasing everywhere and $f(y_1) = f(x_1) \geq \sup\{g(z_2) : z_2 \in \mathbb{R}\} = 1$. Thus, $f(y_1) = f(x_1) = 1$. Since $\succsim$ is continuous and monotonic, $\{z_1 : f(z_1) = 1\} = [m, \infty)$ for some m.

If $m = -\infty$, then $f \equiv 1$ and for any x_1, x_2, $U(x_1, x_2) = g(x_2)$. Therefore, all indifference curves are horizontal. Let h be a strictly increasing function such that $h(z_1) > 1$ everywhere. Then $\min(h(x_1), g(x_2)) = g(x_2) = U(x_1, x_2)$ is the required representation.

If $m > -\infty$, then define $h(z_1) = f(z_1) + (z_1 - m)^+$. This function is strictly increasing since, for $z_1 \geq m$, the function $(z_1 - m)^+$ is strictly increasing and $f(z_1)$ is weakly increasing, and for $z_1 < m$, we have that $(z_1 - m)^+ = 0$ and $h(z) = f(z)$ is strictly increasing. Thus, h and g form the required representation of $\succsim$.   □

3. A MAXMIN REPRESENTATION OF CONVEX PREFERENCES

We now turn to the main analytical result—a Λ-maxmin utility representation of Λ-convex preferences. Our Λ-maxmin utility representation is of the form

$$U(x) = \min\{U_k(x) | U_k \text{ is a utility representation of } \geq_k \in \Lambda\}.$$

The importance of this presentation is twofold. First, it is an attractive procedure for comparing two alternatives. In the hiring example, each candidate receives a score on research, teaching, and charm, and a candidate is evaluated by his/her worst score. Second, it relates to the previous literature that explores other maxmin representations (see Section 5 for a detailed comparison).

DEFINITION 2. A preference relation $\succsim$ over X has a Λ-*maxmin representation* if, for each $\geq_k$ in Λ, there is a utility representation U_k such that $\min_k U_k(x)$ represents $\succsim$.

Example 3 provided such a representation for a particular context. The existence of such a representation means that we can identify each element in the set X with a vector of numbers in $\mathbb{R}^\Lambda$ such that the following statements hold:

  (i) For each primitive ordering, the values that are attached to the elements in X at the corresponding coordinate respect the primitive ordering's ranking.

  (ii) The minimum value that is attached to an alternative across the different dimensions specifies how the alternatives are ranked.

We first verify that any preference relation that has a Λ-maxmin representation is Λ-convex.

PROPOSITION 1. *If $\succsim$ has a Λ-maxmin representation, then $\succsim$ is Λ-convex.*

PROOF. Suppose that for every primitive ordering $\geq_k$, there is a $y_k \neq b$ such that $b \geq_k y_k$ and $y_k \succsim a$. Then, for every $\geq_l$, $U_l(b) \geq U_l(y_l) \geq \min_k U_k(y_l) \geq \min_k U_k(a)$ and, therefore, $\min_k U_k(b) \geq \min_k U_k(a)$, which implies $b \succsim a$. □

The converse requires more than Λ-convexity. For example, the total indifference is always Λ-convex, but typically does not have a Λ-maxmin representation. Proposition 2 shows that Λ-strict convex preferences have Λ-maxmin representations. As preparation, we need one additional concept. Recall the familiar Euclidean property that for any strictly-convex preference relation and any point x, there is a tangent hyperplane that touches x's indifference curve only at x. This motivates the following definition: Given a preference relation $\succsim$, the set Critical(z) contains every ordering $\geq_k \in \Lambda$ that satisfies the condition "for every $y \neq z$, if $y \succsim z$, then $y >_k z$." Define $C_k = \{z \mid \geq_k \in \text{Critical}(z)\}$.

In the Euclidean context, the set Critical(z) is analogous to the subdifferential of a utility representation of $\succsim$. Recall that a standard strictly-convex preference relation has a nonempty subdifferential at every point. Analogously, if $\succsim$ is a Λ-strictly-convex preference relation, then Critical(z) $\neq \varnothing$ for all z. To see why, if Critical(z) $= \varnothing$ for some $z \in X$, then for every $\geq_k$ there would exist $y_k \in X - \{z\}$ such that $z \geq_k y_k$ and $y_k \succsim z$, which violates trait (III).

PROPOSITION 2. *Let X be a finite set. Any Λ-strictly-convex preference relation $\succsim$ on X has a Λ-maxmin representation.*

PROOF. First notice that the elements of C_k are strictly ordered identically by both $\geq_k$ and $\succsim$: given any two distinct elements $x, y \in C_k$, where $x \succsim y$, we have $x >_k y$ since $y \in C_k$. Moreover, if $x >_k y$, then it must be that $x \succ y$ since $x \in C_k$.

Let U be a utility function representing $\succsim$. For every $\geq_k$, define $U_k(z) = U(z)$ for all $z \in C_k$. Since $\succsim$ and $\geq_k$ give exactly the same ranking over C_k, the function U_k represents $\geq_k$ on C_k.

So as to expand the definitions of U_k to the entire set X, count the elements of C_k as $c_1 >_k \ldots >_k c_L$ and consider the following partitions of X: $D_0 = \{x | x >_k c_1\}$, $D_i = \{x | c_i \geq_k x >_k c_{i+1}\}$, and $D_L = \{x | c_L \geq_k x\}$. Notice that for all $z \in D_i \backslash \{c_i\}$, $c_i \succ z$ since if $c_i \geq_k z$, then $c_i \succ z$ since $c_i \in C_k$. Expand U_k on $D_i \backslash \{c_i\}$ to represent $\geq_k$ with values taken from the interval $(\max\{U(z) : c_i \succ z\}, U(c_i)]$. Therefore, U_k represents $\geq_k$ for all k.

For all $x \in C_k$, $U_k(x) = U(x)$, and for all $x \notin C_k$, $U_k(x) > U(x)$. Since, Critical(x) is always nonempty, it follows that $X = \bigcup C_k$, and so $\min_k U_k(x) = U(x)$ for all x. Recall that U represents $\succsim$ and, thus, $\succsim$ has a Λ-maxmin representation. $\square$

Notice that for weak preferences, Propositions 1 and 2 do not form a complete if-and-only-if characterization because Proposition 1 demonstrates Λ-convexity and Proposition 2 requires Λ-strict-convexity. However, recall that for strict preferences, the concepts of Λ-convexity and Λ-strict-convexity are equivalent (VIII). Thus, for strict preferences, Propositions 1 and 2 together provide an exact equivalence between Λ-convexity and the existence of a Λ-maxmin representation.

EXAMPLE 4 (Monotonic Preferences Over Menus). Let Z be a finite set of alternatives and let X be the set of all nonempty menus of Z. Given a utility function over alternatives $u : Z \to \mathbb{R}$, the preference relation $\geq_u^{\max}$ is defined over X by $A \geq_u^{\max} B$ if $\max_{z \in A} u(z) \geq \max_{z \in B} u(z)$. In words, each menu is evaluated by its u-best alternative. Let $\Gamma^{\max}$ consist of all such induced orderings over X. $\Diamond$

OBSERVATION. *(i) A preference on X, $\succsim$, is $\Gamma^{\max}$-strictly-convex if and only if it is strictly monotonic in the sense that $B \supset A$ implies $B \succ A$.*

(ii) If $\succsim$ is a strictly monotonic preference over menus, then there exists a set U of functions from Z to $\mathbb{R}$ such that

$$A \succsim B \quad \text{if and only if} \quad \min_{u \in U} \max_{z \in A} u(z) \geq \min_{u \in U} \max_{z \in B} u(z).$$

PROOF. (i) Let $\succsim$ be a $\Gamma^{\max}$-strictly-convex preference relation. For any two nested menus $B \supset A$, it is the case that $B \geq_k A$ for every $\geq_k \in \Gamma^{\max}$ and, thus, $B \succ A$ (by (II), the strong Pareto property).

For the other direction, let $\succsim$ be a strictly monotonic preference on X, and let A and B be two menus. Suppose that for every $\geq_k \in \Gamma^{\max}$, there exists $Y_k \neq B$, such that $B \geq_k Y_k$ and $Y_k \succsim A$. Take $\geq_u^{\max} \in \Gamma^{\max}$, where $u(z) = 0$ for all $z \in B$ and $u(z) > 0$ for all $z \notin B$. This ordering bottom-ranks B and all of its subsets and ranks all other sets above it. Thus, $B \geq_u^{\max} Y_u$ implies that $B \supset Y_u$ (inclusion is strict because $Y_u \neq B$) and $B \succ Y_u$ by the strict monotonicity of $\succsim$. Therefore, $B \succ Y_u \succsim A$, which implies that $B \succ A$.

(ii) By part (i), $\succsim$ is $\Gamma^{\max}$-strictly-convex. Then, for every ordering $\geq_k$ in $\Gamma^{\max}$, pick one utility function u_k on Z that represents it. By Proposition 2, there exist a strictly increasing function f_k such that $\min_k (f_k (\max_{z \in A} u_k(z))$ represents $\succsim$. But then $f_k \circ u_k$ also represents $\geq_k$ and $\min_k \max_{z \in A} f_k \circ u_k(z)$ represents $\succsim$. $\square$

This kind of representation can be thought of as a state-dependent maxmin utility. The agent currently does not know his future preferences over Z, but will know them

when he chooses from the menu. He evaluates each menu by its worst possible state. This conclusion was proved by Gorno and Natenzon (2018), who in fact show that any weakly monotonic menu preference $\succsim$ can be represented in this manner. Notice the difference from Kreps (1979), who requires weak monotonicity and an additional submodularity axiom to derive a representation of the form $\sum_u \pi(u) \max_{z \in A} u(z)$, where π is a distribution over utility functions.

EXAMPLE 5 (Betweenness Preferences Over Menus). Let Z be a finite set of alternatives and let X be the set of nonempty menus of Z. Given a function $u : Z \to \mathbb{R}$, the preference relation $\geq_u^{\text{avg}}$ over menus is defined by $A \geq_u^{\text{avg}} B$ if $\text{avg}_{z \in A} u(z) \geq \text{avg}_{z \in B} u(z)$. Let Γ^{avg} consist of all such induced orderings over X.

A nonempty sequence of proper subsets of A (the sequence may contain repetitions) is an *equal cover of* A if there is some positive number m such that each alternative in A is contained in exactly m of the subsets. We say that a preference $\succsim$ satisfies the *equal covering property* if for every equal cover of A, at least one of the sets in the sequence is strictly inferior to A. Clearly, the monotonicity property of Example 4 implies the equal covering property. $\diamond$

OBSERVATION. *For any preference $\succsim$ over X, the following statements hold:*

(i) The preference $\succsim$ is Γ^{avg}-strictly-convex if and only if it satisfies the equal covering property.

(ii) If $\succsim$ satisfies the equal covering property, there exists a set $\{U_k\}$ of functions from Z to $\mathbb{R}$ and a set of increasing functions $\{V_k\}$ such that

$$A \succsim B \quad \text{if and only if} \quad \min_{u_k \in U} V_k \left(\underset{z \in A}{\text{avg}}\, u_k(z) \right) \geq \min_{u_k \in U} V_k \left(\underset{z \in B}{\text{avg}}\, u_k(z) \right).$$

(iii) If $\succsim$ satisfies Gul and Pessendorfer's (2001) set-betweenness axiom [$\forall A, B \subseteq Z$ such that $B \succsim A$, and it is the case that $A \cup B \succsim A$ and $B \succsim A \cup B$], then it is Γ^{avg}-convex.

PROOF. (i) Assume that $\succsim$ is Γ^{avg}-strictly-convex. To show that it satisfies the equal covering property, let $(A_1, \ldots, A_n)$ be an equal cover of a set A and WLOG assume that $A_2, \ldots, A_n \succsim A_1$. Then $\text{avg}\, u(A)$ is a convex combination of $\{\text{avg}\, u(A_i)\}$. (Let $m = |\{i : x \in A_i\}|$. Then, for any given u, $m|A| \text{avg}\, u(A) = m \sum_{a \in A} u(a) = \sum_{i=1}^n \sum_{a \in A_i} u(a) = \sum_{i=1}^n |A_i| \text{avg}\, u(A_i)$ and $\sum_{i=1}^n \frac{|A_i|}{m|A|} = 1$.) Thus, $\text{avg}\, u(A) \geq \text{avg}\, u(A_i)$ for at least one A_i, so $A \geq_u^{\text{avg}} A_i$ and $A_i \succsim A_1$. Therefore, by Γ^{avg}-strict-convexity, $A \succ A_1$.

For the other direction, let $\succsim$ be a preference satisfying the equal covering property. Suppose that for every $\geq_k \in \Gamma^{\text{avg}}$, there exists $Y_k \neq B$, such that $B \geq_k Y_k$ and $Y_k \succsim A$. Our goal is to now show that $B \succ A$. Index all elements in Z as $z_1, \ldots, z_{|Z|}$ and attach to each set $A \subseteq X$, a vector $v(A)$, where $v(A)_i = 1/|A|$ if $z_i \in A$ and $v(A)_i = 0$ otherwise. Notice that for any utility function u, $\text{avg}_{z \in A} u(z) = \hat{u} \cdot v(A)$, where $\hat{u} = (u(z_1), \ldots, u(z_{|Z|}))$. Thus, $B \geq_u^{\text{avg}} Y_u$ means that $\hat{u} \cdot v(B) \geq \hat{u} \cdot v(Y_u)$, and since this holds for every $\hat{u}$, it must be that $v(B)$ is a convex combination of $\{v(Y_k)\}$. All of the $v(Y_k)$ are rational vectors and by a theorem of the alternative (Fishburn (1971), Theorem A), B can be equally covered

Preferences	Underlying utilities
$\{b\} \succ \{a, b\} \succ \{a\}$	$u(a) = 1, u(b) = 2$
$\{a, b\} \succ \{a\} \sim \{b\}$	$u(a) = 1, u(b) = 2; v(a) = 2, v(b) = 1$

TABLE 1. Preferences satisfying the equal covering property and their min avg representation

by a sequence of the Y_k (possibly with repetitions). Therefore, by the equal covering property for at least one Y_k, $B \succ Y_k$ and, thus, $B \succ A$.

(ii) By part (i), $\succsim$ is Γ^{avg}-strictly-convex. Then, for every ordering $\geq_k$ in Γ^{avg}, pick one utility function u_k on Z such that $\text{avg } u_k$ represents it. By Proposition 2, there exist a strictly increasing function V_k such that $\min_k(V_k(\text{avg}_{z \in A} u_k(z))$ represents $\succsim$.

(iii) By induction, the first half of set betweenness implies the following stronger condition: For any sequence of proper subsets of A that covers A (not necessarily an equal cover), A is weakly preferred to at least one of the subsets.   □

To demonstrate the above min avg representation, here are two preferences that satisfy the equal covering property and their representations with $Z = \{a, b\}$ (see Table 1).

4. A MAXMIN REPRESENTATION THEOREM FOR COMPACT METRIC SPACES

In Proposition 2 we proved that when X is finite, any Λ-strictly-convex preference relation has a Λ-maxmin representation. Proposition 4 below is an analogous result (with additional continuity-type restrictions) for compact metric spaces. That result requires a significant amount of technical machinery and, therefore, we first present Proposition 3, which illustrates some of the key ideas in a simpler two-dimensional Euclidean setting.

PROPOSITION 3. *Let X be a compact convex subset of $\mathbb{R}^2$ and let Ψ be the set of all algebraic linear orderings on X. If $\succsim$ is a continuous Ψ-strictly-convex preference relation (not necessarily monotonic), then it has a Ψ-maxmin representation.*

PROOF. We first need to derive some properties of the set C_k—the set of critical points of $\geq_k$.

(i) The set C_k contains an element a_k, which is a $\succsim$-maximal element in the set of $\geq_k$-minimal elements. To see it, note that the set of $\geq_k$-minimal elements is convex and compact, and since $\succsim$ is strictly convex and continuous, an element a_k that is a $\succsim$-maximal element in the set of $\geq_k$-minimal elements exists and is strictly preferred to all other $\geq_k$-minimal elements. Therefore, for any different $z \succsim a_k$, it is the case that $z >_k a_k$ and, thus, a_k belongs to C_k.

(ii) There are no two distinct $x, y \in C_k$ such that $x \sim_k y$ (WLOG $x \succsim y$ and then $y \in C_k$ leads to $x >_k y$).
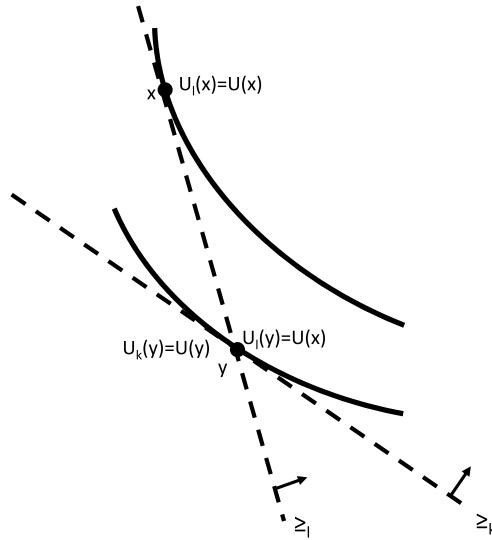
(iii) The set C_k is connected.

FIGURE 1. The construction of U_k, U_l.

(iv) Define $\overline{C_k}$ to be the set of all x such that $x \sim_k y$ for some $y \in C_k$. Any element $x \notin \overline{C_k}$ satisfies $x >_k a_k$ by definition of a_k and satisfies $x >_k y$ for all $y \in C_k$ (since C_k is connected and $\geq_k$ is continuous).

We now define for each k, a function U_k that represents $\geq_k$. In the construction, we use $U : X \to [0, 1]$, a continuous utility representation of $\succsim$ (whose existence is guaranteed by the continuity of $\succsim$). For every $x \in \overline{C_k}$, define $U_k(x) = U(y)$, where y is the unique element in C_k for which $x \sim_k y$. The function U_k represents $\geq_k$ on $\overline{C_k}$. (Let $x, y \in \overline{C_k}$ and let $\hat{x}, \hat{y} \in C_k$ satisfying $x \sim_k \hat{x}$, $y \sim_k \hat{y}$. If $x >_k y$, then $\hat{x} >_k \hat{y}$ and since $\hat{x} \in C_k$, it must be that $\hat{x} \succ \hat{y}$ and, therefore, $U_k(x) = U(\hat{x}) > U(\hat{y}) = U_k(y)$. If $x \sim_k y$, then by (ii), $\hat{x} = \hat{y}$ and $U_k(x) = U(\hat{x}) = U_k(y)$.) For $X - \overline{C_k}$, extend U_k to represent $\geq_k$ with values above 1. Figure 1 illustrates the construction.

It remains to be shown that for every $x \in X$, $\min_k U_k(x) = U(x)$. As mentioned earlier, for each $x \in X$, there is some $\geq_k$, such that $x \in C_k$, and for this ordering $U_k(x) = U(x)$. For any l such that $x \in \overline{C_l} \backslash C_l$, $x \sim_l \hat{x}$ for some $\hat{x} \in C_l$. Since $x \neq \hat{x}$ and $x \sim_l \hat{x}$, it must be that $U(\hat{x}) > U(x)$ and, therefore, $U_l(x) = U(\hat{x}) > U(x)$. Finally, for any l such that $x \notin \overline{C_l}$, $U_l(x) > 1 \geq U(x)$. Thus, $\min_k U_k(x) = U(x)$. □

As mentioned in Example 2, when Λ is the set of all algebraic linear orderings, a continuous preference relation is Λ-strictly-convex if and only if it is strictly convex in the standard sense. Therefore, Proposition 3 demonstrates that any continuous strictly-convex preference relation on a compact convex subset of $\mathbb{R}^2$ has a Λ-maxmin representation of the form $\min_k F_k(x \cdot e_k)$, where e_k is a vector in $\mathbb{R}^2$ that points in the direction of $\geq_k$ and F_k is a strictly increasing function. For Euclidean settings with the standard convexity, Cerreia-Vioglio et al. (2011) establish a similar result that any continuous convex preference relation (not necessarily strict) has a representation using

weakly increasing F_k. That is, they represent a larger class of preferences for $\mathbb{R}^N$, but they are not Λ-maxmin representations since the F_k are $\geq_k$-weakly increasing but do not represent $\geq_k$.

We now prove the existence of a Λ-maxmin representation when X is a compact metric space and (X, Λ) satisfies the following *betweenness* condition: For every $x, y \in X$ and ordering $\geq_l \in \Lambda$, if $y >_l x$, then there exists $z \in X$ such that (i) $y >_l z >_l x$ and (ii) $z \geq_k x$ or $z \geq_k y$ for all other $\geq_k$.

This condition is inspired by the Euclidean setting. In any convex subset of Euclidean space with any collection of linear orderings Λ, an even stronger property holds: For any x and y, and any point z on the line segment between them, z is sandwiched between x and y according to every algebraic linear ordering. An example of a nonconvex set that satisfies the betweenness condition with $\Lambda = \Psi_+$ is a hollow square. The only closed sets in $\mathbb{R}^2$ that satisfy betweenness with Ψ-convexity are the standard convex sets.

PROPOSITION 4. *Let X be a compact metric space and let Λ be a set of continuous primitive orderings that satisfy betweenness. Then any continuous Λ-strictly-convex preference relation $\succsim$ has a Λ-maxmin representation.*

See the Appendix for the proof.

5. COMMENTS

(i) *Maxmin models.* Maxmin functions have a long history, originating with Wald (1950). It is interesting to compare our maxmin representation with the familiar but different maxmin representation of Gilboa and Schmeidler (1989). Let S be a finite set of states and let Z be a set of outcomes. An act is a function $f : S \to Z$. Gilboa and Schmeidler (1989) prove that if a preference relation over the set of acts satisfies certain axioms, then there is a function $u : Z \to \mathbb{R}$ and a set C of probability measures (priors) over S such that the preference relation is represented by $U(f) = \min_{p \in C}\{p \times [u \circ f]\}$. By this approach, an act is transformed subjectively into a point $u \circ f \in \mathbb{R}^S$. Each $p_k \in C$ can be thought of as the algebraic linear function $p_k \times f$ over $\mathbb{R}^S$, and the utility of an act is the minimal value it receives according to these functions.

To obtain a related but different representation in our framework, one can take the alternatives to be objective vectors $f \in \mathbb{R}^S$ and take the set Λ to be a set of orderings represented by functions of the type $p_k \times f$. Then, by Proposition 4, an agent's Λ-strictly-convex preferences can be thought of as him having in mind a set of increasing functions $\{U_k\}$ that he applies to the values $p_k \times f$ and then judges alternatives by $\min_k\{U_1(p_1 \times f), \ldots, U_K(p_K \times f)\}$.

Thus, in our setting, the set of probability measures C is taken as given, in contrast to Gilboa and Schmeidler's (1989) framework. However, this is not the key difference, since any utility function U_k can be taken above the minimum to render the associated probability measure ineffective. The main difference between these two representations is the order in which the functions U and $p_k \times f$ are applied. More importantly, we study a general notion of convex preferences according to which the primitive orderings are

not necessarily algebraic linear functions and where the set of alternatives need not be Euclidean.

(ii) *Social Choice.* Methods for constructing preference relations are the focus of social choice theory, where the social preferences are determined as a function of the individuals' preferences (Arrow and Raynaud, 1986). The notion of Λ-convexity can also be thought of as a social welfare function (SWF) requirement. We say that the social welfare function F is *convex* if, for every profile P, the social preference $F(P)$ is Λ-convex, where Λ consists of all preferences that appear in the profile P. Note that the concept of Λ-convexity is an intra-profile condition. Thus, our analysis can be thought of as being within the single-profile approach in social choice, where a preference relation is built on a specific profile of preference relations without requiring consistency in its definition across various profiles.

Recall that for finite X and strict primitive rankings, Λ-convexity requires that at the bottom of the social ranking lies an element that is at the bottom of one of the individuals' rankings (see (V)). However, the principle by which a convex SWF picks one of the bottom-ranked elements may vary from one profile to another. Two convex social welfare functions that additionally satisfy the standard neutrality, monotonicity, and anonymity conditions are the following.

(a) The *uniform maxmin* SWF is defined by $U(x) = \min_i u_i(x)$, where $u_i(x) = -\operatorname{rank}(x, \geq_i)$. This SWF bottom-ranks all elements that are ranked last by at least one individual. Above them it places those that are ranked next to last by at least one individual but were not ranked last by any agent, and so on.

(b) A *recursive bottom element* SWF: Let $X_1 = X$ and define inductively $M_j = \{x \in X | \text{there is an individual } i \text{ for whom } x \text{ is } \succ_i \text{ -minimal in } X_j\}$ and let $X_{j+1} = X_j - M_j$. Define class$(x) = l$ if $x \in M_l$. The SWF ranks x at least as high as y if class$(x) \geq$ class(y).

This SWF bottom-ranks all elements that are ranked last by at least one individual, then above them it places all the elements that are ranked last by at least one individual among the remaining alternatives, and so on.

Note that the Borda rule is a typical SWF that is not necessarily convex.

(c) Λ-*Concavity.* Dual to our notion of Λ-convexity is the following concept, which we call Λ-*concavity*. A preference relation $\succsim$ on X is Λ-*concave* (Λ-*strict concave)* if for every $a, b \in X$ the following condition holds: If for every $\geq_k \in \Lambda$, there is a $y_k \neq a$ such that $b \succsim y_k$ and $y_k \geq_k a$, then $b \succsim a$ ($b \succ a$).

Recall that the "persuading argument" for $b \succsim a$ that lies behind the notion of Λ-convexity is the existence for any criterion of an alternative that is ranked weakly below b by the criterion and still is weakly superior to a. The persuading argument behind the notion of Λ-concavity is the existence for each criterion of an alternative that is ranked weakly above a by the criterion and still is weakly inferior to b. Both arguments are sound, but apparently it is the former that fits the standard notion of convexity. In the context of choice, the Λ-convexity conditions are arguments for choosing b, whereas Λ-concavity provides arguments for not choosing a.

The reader will now be expecting an attempt to connect the notion of Λ-*strict concavity* to dual representations in the spirit of Propositions 1–4, and we shall not disappoint. For simplicity, we only do so for Proposition 2. We say that a preference relation $\succsim$ over X has a Λ-*maxmax representation* if $\max_k U_k(x)$ represents $\succsim$, where U_k is a utility representation of $\geq_k$.

PROPOSITION 2 (Dual). *Let X be a finite set. Any Λ-strictly-concave preference relation $\succsim$ on X has a Λ-maxmax representation.*

PROOF. For any binary relation R, define the converse binary relation R^T, as $bR^T a$ if aRb. If $\succsim$ is Λ-strictly-concave, then $\succsim^T$ is Λ^T-strictly-convex, where $\Lambda^T = \{\geq_k^T : \geq_k \in \Lambda\}$. By Proposition 2, there exists $\{V_k\}$ such that V_k represents $\geq_k^T$ and $V(x) = \min V_k(x)$ represents $\succsim^T$. Therefore, for every $x \in X$, $-V(x) = -\min V_k(x) = \max -V_k(x)$, and $-V_k$ and $-V$ represent $\geq_k$ and $\succsim$, respectively. Thus, $\succsim$ has a Λ-maxmax representation. $\quad\square$

APPENDIX

PROPOSITION 4. *Let X be a compact metric space and let Λ be a set of continuous primitive orderings satisfying betweenness. Then any continuous Λ-strictly-convex preference relation $\succsim$ has a Λ-maxmin representation.*

PROOF. Let U and V_k be continuous functions representing $\succsim$ and $\geq_k$, respectively, each with a range of $[0, 1]$. Recall that for every $\geq_k \in \Lambda$, the set C_k is defined as $C_k = \{x \mid \geq_k \in \text{Critical}(x)\}$. Notice that there cannot be $x \sim_k y$ such that $x, y \in C_k$. To see why, WLOG suppose $x \succsim y$. Then $x >_k y$ because $y \in C_k$. Let $\text{cl}(C_k)$ denote the topological closure of C_k and define $\overline{\text{cl}(C_k)} = \{y : y \sim_k x$ for some $x \in \text{cl}(C_k)\}$.

We now define, for each k, a function U_k that represents $\geq_k$, such that $\min_k U_k(x)$ represents $\succsim$.

Step 1: Defining U_k on $\overline{\text{cl}(C_k)}$. For each $x \in \overline{\text{cl}(C_k)}$, define $U_k(x) = \max\{U(y) : y \sim_k x\}$. Notice that this definition implies that for all $x \in C_k$, $U_k(x) = U(x)$. To see this, note that since $x \sim_k x$, we have $U_k(x) \geq U(x)$, and by the definition of C_k, there is no y such that $U(y) > U(x)$ and $y \sim_k x$.

Step 2. If $x >_k y$, where $x \in \text{cl}(C_k)$ and y is arbitrary, then $x \succ y$. Suppose to the contrary that $y \succsim x$. By betweenness, there exists w such that $x >_k w >_k y$ and for all other l, $w \geq_l x$ or $w \geq_l y$.

Case (i): $x \succsim w$. Then $y \succsim x \succsim w$ and for every l, either $w \geq_l x$ or $w \geq_l y$, which implies by strict convexity that $w \succ w$, a contradiction.

Case (ii): $w \succ x$. Since $x \in \text{cl}(C_k)$, take a sequence $x_n \to x$ such that $x_n \in C_k$. By the continuity of $\succsim$ and $\geq_k$, for n large enough, it is true that $w \succ x_n$ and $x_n >_k w$, violating $x_n \in C_k$.

Step 3: U_k represents $\geq_k$ on $\overline{\text{cl}(C_k)}$. Let $x, y \in \overline{\text{cl}(C_k)}$. By definition, if $x \sim_k y$, then $U_k(x) = U_k(y)$. Now suppose that $x >_k y$, and consider $w \in \text{cl}(C_k)$ such that $w \sim_k x$ and some $z \in X$ such that $z \sim_k y$ and $U_k(y) = U(z)$. Then $w >_k z$. Since $w \in \text{cl}(C_k)$, then by Step 2, $w \succ z$. Thus, $U_k(x) = U_k(w) \geq U(w) > U(z) = U_k(y)$.

Step 4: Extension of U_k for $x \notin \overline{\text{cl}(C_k)}$. Since $\text{cl}(C_k)$ is a closed subset of a compact set and V_k is continuous, the set of numbers $V_k(\text{cl}(C_k))$ is also closed and $V_k(\overline{\text{cl}(C_k)}) = V_k(\text{cl}(C_k))$ is, therefore, closed as well. Thus, the set $[0, 1] \backslash V_k(\text{cl}(C_k))$ is a collection I_k of disjoint open intervals of the form (a, b), $[0, b)$, or $(a, 1]$.

Case (i). Take $x \notin \overline{\text{cl}(C_k)}$, which according to $\geq_k$ is neither strictly above nor strictly below all members of $\text{cl}(C_k)$. Then $V_k(x)$ lies on a member of I_k of the type $(a, b) = (V_k(\alpha), V_k(\beta))$, where $\alpha, \beta \in \text{cl}(C_k)$. Define $W(V_k(x)) = \max\{U(y) : x \sim_k y\}$. Let $\overline{W} : (a, b] \to (U_k(\alpha), U_k(\beta)]$ be the upper convex envelope of W on $[a, b]$. To see that $\overline{W}$ is strictly increasing, since W is upper hemicontinuous by the theorem of the maximum, it suffices to show that if $\beta >_k x >_k \alpha$, then $U_k(\beta) > W(V_k(x))$. To see this, take y such that $x \sim_k y$ and $W(V_k(x)) = U(y)$. As $\beta \in \text{cl}(C_k)$ and $\beta >_k x \sim y$, then by Step 2, $U(\beta) > U(y)$. Therefore, $U_k(\beta) \geq U(\beta) > U(y) = W(V_k(x))$. Define $U_k(x) = \overline{W}(V_k(x))$. The function U_k represents $\geq_k$ for any x, y, such that $b \geq V_k(x), V_k(y) \geq a$ since U_k is a strict monotonic transformation of V_k. Furthermore, $U_k(x) = \overline{W}(V_k(x)) \geq W(V_k(x)) \geq U(x)$.

Case (ii). There is no $x \notin \overline{\text{cl}(C_k)}$, which according to $\geq_k$ is strictly below all members of $\overline{\text{cl}(C_k)}$. This is because a $\succsim$-maximal element of $V_k^{-1}(0)$ is necessarily in C_k.

Case (iii). Consider $x \notin \overline{\text{cl}(C_k)}$ which according to $\geq_k$ is strictly above all members of $\overline{\text{cl}(C_k)}$. For interval $(a, 1] \in I_k$, we can simply define $W(1) = 2$ and then allow $\overline{W}$ to be the upper convex envelope of W on $[a, 1]$, where W is defined as before. Define $U_k(x) = \overline{W}(V_k(x))$. Since $1 \geq W$, the function $\overline{W}$ is strictly increasing and, therefore, U_k represents $\geq_k$ for x, y, such that $1 \geq V_k(x), V_k(y) \geq a$, and $U_k(x) = \overline{W}(V_k(x)) \geq W(V_k(x)) \geq U(x)$.

Step 5: $U(x) = \min_k U_k(x)$. By construction, for all k, $U_k(x) \geq U(x)$. Because $\succsim$ is Λ-strictly-convex, for every x there is an ordering $\geq_k$ such that $x \in C_k$ and $U_k(x) = U(x)$. Consequently, for all x, $U(x) = \min_k U_k(x)$.   □

REFERENCES

Arrow, Kenneth J. and Herve Raynaud (1986), *Social Choice and Multicriterion Decision-Making*. MIT Press. [1180]

Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci, and Luigi Montrucchio (2011), "Complete monotone quasiconcave duality." *Mathematics of Operations Research*, 36, 321–339. [1178]

Edelman, P.H. and Jamison R.E. (1985), *Geom Dedicata*, 19, 247. [1170]

Fishburn, Peter C. (1971), "The theorem of the alternative in social choice theory." *Operations Research*, 19, 1323–1330. [1176]

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin expected utility with a non-unique prior." *Journal of Mathematical Economics*, 18, 141–153. [1179]

Gorno, Leandro and Paulo Natenzon (2018), "Subjective ambiguity and preference for flexibility." *Journal of Economic Behavior & Organization*, 154, 24–32. [1176]

Gul, Faruk and Wolfgang Pesendorfer (2001), "Temptation and self-control." *Econometrica*, 69, 1403–1435. [1176]

Kreps, David M. (1979), "A representation theorem for 'preference for flexibility'." *Econometrica*, 47, 565–577. [1176]

Richter, Michael and Ariel Rubinstein (2015), "Back to fundamentals: Equilibrium in abstract economies." *American Economic Review*, 105, 2570–2594. [1170]

Rubinstein, Ariel (1978), *Definable Preference Relations—Three Examples.* Research Memorandum 31, Center of Research in Mathematical Economics and Game Theory, The Hebrew University, Jerusalem. [1170]

Rubinstein, Ariel (1998), "Definable preferences: An example." *European Economic Review*, 42, 553–560. [1170]

Wald, Abraham (1950), *Statistical Decision Functions*. Wiley, Oxford, England. [1179]