

Differentiability of preferences without derivatives

MICHAEL RICHTER

Department of Economics, Baruch College, CUNY and Department of Economics, Royal Holloway, University of London

ARIEL RUBINSTEIN

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

ÁRON TÓBIÁS

Maxwell School of Citizenship and Public Affairs, Syracuse University

A new definition of differentiable convex preferences is proposed for spaces that require no algebraic structure. Given a set of primitive orderings Λ , preferences are Λ -convex-differentiable if, at each non-maximal alternative, there is a unique ordering in Λ whose advancement is necessary for preference improvement and whose reduction is sufficient for preference deterioration. As an application, it is shown that differentiability of preferences guarantees the second welfare theorem in a two-agent exchange economy with a finite set of distinct indivisible goods.

KEYWORDS. Preferences, differentiable preferences, convex preferences.

JEL CLASSIFICATION. C6, D1, D6.

1. INTRODUCTION

It appears that many assumptions in economics are adopted for technical convenience rather than because they express aspects of human nature. This observation motivates replacing standard definitions of properties of preference relations with formulations that admit a more natural interpretation. It also leads us to seek formulations that extend the standard framework while also applying to settings in which the underlying grand set is not Euclidean and that do not rely on utility representations.

The approach we use originated in Richter and Rubinstein (2015), and was refined in Richter and Rubinstein (2019, 2024). The *grand set* of alternatives X is enriched with a collection Λ of primitive orderings on X , each representing a relevant consideration. Those papers introduced the concept of Λ -convex preferences: a preference relation on X is Λ -convex if, for every element of X , there exists at least one ordering in Λ , called a *critical ordering*, for which moving up in the ordering is necessary for improvement.

Michael Richter: michael.dan.richter@gmail.com

Ariel Rubinstein: rariel@tauex.tau.ac.il

Áron Tóbiás: adtobias@maxwell.syr.edu

We thank C. Sinan Güntürk for a useful discussion of the topic. Michael Richter gratefully acknowledges funding from the Leverhulme Trust through the Leverhulme International Professorship.

As an example, we considered a reviewer of candidates who vary in research output, teaching ability, and charisma. If the preferences over candidates are Λ -convex, then for any candidate x , there is a criterion, say research, with the property that for a candidate to be superior to x , the candidate must be a better researcher. Such a criterion is a critical ordering but research need not be the only such criterion.

This definition of convex preferences requires no algebraic structure on the space of alternatives, and it captures a sense in which convexity is related to the background considerations used to form the preferences. In [Richter and Rubinstein \(2019\)](#), it was shown that this definition encompasses the standard definition of convex preferences in a Euclidean space by taking Λ to be the set, denoted Ψ , of all “algebraic linear orderings”—those represented by a utility function $\alpha \cdot x$ for some non-zero vector α .

In this paper, we continue this research agenda by offering a new definition of “convex-differentiable (CD) preferences” that does not rely on any algebraic structure. The definition is related to the standard requirement that preferences be represented by a differentiable and quasi-concave utility function with a non-zero gradient at every non-maximal point, but it will be shown to be different. The concept requires that, for each non-maximal alternative, there is a *unique* primitive ordering in Λ that is both necessary for improvement and sufficient for deterioration according to the preferences. We call this ordering the *gradient* of the preferences at that alternative.

Notice that Λ -convex-differentiability adds the following two requirements to Λ -convexity: (i) moving down in the ordering harms preferences, and (ii) there is no other such ordering. Returning to our previous example of ranking candidates, given some Λ -CD preferences over the candidates, if research is the gradient at candidate x , then not only is being a better researcher necessary for being judged above x , but also (i) any candidate who is inferior to x in research is judged to be inferior and (ii) there are no other such criteria.

As will be shown, these properties hold for the standard gradient in the Euclidean setting. However, the standard gradient has another property, namely that small local moves in the direction of the gradient are strictly improving, which we do not require. Given that the space is abstract, often there is no notion of local moves.

Our interest in a new definition of differentiable preferences emerged as a purely conceptual discussion. Accordingly, we study some basic properties of the concept and characterize the set of Λ -convex-differentiable preferences in a number of settings, both finite and infinite. As a demonstration of its potential within a classical economic setting, we consider an exchange economy with two agents and a finite set of distinct indivisible goods. It is well known that in such economies, the second welfare theorem can fail ([Shapley and Scarf, 1974](#)). We let Ω consist of all linear pricing orderings and show that for the two-agent exchange economy, if both agents’ preferences satisfy a variant of Ω -convex-differentiability, then the second welfare theorem holds: any Pareto-optimal allocation is also a competitive equilibrium outcome with linear prices.

Our definition differs from the standard definition of differentiability even when X is a Euclidean space and the possible directions of improvement are linear. For example, the preferences need not be continuous (see Section 3). For that case our approach is closely related to that of [Rubinstein \(2006\)](#), who proposed an alternative definition of differentiable preferences on Euclidean spaces without using the term “utility”. We compare the two definitions in

Section 3. However, our definition is not restricted to Euclidean spaces, depends crucially on the choice of possible directions of improvement and does not rely on a notion of closeness.

There are several conceptual aspects of our notion of differentiability that warrant further discussion and motivation.

Convexity of preferences. Our formal notion of differentiability applies only to preferences that are Λ -convex. Why do we use convexity in order to describe differentiability of preferences? One reason is that convexity of preferences is a ubiquitous assumption in economics and we feel that focusing on convex preferences would serve the purpose of the paper better.

More importantly, we think of Λ as a set of building blocks in the formation of preferences. It provides a language through which we can express features of the alternatives (geometric or otherwise) and properties of preferences. The concept of Λ -convexity includes the standard notion of convex preferences in Euclidean spaces when Λ is taken to be the set of algebraic linear orderings. But with the right choice of primitive orderings, it can express other properties as well: monotonicity, concavity, kinkedness, and much more. Admittedly, it was probably a mistake to use the term Λ -convexity; we did so at the time because it seemed to be a natural extension of the standard notion of convex preferences.

The notion is also applicable to settings with a finite set of discrete goods, which means that we can use Λ -convex-differentiability to study, for example, economies with indivisible goods (see Section 8). This stands in contrast to the usual approach which convexifies the space of alternatives by adding, for example, lotteries, in order to use the standard convexity toolkit. Using the Λ -convexity language, we were able to derive new welfare theorems in many settings (Richter and Rubinstein, 2015). In the current paper, we will see that a variant of Λ -convex-differentiability is powerful and derive a novel second welfare theorem for a two-agent exchange economy with indivisible goods.

Why these particular definitions? The most important reason for our central definition is that it has a simple verbal expression: at each non-maximal point, there is a unique primitive consideration that governs the preferences. The second reason is that, as we will soon see, the definition extends the standard definitions in a wide range of cases.

The reader may think of other definitions of differentiable preferences for abstract spaces. In general, we do not believe that there should necessarily be a single definition of differentiable preferences, and we expect different, related definitions will be useful in different contexts. Indeed, we use two. The central definition focuses on one appealing aspect of differentiability: the uniqueness of the gradient. However, our notion of a gradient can break down at the boundaries, and so we also introduce a weaker notion, especially relevant for bounded spaces (see Sections 7–8).

Turning to the standard framework, recall that the toolkit of mathematical analysis also contains a rich variety of differentiability concepts for multivariate real-valued functions defined on Euclidean spaces. All require the existence of partial derivatives. Partial differentiability requires only the existence of derivatives along the axes; directional differentiability requires the existence of derivatives in every direction; and standard differentiability requires that those derivatives at a point be connected by a linear formula. No single definition is the right one; each captures a particular property.

2. DIFFERENTIABLE CONVEX PREFERENCES

Let X be a non-empty set and let Λ be a set of primitive orderings, each being a preference relation (namely, a complete, transitive, and reflexive binary relation) on X . We denote a generic primitive ordering by $\succeq \in \Lambda$.

Following [Richter and Rubinstein \(2019\)](#), we first provide a definition of Λ -convexity. The concept of Λ -convexity links the underlying set of primitives, Λ , and the preferences, $\succeq$. Suppose that a buyer is considering cell phones that vary in price, resolution, and attractiveness. The notion of Λ -convexity says that for any cell phone c , there must be at least one critical ordering such that every phone which is preferred to c is superior by that ordering. That is, if price is a critical ordering at c , then being cheaper is necessary for a cell phone to be preferred to c , but it does not mean that the buyer prefers every cheaper phone to c .

Formally, given a preference relation $\succeq$ on X we say that $\succeq \in \Lambda$ is *critical* at x if for every y , it is necessary for $y \succ x$ that $y \triangleright x$ (where $\succ$ and $\triangleright$ are the strict parts of $\succeq$ and $\succeq$, respectively). The relation $\succeq$ is Λ -convex if for every $x \in X$ there exists at least one critical ordering in Λ . Note that, if $\succeq$ is critical at x , then any alternative that is weakly $\succeq$ -lower must be weakly $\succeq$ -dispreferred.

Turning to our differentiability notion, we first need a notion of a *gradient*, which is a strengthening of the above criticality property. Again, suppose that a buyer is considering cell phones which vary in their price, resolution, and attractiveness. A gradient is an ordering such that every phone which is preferred to c is superior by that ordering *and* every cell phone which is lower than c by the gradient is *strictly* worse than c . That is, if price is a critical ordering at c , then being cheaper is necessary for a cell phone to be preferred over c *and* being more expensive is sufficient for a cell phone to be strictly dispreferred to c . The change is that for an ordering to be critical, strictly lower-ordered alternatives are weakly dispreferred, but for a gradient, they are strictly dispreferred.

Formally, an ordering $\succeq \in \Lambda$ is a *gradient* of $\succeq$ at x if

- (i) if $y \succ x$, then $y \triangleright x$; and
- (ii) if $x \triangleright y$, then $x \succ y$.

In other words, let $U_{\succ}(x)$ and $L_{\succ}(x)$ denote the *strict upper contour set* and the *strict lower contour set* at x with respect to $\succeq$. A gradient at x is a primitive ordering $\succeq \in \Lambda$ such that $U_{\triangleright}(x) \supseteq U_{\succ}(x)$ and $L_{\triangleright}(x) \subseteq L_{\succ}(x)$.

We now come to the central definition of our differentiability notion. Given the already established framework, it is very simple: at any non-maximal point, the preference relation has a unique gradient. As discussed earlier, uniqueness of the gradient is an appealing and important property of standard differentiable preferences in Euclidean spaces. Formally, a preference relation $\succeq$ is Λ -convex-differentiable (Λ -CD) if, for any non-maximal element $x \in X$, there is a unique gradient $\nabla_{\succeq}(x)$ in Λ . Denote by $D(\Lambda)$ the set of Λ -CD preference relations.

DISCUSSION. What is special about maximal alternatives? In the standard setting, typically, the gradient at a point indicates the “direction of improvement”. But at interior maximal points, no such direction exists, and the gradient is simply the zero vector. Our notion of Λ -CD

preferences is direction-based, and at maximal points there is no unique direction of improvement. We could have augmented our definition by adding to Λ a special zero direction that would serve as the gradient at maximal points, but rather than complicating the structure, we prefer to simplify the definition and require unique gradients only at non-maximal points.

3. Λ -CD AND STANDARD DIFFERENTIABILITY

In the standard Euclidean setting, a preference relation is differentiable if it is representable by a differentiable utility function.

We now show that the definition of Λ -CD extends the standard definition on Euclidean spaces.

Let $X = \mathbb{R}^K$, and Ψ be the set of algebraic linear orderings, that is, $\Psi = \{\succeq_\nu \mid \nu \in \mathbb{R}^K \setminus \{0\}\}$ where $\succeq_\nu$ is the ordering represented by the function $\nu \cdot x$. The following proposition shows that the Ψ -CD notion extends standard differentiability.

PROPOSITION 1. (i) *If a preference relation $\succsim$ on $\mathbb{R}^K$ is monotonic, convex and represented by a differentiable utility function u with a non-zero gradient at all x , then $\succsim$ is Ψ -CD and the gradient ordering is $\succeq_{\nabla u(x)}$.*

(ii) *There are Ψ -CD preferences on $\mathbb{R}^K$ which are monotonic and convex, but are not representable by a continuous utility function u .*

PROOF. (i) Take a preference relation $\succsim$ that is monotonic, convex and represented by a differentiable utility function u with a non-zero gradient at every x .


At any x which is $\succsim$ -non-maximal it needs to be seen that (a) $\succeq_{\nabla u(x)}$ is a gradient ordering and (b) there is no other.

To see (a), by the convexity of the preferences, $U_{\succsim}(x)$ is convex. For every $z \in U_{\succsim}(x)$ and every $\varepsilon \in (0, 1)$, $u(x + \varepsilon(z - x)) \geq u(x)$. Differentiability therefore implies $\nabla u(x) \cdot (z - x) \geq 0$. Thus, since $\nabla u(x) \neq 0$, the hyperplane through x with normal vector $\nabla u(x)$ supports $U_{\succsim}(x)$ at x . If $u(y) > u(x)$, then continuity implies that y is an interior point of $U_{\succsim}(x)$. Hence, $\nabla u(x) \cdot (y - x) > 0$, since an interior point cannot lie on a proper supporting hyperplane. Therefore, $y \succ_{\nabla u(x)} x$. Conversely, if $\nabla u(x) \cdot (z - x) < 0$, then $z \notin U_{\succsim}(x)$, and therefore $u(z) < u(x)$. Thus, $\succeq_{\nabla u(x)}$ is a gradient ordering at x .

To see (b) suppose that there is another gradient $\succeq'$. Notice that for any such gradient, there must be some z such that $x \succ' z$ and yet $z \succ_{\nabla u(x)} x$. By the definition of $\nabla u(x)$, it must be that $\varepsilon z + (1 - \varepsilon)x \succ x$. But this contradicts $\succeq'$ being a gradient because $x \succ' \varepsilon z + (1 - \varepsilon)x$.

(ii) While standard differentiable preferences need to be continuous, this is not true of Ψ -CD preferences: there are Ψ -CD, monotonic, and convex preferences that are *not* continuous. Figure 1 provides an example of such a preference relation on $\mathbb{R}^2$, which even has a utility representation. $\square$

The proposition shows that the set of monotonic, convex, continuous, and standard differentiable preference relations is subsumed by our definition. However, the spirit of the definition is not identical. The standard definition also requires that $x + \varepsilon d \succ x$ for any d satisfying $d \cdot \nabla u(x) > 0$ and sufficiently small $\varepsilon > 0$, a requirement that has no analogue in our definition.



$$U(x_1, x_2) = \begin{cases} x_1 & \text{if } x_1 \leq 0 \\ e^{x_2} & \text{if } x_1 > 0 \end{cases}$$

FIGURE 1. Ψ -CD preferences that are not continuous.

Relationship to Rubinstein's (2006) notion of differentiability. Rubinstein (2006) proposed an alternative definition of differentiable preferences on Euclidean spaces. A revised formulation of the definition appeared in the 2012 version (pp. 55–56).¹

A preference relation $\succsim$ on X is *R-differentiable* if, for every $x \in X$, there exists a vector $v(x) \in \mathbb{R}_+^K \setminus \{0\}$ such that

$$\{z \in \mathbb{R}^K \mid x + z \in X \text{ and } v(x) \cdot z > 0\} = \left\{ z \in \mathbb{R}^K \mid \begin{array}{l} x + z \in X \text{ and } \exists \lambda^* > 0 \text{ such that} \\ x + \lambda z \succ x \text{ for every } \lambda \in (0, \lambda^*) \end{array} \right\}.$$

In other words, at every point there is a vector of local values such that a direction is locally improving precisely when it has positive value according to this vector.

Both Λ -CD and *R-differentiability* seek to extend the standard notion of differentiability without any reference to utility functions. The main difference between them is that Λ -CD is an abstract notion applicable to general domains, whereas *R-differentiability* is formulated for Euclidean spaces. However, even when attention is restricted to convex preferences on Euclidean spaces and $\Lambda = \Psi$, the notions differ in several respects. Most importantly, *R-differentiability* identifies a vector such that a move's having positive value according to this vector is locally both necessary and sufficient for improvement. By contrast, Λ -CD identifies a *unique* gradient ordering such that advancement according to this ordering is necessary for improvement, whereas movement against it is sufficient for deterioration.

The two concepts agree on the interior of X for monotonic, convex and continuous preferences,² but might differ on the boundary. For example, let $X = \mathbb{R}_+^2$. The preference relation represented by $U(x_1, x_2) = x_1 + x_2$ is *R-differentiable* at the origin but is not Λ -CD there, because the gradient ordering is not unique at this boundary point. Our later notion of Λ -TCD addresses this boundary issue.

¹Renou and Schlag (2014) also propose a notion of ordientability that coincides with *R-differentiability* for convex, continuous, and monotone relations on $\mathbb{R}^K$. Diasakos and Gerasimou (2022) discuss the connection between *R-differentiability* and smooth indifference sets.

²That *R-differentiability* implies Ψ -CD is because monotonicity, convexity, and continuity ensure that the vector characterizing local improvements is also a gradient, and interiorness ensures that the gradient is unique. The reverse implication follows from a separating-hyperplane argument.

4. PROPERTIES OF Λ -CD PREFERENCES

(I) *Existence of non-trivial Λ -CD preferences.* Total indifference is trivially a Λ -CD preference because every alternative is maximal. The existence of a non-constant Λ -CD preference is not guaranteed. For example, if X has at least three elements and Λ contains all strict orderings on X , then no non-constant preference relation is Λ -CD since at any non-maximal element there are multiple gradients (including all primitive orderings which rank that element at the bottom).

(II) *Differentiability of primitive orderings.* Natural candidates for Λ -CD preferences are the members of Λ themselves. However, the above example demonstrates that members of Λ need not be Λ -CD. A condition which guarantees that the primitive orderings are themselves Λ -CD is non-nested upper contours: for all x , there are no two orderings $\succeq, \succeq' \in \Lambda$ such that $U_{\succeq}(x) \supseteq U_{\succeq'}(x) \neq \emptyset$. This condition is satisfied by the set of algebraic primitive orderings on $\mathbb{R}^K$.

(III) *A Pareto property.* As noted in Richter and Rubinstein (2019), a weak Pareto property holds even if the preferences $\succsim$ are just Λ -convex. If $a \succeq b$ for all $\succeq \in \Lambda$, then $a \succsim b$ since otherwise there would be no critical ordering at a .

If $\succsim$ is Λ -CD, then a strict Pareto property also holds: if $a \succ b$ for all $\succeq \in \Lambda$ and b is not $\succsim$ -maximal, then $a \succ b$. This is because if a is $\succsim$ -maximal, then we are done. If not, then there is a gradient at a , and b is lower than a by that gradient, so by property (ii) of the definition of Λ -CD preferences, it must be that $a \succ b$.

(IV) *Λ -CD preferences with a unique primitive ordering.* If $\Lambda = \{\triangleright\}$ and $\triangleright$ is strict, then $\triangleright$ is the unique strict Λ -CD preference relation: for any preference relation, if $b \succ a$ and $a \triangleright b$, then there is no gradient at a . The only other Λ -CD preference relations are those formed by creating indifferences at the top and continuing the strict preference thereafter. (Suppose $a \sim b$ are not $\succsim$ -maximal and without loss of generality $a \triangleright b$. Then, the gradient at a has to be $\triangleright$ and therefore $a \succ b$ by property (ii) of the gradient, a contradiction.)

(V) *$D(\Lambda_1)$ and $D(\Lambda_2)$ need not be nested when $\Lambda_2 \supset \Lambda_1$.* We have seen that if Λ_1 consists of a single strict ordering it is Λ_1 -CD, but it is not Λ_2 -CD when Λ_2 contains all orderings. In addition, if Λ_1 consists of a single strict ordering and Λ_2 also includes its opposite, then both orderings in Λ_2 are Λ_2 -CD, but only the original one is Λ_1 -CD.

5. DIFFERENTIABILITY AND OPTIMALITY

Given any primitive ordering $\succeq$ and an alternative e , define $B(e, \succeq) = \{x \in X \mid e \succeq x\}$ to be a *budget set*. We refer to $\succeq$ as the *expenditure ordering* and $x \succeq y$ means that x is (weakly) “more expensive” than y . We make the following analogy to a standard “budget frontier” condition: given any expenditure ordering $\succeq$, any $\succsim$ -maximal alternative from any budget set $B(e, \succeq)$ is $\succeq$ -maximal. The following claim is analogous to a standard price theory result:

CLAIM 1. *Let $\succsim$ be Λ -CD preferences that satisfy the budget frontier analogy and $e \in X$ be not $\succsim$ -maximal. The alternative e is $\succsim$ -maximal in $B(e, \succeq)$ if and only if $\nabla_{\succsim}(e) = \succeq$.*

PROOF. Assume that $\triangleright$ is the gradient at e . If $y \succ e$, then $y \triangleright e$ (by property (i) of the gradient) and thus $y \notin B(e, \triangleright)$. Therefore, e is $\succsim$ -maximal in the budget set $B(e, \triangleright)$.

For the other direction, assume that e is $\succsim$ -maximal from $B(e, \triangleright)$. Then, $\triangleright$ is a gradient at e by the budget frontier analogy, and since the preferences are Λ -CD it is the only gradient at e . $\square$

Claim 2 states the price-theory duality result according to which “expenditure minimization” and “preference maximization” are “equivalent” when preferences are differentiable. Let $e \in X$ be the individual’s “endowment”. Suppose that x^* is $\succsim$ -maximal in $B(e, \triangleright)$ but not $\succsim$ -maximal in X . Define $U_{\succsim}(x^*) = \{y \in X \mid y \succsim x^*\}$ to be the weak upper contour of x^* and suppose that x^{**} is $\triangleright$ -minimal in $U_{\succsim}(x^*)$.

CLAIM 2. *If $\succsim$ is a Λ -CD preference relation satisfying the budget frontier analogy, then x^* and x^{**} are (i) $\succsim$ -indifferent and (ii) $\triangleright$ -indifferent.*

PROOF. (i) By the definition of x^{**} , it holds that $x^{**} \succsim x^*$. Since $e \triangleright x^* \triangleright x^{**}$, it holds that $x^{**} \in B(e, \triangleright)$, and therefore $x^* \succsim x^{**}$.

(ii) Since x^{**} is $\triangleright$ -minimal in $U_{\succsim}(x^*)$, it holds that $x^* \triangleright x^{**}$. Since x^* is $\succsim$ -maximal from $B(e, \triangleright)$, it is also $\succsim$ -maximal in $B(x^*, \triangleright)$. By Claim 1, the gradient at x^* is $\triangleright$. Thus, if $x^* \triangleright x^{**}$, then $x^* \succ x^{**}$, contradicting $x^* \sim x^{**}$. Therefore, x^* and x^{**} must be $\triangleright$ -indifferent. $\square$

6. DIFFERENTIABILITY WITH TWO PRIMITIVES

To better understand the Λ -CD notion, we now consider a case with only two primitives: the horizontal primitive $\rightarrow$ and the vertical primitive $\uparrow$, both on a finite grid and on a square. For the grid, standard differentiability is not well-defined because the space is a finite grid of isolated points. Nevertheless, there is still meaning to phrases such as “direction of improvement” and the analysis will demonstrate how the Λ -CD notion operates on such spaces where standard differentiability cannot. The continuous square case is analyzed to demonstrate the difference between the standard differentiability definition and ours. The only preferences on a square with a differentiable utility function that always has either a rightward or an upward gradient are those represented by either $U(x_1, x_2) = x_1$ or $U(x_1, x_2) = x_2$. In contrast, there is a rich family of Λ -CD preferences according to our definition.

Horizontal and vertical primitives on a finite grid. Let $X = \{0, 1, 2, \dots, K\} \times \{0, 1, 2, \dots, K\}$ and $\Lambda = \{\rightarrow, \uparrow\}$, where $(a_1, a_2) \rightarrow (b_1, b_2)$ if $a_1 \geq b_1$ and $(a_1, a_2) \uparrow (b_1, b_2)$ if $a_2 \geq b_2$. A Λ -CD preference relation with a unique maximum is illustrated in Table 1: the key feature of the example is that for each non-maximal alternative, there is a unique direction in which moving is a strict improvement while moving in the other direction preserves indifference.

CLAIM 3. *Any Λ -CD preferences $\succsim$ with a unique maximum are characterized by a path $p_0 = (0, 0), \dots, p_{2K} = (K, K)$.*

For each $i < 2K$, if $\nabla_{\succsim}(p_i) = \uparrow$ (and analogously if $\nabla_{\succsim}(p_i) = \rightarrow$), then

1. $p_{i+1} = p_i + (0, 1)$

TABLE 1. Illustration of Λ -CD preferences on a two-dimensional grid and two primitive orderings.

	4	1 $\rightarrow$	4 $\rightarrow$	5 $\rightarrow$	6 $\rightarrow$	8
	3	1 $\rightarrow$	4 $\rightarrow$	5 $\rightarrow$	6 $\rightarrow$	7 $\uparrow$
	2	1 $\rightarrow$	3 $\uparrow$	3 $\uparrow$	3 $\uparrow$	3 $\uparrow$
a_2	1	1 $\rightarrow$	2 $\uparrow$	2 $\uparrow$	2 $\uparrow$	2 $\uparrow$
	0	0 $\uparrow$	0 $\uparrow$	0 $\uparrow$	0 $\uparrow$	0 $\uparrow$
		0	1	2	3	4
				a_1		

Note: The numbers in the entry (a_1, a_2) are utility numbers which represent the preferences. The arrow indicates the gradient at each entry. The colored path from the worst alternative to the best uniquely determines the preferences.

2. every point to the right of p_i is indifferent to p_i and
3. every point weakly to the right and strictly above p_i is strictly preferred to p_i .

PROOF. Let $\succsim$ be Λ -CD preferences. By the weak Pareto property of Λ -convexity, the preferences are weakly monotonic. Let $p_0 = (0, 0)$ and assume without loss of generality that its gradient is $\uparrow$.

We will show that every point above it is strictly preferred and every point to the right of it is indifferent. For any $a > 0$, it holds that $(a, 0) \succsim (0, 0)$ by weak monotonicity and $(a, 0) \prec (0, 0)$ by the gradient at $(0, 0)$ being $\uparrow$. Therefore $(0, 0) \sim (a, 0)$ for every $a > 0$.

It cannot be that $(0, 0) \sim (0, b)$ for every b since in that case $\rightarrow$ would also be a gradient at $(0, 0)$ (which is unique by Λ -CD). Thus, there must be some b for which $(0, b) \succ (0, 0)$. If $(0, 1) \sim (0, 0)$, then the gradient at $(0, 1)$ is $\uparrow$ because $(0, b) \succ (0, 0) \sim (0, 1)$. However in that case $(0, 1) \succ (0, 0)$ by property (ii) of the gradient, a contradiction. Thus, $(a, b) \succ (0, 0)$ for all $b > 0$ by weak monotonicity.

Remove the bottom row and the proof continues by induction on the remaining set of alternatives with $p_1 = (0, 1)$. $\square$

To summarize, there is an “increasing” and “continuous” path from $(0, 0)$ to (K, K) such that at every point on it, one direction preserves indifference and the other specifies strict improvement.

The set of Λ -CD preferences is strictly smaller than the set of Λ -convex preferences for this setting. For example, the preferences represented by $\min\{x_1, x_2\}$ are Λ -convex but not Λ -CD because at the point $(1, 1)$ (for example) there is no unique gradient because both primitive orderings are gradients.

The procedure described above to construct the Λ -CD preferences has a natural interpretation when the two orderings reflect the preferences of two agents, 1 and 2. The preferences are constructed sequentially. First one agent, say i_1 , is selected and all i_1 -worst alternatives are assigned to the same indifference class and are placed at the bottom of the ranking. Then, an agent i_2 (who may be the same as i_1) is selected and all i_2 -worst alternatives among the as-yet-unranked alternatives are assigned to be indifferent, strictly above those already ranked. And so on, until all alternatives are ranked. In other words, the preferences can be thought of as the outcome of a sequential procedure according to which, at each stage, rather than stating what is best, an agent is selected and specifies his worst alternatives, which are then placed above all previously ranked alternatives.




FIGURE 2. Λ -CD preferences. The boundary between the regions U and R is indicated in blue and the preferences are strictly increasing along this path. The partition U, R indicates the gradient (if $x \in U$ the gradient is $\uparrow$, otherwise it is $\rightarrow$), which is displayed as arrows. Indifference lines are indicated by dashed lines and they are ranked by where they touch the path. If two indifference lines touch the path at the same point, then they are ranked by that point's membership in U or R (the red dashed indifference lines at a point are preferred to the gray).

Horizontal and vertical primitives in the plane. Let $X = [0, 1]^2$ and $\Lambda = \{\rightarrow, \uparrow\}$. Claim 4 characterizes all Λ -CD preferences. It demonstrates that the set of Λ -CD preferences is much richer than the set of standard differentiable preferences for which the gradient at each point is either $(1, 0)$ or $(0, 1)$ which includes only the two preferences represented by the utility functions x_1 and x_2 . Figure 2 illustrates such a preference relation.

CLAIM 4. *Let $\succsim$ be a monotonic preference relation on X with a unique maximum. Then $\succsim$ is Λ -CD if and only if $X \setminus \{(1, 1)\}$ can be partitioned into two sets U and R such that:*

- (i) *if $x \in U$, then any point weakly below and to the right of x is also in U ; equivalently, if $x \in R$, then any point weakly above and to the left of x is also in R ;*
- (ii) *any point $(a, b) \in U$ lies on a horizontal indifference line, denoted H_b , connecting it to $(1, b)$ and any point $(a, b) \in R$ lies on a vertical indifference line, denoted V_a , connecting it to $(a, 1)$;*
- (iii) *for two horizontal indifference lines H_a and H_b the former is preferred if and only if $a > b$, and likewise for any two vertical indifference lines V_a and V_b ; to compare V_a with H_b , consider the point (a, b) : if it is in U , then $V_a \succ H_b$ and if it is in R , then $H_b \succ V_a$.*

PROOF. Obviously the preferences with the structure described in the claim are Λ -CD with $\nabla_{\succsim}(x) = \uparrow$ if $x \in U$ and $\nabla_{\succsim}(x) = \rightarrow$ if $x \in R$. In the other direction, consider a Λ -CD preference relation $\succsim$. Let $R = \{x \in X \setminus \{(1, 1)\} \mid \nabla_{\succsim}(x) = \rightarrow\}$ and $U = \{x \in X \setminus \{(1, 1)\} \mid \nabla_{\succsim}(x) = \uparrow\}$. Notice that it is impossible for there to be points $x \in R$ and $y \in U$ such that y is weakly up and to the left of x . If such an x and y existed, then by property (i) of the gradient, $y \succsim x$ because $\nabla_{\succsim}(y) = \uparrow$ and likewise $x \succsim y$ because $\nabla_{\succsim}(x) = \rightarrow$. But by property (ii) of the gradient, at least one preference must be strict. Therefore, if $x \in U$, then any point below and to the right of x is in U (and analogously for $x \in R$).

For any x , if $\nabla_{\succsim}(x) = \uparrow$ then any point y to the right of x must be indifferent to x (it must be that $y \succsim x$ by the weak Pareto property and it must be that $y \precsim x$ by $\nabla_{\succsim}(x) = \uparrow$). An analogous argument applies for any x with $\nabla_{\succsim}(x) = \rightarrow$.

Finally, suppose that $(1, b) \in U$ and $(a, 1) \in R$ where $a, b \neq 1$. Without loss of generality, assume that $(a, b) \in U$. By monotonicity, $(a, 1) \succ (a, b)$. If $(a, 1) \sim (a, b)$, then all of the points on the L-shaped broken line segment traced out by the points $(1, b)$, (a, b) , and $(a, 1)$ are indifferent. In this case $\rightarrow$ would also be a gradient at (a, b) , contradicting uniqueness. Therefore, $(a, 1) \succ (a, b)$, and hence $V_a \succ H_b$. The rankings of two horizontal lines and of two vertical lines follow from monotonicity and property (ii) of the relevant gradients. $\square$

7. A MODIFIED DEFINITION (Λ -TCD)

For the rest of the paper, we work with a variant of the notion of differentiable preferences. The motivation for this modification is that, in some cases, a preference relation may have *multiple* gradients, but among these gradients there is at least one that is tighter than all the others, in the sense that it has the smallest upper contour set. We call such a preference relation *tight convex differentiable*, denoted TCD, and refer to such a gradient as a *tight gradient*.

Formally, given a preference relation $\succsim$, a gradient $\triangleright \in \Lambda$ at x is a *tight gradient* of $\succsim$ at x if, for every other gradient $\triangleright' \in \Lambda$ at x , we have that $y \triangleright x$ implies $y \triangleright' x$. A Λ -*tight-convex-differentiable preference* (Λ -TCD preference) is a preference satisfying the requirement that, for every non-maximal x , there is a tight gradient. Obviously, every Λ -CD preference is Λ -TCD. Also, notice that any $\triangleright \in \Lambda$ is Λ -TCD, where at every point the tight gradient of $\triangleright$ is $\triangleright$ itself.

To demonstrate the concept, consider the preferences on $\mathbb{R}_+^2$ represented by $U(x_1, x_2) = x_1 + x_2$. These preferences are standard differentiable, but if Ψ is the set of algebraic linear orderings, then they are not Ψ -CD, since at the point $(0, 1)$ any ordering $\triangleright_{(v_1, v_2)}$, where $0 < v_2 \leq v_1$, is a gradient. However, the preferences are Ψ -TCD since $\triangleright_{(1, 1)}$ is a tight gradient at every point.

We now show that continuous Ψ -TCD preferences need not be representable by a standard differentiable function.

CLAIM 5. *A continuous, monotonic and convex preference relation $\succsim$ on $\mathbb{R}_+^2$ may be Ψ -TCD and yet not be representable by a function u that is standard differentiable and satisfies*

$$y \succ x \implies \nabla u(x) \cdot (y - x) > 0.$$

PROOF. Consider the following example, based on Dekel (1986) and Neilson (1991), of a preference relation on $X = \mathbb{R}_+^2$ that is convex, continuous, and Ψ -TCD. The example uses a function $f : [1, 2] \rightarrow [1, 2]$ constructed by Billingsley (1986, Example 31.1), which is continuous and strictly monotonic, satisfies $f(1) = 1$ and $f(2) = 2$, and has derivative 0 almost everywhere. Define $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by $g(z) = f(z)$ if $z \in [1, 2]$ and $g(z) = z$ otherwise. Let $\succsim$ be a monotonic preference relation on X for which $(z, 0) \sim (0, g(z))$ for every z , with an indifference line connecting any two such points; see Figure 3 for an illustration.

The preference relation is continuous and Ψ -TCD: At the origin, every gradient ranks every nonzero point in $\mathbb{R}_+^2$ above the origin. Hence, all gradients have the same upper contour set




FIGURE 3. A continuous Ψ -TCD preference relation need not be representable by a differentiable utility function with a gradient that points in the direction of improvement.

there and are tight. Every other point lies on some indifference line $(z, 0) - (0, g(z))$ and the tightest gradient is the normal to the indifference line. Suppose that $\succsim$ is representable by a differentiable function u , and denote by u_i the partial derivative with respect to dimension i . Since $u(z, 0) = u(0, g(z))$ for every z , it follows that $u_1(z, 0) = u_2(0, g(z))g'(z)$.

Thus, $u_1(z, 0) = 0$ for almost every $z \in [1, 2]$. Choose $z \in [1, 2)$ such that $g'(z) = 0$, and let $z' > z$. Then $(z', 0) \succ (z, 0)$, whereas $\nabla u(z, 0) \cdot (z' - z, 0) = u_1(z, 0)(z' - z) = 0$. Thus, the gradient of u does not point strictly in every direction of improvement. $\square$

8. APPLICATION: COMPETITIVE EQUILIBRIUM IN A BUNDLES-OF-ITEMS ECONOMY

In this section we apply our approach in order to shed light on an old problem: what, if any, are the competitive equilibrium outcomes of an exchange economy with a finite set of distinct and indivisible goods? Equilibrium existence results were achieved in the case in which agents maximize the utility of the bundle minus the cost of the bundle when the agents' utilities satisfy the "gross substitutes" condition in [Kelso and Crawford \(1982\)](#). Generalizations of this proposition admitting both substitutes and complements were proven in [Baldwin and Klemperer \(2019\)](#) and [Baldwin, Jagadeesan, Klemperer, and Teytelboym \(2023\)](#).

In a setting like ours, without utility for money, [Babaioff, Nisan, and Talgam-Cohen \(2021\)](#) establish a second welfare theorem for the case of two agents with linear ordinal preferences. We extend this result for any two-agent economy with a finite set of distinct and indivisible goods under the assumption that the agents' preference relations are strict and Ω -TCD (where Ω will be the set of strict linear orderings as below). Such preference relations need not be linear.

Consider a two-agent exchange economy where there is a finite set of distinct items Z and the space of bundles is $X = 2^Z$, that is, all subsets of Z . There are two agents, 1 and 2, and each of them has *strict* preferences over X , denoted $\succsim_1$ and $\succsim_2$. A pricing function $p : Z \rightarrow \mathbb{R}_{++}$ gives rise to a price for any bundle $A \in X$ by $p(A) = \sum_{a \in A} p(a)$. The ordering $\succeq_p$ on X is the "more expensive" relation according to p , that is, $A \succeq_p B$ if $p(A) \geq p(B)$. Let Ω be the set of all strict $\triangleright_p$ (that is, there are no two equally priced bundles according to the price vector p). Since the orderings are strict and agents have strict preferences, for $\triangleright_p$ to be a gradient at A , it is

TABLE 2. An economy with two Ω -convex preferences for which a Pareto-efficient allocation (in green) is not an equilibrium outcome.

$\succsim_1$	$\succsim_2$
ac	ad
bd	bc
ab	cd
ad, bc, cd	ab, ac, bd

necessary and sufficient that $B \succ A$ implies $p(B) > p(A)$. That is, condition (i) of the gradient implies condition (ii) of the gradient.

An *equilibrium* of this economy is a partition (A_1, A_2) of Z and a price vector p such that for each agent i and each bundle B , if $B \succ_i A_i$, then $p(B) > p(A_i)$. That is, any strictly preferred bundle is unaffordable.

Richter and Rubinstein (2015) showed that the second welfare theorem can fail for Ω -convex preferences (which were shown to be precisely the weakly monotonic preferences). That is, there can be Pareto-efficient allocations which are not supported as a linear price equilibrium. Table 2 shows such an example. The table reports only the rankings among doubletons. Complete each preference on X by ranking every bundle of larger cardinality strictly above every bundle of smaller cardinality. Within the singleton and triple classes, complete the rankings arbitrarily.

The allocation (ab, cd) is Pareto efficient, but it is not an equilibrium outcome: For p to be linear equilibrium prices they must satisfy $p(a) > p(c)$ (because $ad \succ_2 cd$), $p(c) > p(b)$ (because $ac \succ_1 ab$), and $p(b) > p(d) > p(a)$ for similar reasons. But that is an impossible cycle.

The set of Ω -TCD preferences includes Ω , but also much more. For example, consider the preferences: $abc \succ ab \succ ac \succ bc \succ c \succ b \succ a \succ \emptyset$. These preferences are not in Ω because $c \succ b$ and yet $ab \succ ac$. In this example, b is more of a complement to a than c is. To verify that the preferences are Ω -TCD: For any doubleton, the pricing function $p(a)=6$, $p(b)=5$, $p(c)=4$ is a tight gradient. For any singleton, the pricing function $p(a)=4$, $p(b)=5$, $p(c)=6$ is a tight gradient. Every gradient is tight at $\emptyset$.

We will now show that for two-agent economies with Ω -TCD preferences and indivisible distinct goods, the second welfare theorem does hold.

CLAIM 6. *Let $\succ_1, \succ_2$ be strict Ω -TCD preferences. Then, any Pareto-efficient allocation is a price equilibrium outcome.*

PROOF. Given a Pareto-efficient partition (A_1, A_2) , for each i , denote a tight gradient of $\succ_i$ at A_i by $\triangleright_{p_i}$. We will show that each $\triangleright_{p_i}$ constitutes an equilibrium ordering. We call V a *violating set for agent i* if $V \succ_i A_i$ and $p_j(A_i) > p_j(V)$ (for $j \neq i$). If a violating set exists, then p_j is not an equilibrium price ordering because agent i has a set V which he prefers to A_i and which is cheaper by p_j . If no violating set for agent i exists, then $\triangleright_{p_j}$ is an equilibrium ordering (agent j does not prefer any strictly p_j -lower bundle because $\triangleright_{p_j}$ is a gradient of $\succ_j$).

Denote by $SD(A, B)$ the size of the symmetric difference between A and B . There is no violating V for any i with $SD(A_i, V) = 1$. If $V = A_i \setminus \{b\}$ for some b , then it is impossible that $V \succ_i A_i$, and if $V = A_i \cup \{b\}$ for some b , then it is impossible that $p_j(A_i) > p_j(V)$. It remains

to be shown that for any violation V for agent i with $SD(A_i, V) > 1$, there is a violation V' for agent j with $SD(A_j, V) > SD(A_j, V')$.

Suppose that V is a violation (without loss of generality) for agent 1, that is, $V \succ_1 A_1$ and $p_2(V) < p_2(A_1)$. Since $V \succ_1 A_1$, Pareto optimality implies that $A_2 \succ_2 V^c$. Since $p_2(V) < p_2(A_1)$ and $Z = A_1 \cup A_2 = V \cup V^c$, it holds that $p_2(V^c) > p_2(A_2)$.

To construct a violation with smaller symmetric difference, we need our old friend (Farkas' lemma). To utilize it, we represent any bundle S with the indicator vector 1_S .

As mentioned earlier, property (ii) of the gradient condition is automatically implied by property (i) because preferences are strict. So, an ordering $\triangleright_q$ is a gradient of $\succ_2$ at A_2 if and only if $q \cdot 1_S > q \cdot 1_{A_2}$ for every $S \in U_{\succ_2}(A_2)$.

Since the preferences are Ω -TCD and $\triangleright_{p_2}$ is a tight gradient at A_2 , there is no gradient $\triangleright_q$ at A_2 with $q(A_2) > q(V^c)$. Then, there does not exist a positive vector q such that $q \cdot 1_S > q \cdot 1_{A_2}$ for all $S \in U_{\succ_2}(A_2)$ and $q \cdot 1_{A_2} > q \cdot 1_{V^c}$ because then $q + \varepsilon p_2$ would be a gradient which places V^c below A_2 , contradicting p_2 being tight.³ Let Γ be a matrix that has as its columns $1_S - 1_{A_2}$ for every $S \in U_{\succ_2}(A_2)$ as well as one standard unit vector for each good. By the foregoing, there exists no vector q for which each entry of the matrix product $q\Gamma$ is strictly positive, yet $q \cdot (1_{V^c} - 1_{A_2}) < 0$. By a perturbation argument, nor does there exist any vector q for which each entry of $q\Gamma$ is *weakly* positive, yet $q \cdot (1_{V^c} - 1_{A_2}) < 0$.⁴ Farkas' lemma then implies the existence of some non-negative vector α such that $\Gamma\alpha = 1_{V^c} - 1_{A_2}$. Given the structure of the matrix Γ , the entries of the vector α that correspond to the unit vectors in Γ can be considered as non-negative “dummy” variables, so that the following holds:

$$1_{V^c} - 1_{A_2} \geq \sum_{W \in U_{\succ_2}(A_2)} \alpha_W (1_W - 1_{A_2}) \quad (*)$$

for some set of non-negative coefficients (α_W) . Let $U^+(A_2)$ consist of all $W \in U_{\succ_2}(A_2)$ for which $\alpha_W > 0$.

Let $W \in U^+(A_2)$. If $a \in V^c \cap A_2$, then a can be added to W and $(*)$ continues to hold. Thus, without loss of generality $A_2 \setminus W \subseteq A_2 \setminus V^c$. Next, if $s \in W \setminus A_2$, then $s \in V^c \setminus A_2$. Therefore, $W \setminus A_2 \subseteq V^c \setminus A_2$. Thus, $SD(W, A_2) = |A_2 \setminus W| + |W \setminus A_2| \leq |A_2 \setminus V^c| + |V^c \setminus A_2| = SD(V^c, A_2)$. Finally, since $W \succ_2 A_2 \succ_2 V^c$, it must be that $W \neq V^c$, and thus $SD(W, A_2) < SD(V^c, A_2) = SD(V, A_1)$.

So, all $W \in U^+(A_2)$ have a strictly smaller symmetric difference than V^c and $W \succ_2 A_2$. To see that at least one of them constitutes a violation, multiply inequality $(*)$ by p_1 . Notice that $p_1 \cdot (1_{V^c} - 1_{A_2}) = p_1 \cdot (1_{A_1} - 1_V) < 0$ because $V \succ_1 A_1$ and $\triangleright_{p_1}$ is the gradient of $\succ_1$ at A_1 . Therefore, $p_1 \cdot (\sum_{W \in U^+(A_2)} \alpha_W (1_W - 1_{A_2})) < 0$ and so there is at least one $W \in U^+(A_2)$ such that $p_1(W) < p_1(A_2)$. Thus, W is a violation for agent 2 with a smaller symmetric difference than V . $\square$

REMARK. Claim 6 proves the stronger result that, in a two-agent economy with strict, monotonic preferences, a Pareto-efficient allocation is supportable whenever each agent's preference has a tight Ω -gradient at that agent's allocated bundle; tightness need not hold elsewhere.

³The issue is that the vector q may not be a gradient because it prices two bundles the same, but then $q + \varepsilon p_2$ for sufficiently small ε prices every bundle differently and maintains the above strict inequalities, contradicting p_2 being a tight gradient at A_2 .

⁴The perturbation is the same as before, $q + \varepsilon p_2$ for sufficiently small ε .

TABLE 3. Failure of the second welfare theorem in a three-agent economy.

	1	2	3
a	16	8	ε
b	1	1	2
c	2	16	16
d	ε	2	3
e	8	ε	8
f	12	12	4

Note: A Pareto-efficient allocation (in green) that cannot be supported by an equilibrium price vector, even though preferences are linear and strict.

REMARK. The following example shows that the second welfare theorem is no longer guaranteed to hold in bundles-of-items economies with more than two agents, even if agents have linear preferences.

Consider an economy with three agents and six goods. Agent 1 owns the bundle ab , agent 2 owns the bundle cd , and agent 3 owns the bundle ef . The agents' valuations of the individual items are as indicated in Table 3. Fix $0 < \varepsilon < 1$. Each agent's valuations of bundles are linearly additive and the induced preferences are strict.

The allocation is Pareto optimal: In any Pareto improvement, agents 1 and 2 would both have to consume more than one good. If agent 3 got c , then agent 2 has to at least take both a and f , but then the most agent 1 would be left with is bde , leaving agent 1 worse off. Without getting c , agent 3 would need two goods and those must be ef . But, then ab for agent 1 and cd for agent 2 maximizes their total utility, and so no Pareto improvement among them is possible.

However, no linear equilibrium price vector exists: If p were to be one, it would have to satisfy $p(e) + p(f) > p(a) + p(b)$ (since $ef \succ_1 ab$), $p(a) + p(f) > p(c) + p(d)$ (since $af \succ_2 cd$), $p(c) > p(e) + p(f)$ (since $c \succ_3 ef$), and $p(b) + p(d) + p(e) > p(e) + p(f)$ (since $bde \succ_3 ef$). But adding these four inequalities leads to an impossibility.

REMARK. Finally, a natural question arises: what happens when nonlinear prices are permitted? In this case, a larger set of allocations can be sustained as equilibria. [Richter and Rubinstein \(2015\)](#) showed that the second welfare theorem holds with nonlinear prices, but that the first welfare theorem can fail. Thus, nonlinear prices can support all Pareto-efficient allocations, but may also support Pareto-inefficient ones.

REFERENCES

- Babaioff, Moshe, Noam Nisan, and Inbal Talgam-Cohen (2021), "Competitive Equilibrium with Indivisible Goods and Generic Budgets." *Mathematics of Operations Research*, 46(1), 382–403. [12]
- Baldwin, Elizabeth, Ravi Jagadeesan, Paul Klemperer, and Alexander Teytelboym (2023), "The Equilibrium Existence Duality." *Journal of Political Economy*, 131(6), 1440–1476. [12]
- Baldwin, Elizabeth, and Paul Klemperer (2019), "Understanding Preferences: 'Demand Types', and the Existence of Equilibrium with Indivisibilities." *Econometrica*, 87(3), 867–932. [12]

Billingsley, Patrick (1986), *Probability and Measure*, Second Edition. John Wiley & Sons, New York, NY. [11]

Dekel, Eddie (1986), “An Axiomatic Characterization of Preferences under Uncertainty: Weakening the Independence Axiom.” *Journal of Economic Theory*, 40(2), 304–318. [11]

Diasakos, Theodoros M., and Georgios Gerasimou (2022), “Preference Conditions for Invertible Demand Functions.” *American Economic Journal: Microeconomics*, 14(2), 113–138. [6]

Kelso, Jr., Alexander S., and Vincent P. Crawford (1982), “Job Matching, Coalition Formation, and Gross Substitutes.” *Econometrica*, 50(6), 1483–1504. [12]

Neilson, William S. (1991), “Smooth Indifference Sets.” *Journal of Mathematical Economics*, 20(2), 181–197. [11]

Renou, Ludovic, and Karl H. Schlag (2014), “Ordients: Optimization and Comparative Statics without Utility Functions.” *Journal of Economic Theory*, 154, 612–632. [6]

Richter, Michael, and Ariel Rubinstein (2015), “Back to Fundamentals: Equilibrium in Abstract Economies.” *The American Economic Review*, 105(8), 2570–2594. [1, 3, 13, 15]

Richter, Michael, and Ariel Rubinstein (2019), “Convex Preferences: A New Definition.” *Theoretical Economics*, 14(4), 1169–1183. [1, 2, 4, 7]

Richter, Michael, and Ariel Rubinstein (2024), *No Prices No Games!: Four Economic Models*, Second Edition. Open Book Publishers, Cambridge, U.K. [1]

Rubinstein, Ariel (2006), *Lecture Notes in Microeconomic Theory: The Economic Agent*, Second Edition (2012). Princeton University Press, Princeton, NJ. [2, 6]

Shapley, Lloyd, and Herbert Scarf (1974), “On Cores and Indivisibility.” *Journal of Mathematical Economics*, 1(1), 23–37. [2]

Co-editor Mira Frick handled this manuscript.