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

February 2, 2026

ABSTRACT: A new definition of differentiable preferences is proposed for spaces without requiring any algebraic structure. Given a set of primitive orderings Λ, preferences are Λ-differentiable if at each alternative there is a unique ordering in Λ whose advancement is necessary for improvement of the preferences and whose reduction is sufficient for deterioration. As an application, it is shown that the 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

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

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 provides the motivation for 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, but also apply to settings where the underlying grand set is not Euclidean and without any reliance on utility representations.

The approach we use originated in Richter and Rubinstein (2019), and was refined in Richter and Rubinstein (2024). The grand set of alternatives X is enriched with a collection Λ of primitive orderings on X, each representing a relevant consideration. We 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*, whose advancement is necessary for improvement.

As an example, we considered a reviewer of candidates who vary in their 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, which has the property that for a candidate to be superior to x, it is necessary for the candidate to be a better researcher. However, research does not have to be the only such criterion; there could be multiple orderings whose advancement is necessary for improvement.

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." The definition is related to the standard requirement that preferences be represented by a differentiable and quasi-concave utility function with non-zero gradient at any non-maximal point. Rubinstein (2012, pages 55–56) suggested an alternative definition of differentiable preferences on Euclidean spaces, which is formulated directly on the preferences and does not mention the word "utility." That definition was not equivalent to the standard one but captured much of what differentiable preferences do (see also Renou and Schlag, 2014).

We now propose a definition of Λ-differentiable convex preferences without any reliance on algebraic structure. The concept requires that for each alternative which is not maximal, there is a *unique* primitive ordering in Λ which 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-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-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 indivisible and distinct goods. It is well known that in such economies, the second welfare theorem can fail (Shapley and Scarf, 1974). We let Λ consist of all linear pricing orderings and show that for the two-agent exchange economy, if both agents' preferences satisfy a variant of Λ-differentiability, then the second welfare theorem holds (namely, any Pareto-optimal allocation is also a competitive equilibrium outcome with linear prices).

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 $\trianglerighteq \in \Lambda$.

Given a preference relation $\succsim$ on X we say that $\trianglerighteq \in \Lambda$ 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 $\succsim$ and $\trianglerighteq$, respectively). The relation $\succsim$ is Λ-*convex* if for every $x \in X$ there exists at least one critical ordering in Λ. Note that if $\trianglerighteq$ is critical at x, then any alternative that is weakly $\trianglerighteq$-lower must be weakly $\succsim$-dispreferred.

For our definition of Λ-differentiable preferences, we refine the critical ordering notion. An ordering $\trianglerighteq \in \Lambda$ is a *gradient of* $\succsim$ *at* x if
(i) if $y \succ x$ then $y \triangleright x$; and
(ii) if $x \triangleright y$ then $x \succ y$.

A preference relation $\succsim$ is Λ-*convex-differentiable* (Λ-CD) if for any non-maximal element $x \in X$ there is a unique gradient $\nabla_{\succsim}(x)$ in Λ. Denote by $D(\Lambda)$ the set of Λ-CD preference relations. For any preference relation $\succsim$ and $x \in X$, 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 $\succsim$. That is, a preference $\succsim$ is Λ-CD if for each non-$\succsim$-maximal x there is a unique $\trianglerighteq \in \Lambda$ such that $U_{\triangleright}(x) \supseteq U_{\succ}(x)$ and $L_{\triangleright}(x) \subseteq L_{\succ}(x)$. (Note that in the standard setting, the gradient at a maximal point fails to indicate a direction of improvement because no such direction exists. Likewise, in our setting, at a maximal alternative, there may be no unique gradient.)

3. Properties of Λ-CD Preferences

(1) 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).

(2) 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 $\trianglerighteq, \trianglerighteq' \in \Lambda$ such that $U_{\triangleright}(x) \supseteq U_{\triangleright'}(x) \neq \emptyset$. This condition is satisfied by the set of algebraic primitive orderings on $\mathbb{R}^K$.

(3) 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 $a \trianglerighteq b$ for all $\trianglerighteq \in \Lambda$, 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 $a \triangleright b$ for all $\trianglerighteq \in \Lambda$ 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 $a \succ b$.

(4) 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.)

(5) $D(\Lambda_1)$ 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.

4. 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 will 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, i.e., $\Psi = \{\trianglerighteq_v \mid v \in \mathbb{R}^K \setminus \{0\}\}$ where $\trianglerighteq_v$ is the ordering represented by the function $v \cdot x$. If the preference relation $\succsim$ is monotonic, convex and represented by a differentiable utility function u with a non-zero gradient at all x, then $\succsim$ is Ψ-differentiable and the gradient ordering is $\trianglerighteq_{\nabla u(x)}$.

As mentioned earlier, while the set of convex, continuous, and standard differentiable preference relations is subsumed within our definition, the spirit of the definition is not identical. The standard definition also requires that $x + \epsilon d \succ x$ for any d satisfying $d \cdot \nabla u(x) > 0$ and small enough $\epsilon > 0$, a requirement that has no analogy in our definition.

While standard differentiable preferences need to be continuous, that is not the case for 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.

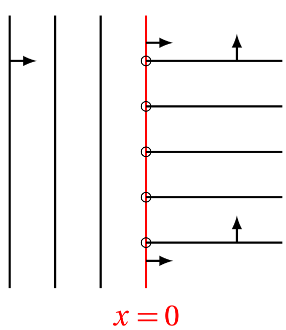

$$U(x,y) = \begin{cases} x & \text{if } x \le 0 \\ e^y & \text{if } x > 0 \end{cases}$$

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

5. Differentiability and Optimality

Given any primitive ordering $\trianglerighteq$ and an alternative e, define $B(e,\trianglerighteq) = \{x \in X | e \trianglerighteq x\}$ to be a *budget set*. We refer to $\trianglerighteq$ as the *expenditure ordering* and $x \trianglerighteq y$ means that x is (weakly) "more expensive" than y. We make the following analogy to a standard "budget frontier" condition: given any expenditure ordering $\trianglerighteq$, any $\succsim$-maximal alternative from any budget set $B(e,\trianglerighteq)$ is $\trianglerighteq$-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 $B(e,\trianglerighteq)$ if and only if $\nabla_{\succsim}(e) = \trianglerighteq$.*

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

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

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 $B(e,\trianglerighteq)$ but not $\succsim$-maximal in X. Define $U_{\succsim}(x^*) = \{y \in X | y \succsim x^*\}$ to be the weak upper contour of x^* and suppose that x^{**} is $\trianglerighteq$-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) $\trianglerighteq$-indifferent.*

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

6. Examples

Example 1: Horizontal and vertical primitives on a finite grid. Let $X = \{0,1,2,3,4\} \times \{0,1,2,3,4\}$ 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: for each non-maximal alternative, there is a unique dimension in which moving upwards is a strict improvement while moving upwards in the other dimension preserves indifference.

4	$1\rightarrow$	$4\rightarrow$	$5\rightarrow$	$6\rightarrow$	8
3	$1\rightarrow$	$4\rightarrow$	$5\rightarrow$	$6\rightarrow$	$7\uparrow$
2	$1\rightarrow$	$3\uparrow$	$3\uparrow$	$3\uparrow$	$3\uparrow$
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**

TABLE 1: Illustration of differentiability in a multidimensional discrete context.

To understand this, consider the point $(0,0)$. By Λ-convexity, it cannot be that both $(a,0) \succ (0,0)$ and $(0,b) \succ (0,0)$ because then there would be no critical ordering at $(0,0)$. By the weak Pareto property, the preferences are weakly monotonic. Accordingly, without loss of generality, $(0,0) \sim (a,0)$ for every $a > 0$ and therefore $\uparrow$ is the gradient at $(0,0)$. 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, by monotonicity, $(0,b) \succ (0,0)$ for all $b > 0$. The result then follows by induction with the bottom row removed.

To summarize, there is an "increasing" and "continuous" path from $(0,0)$ to $(4,4)$ such that at every point on it, one direction preserves indifference and the other is strict improvement.

The Λ-CD preferences have a natural meaning when the two orderings reflect the preferences of two agents, 1 and 2. They are constructed sequentially. First one agent, say i_1 is selected and all i_1-worst alternatives are assigned to be indifferent and are placed at the bottom of the ranking. Then, agent i_2 (who might be i_1) is selected and all i_2-worst alternatives from the as-yet-unranked alternatives are assigned to be indifferent, above those already ranked and below all other alternatives. And so on, until all alternatives are ranked. In other words, the preferences can be thought of as an 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 put above all previously ranked alternatives and below the rest of the alternatives. This analysis and procedure would equally apply to other grand sets, such as $X = \{(x, y) | x \in \{0, \dots, a\}, y \in \{0, \dots, a\}, x + y \leq a\}$ where there is no uniformly top-ranked alternative.

Example 2: Horizontal and vertical primitives in the plane. Let $X = [0,1]^2$ and Λ consist of two primitives: the horizontal ($a \rightarrow b$ if $a_1 \geq b_1$) and the vertical ($a \uparrow b$ if $a_2 \geq b_2$) on X. Claim 3 characterizes all Λ-CD preferences. It implies that the only continuous monotonic Λ-CD preferences are the two primitive orderings themselves.

Claim 3. *Let $\succsim$ be a monotonic preference relation on X and let M denote the set of $\succsim$-maximal points. Then $\succsim$ is Λ-CD if and only if $X \setminus M$ can be partitioned into two sets U and R such that:*
(i) if $x \in U$, then any point weakly to the southeast of x is also in U; and thus if $x \in R$, then any point weakly to the northwest of x is also in R;
(ii) any point $(a, b) \in U$ is on a horizontal indifference line to $(1, b)$ denoted H_b, and any point $(a, b) \in R$ is on a vertical indifference line to $(a, 1)$ denoted V_a; and
(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, look at the point (a, b): if it is in U, then $V_a \succ H_b$ and if it is in R, then $H_b \succ V_a$.

Figure 2 illustrates such a preference relation. The boundary between the regions U and R is indicated in blue and the preferences are strictly increasing on this path. The partition U, R indicates the gradient (if $a \in U$ the gradient is $\uparrow$, otherwise it is $\rightarrow$), which is displayed as arrows in the figure. 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 (in the figure, the red dashed indifference lines at a point are preferred to the gray).

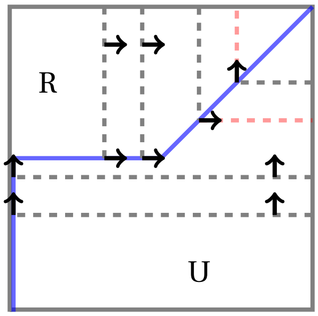

FIGURE 2: Λ-CD preferences in Example 2.

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 $R = \{x \in X \setminus M \,|\, \nabla_{\succsim}(x) = \rightarrow\}$ and $U = \{x \in X \setminus M \,|\, \nabla_{\succsim}(x) = \uparrow\}$. 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 to the southeast of x must also be 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 $(a, 1) \in R$. Without loss of generality, assume that $(a, b) \in U$. By monotonicity, it must be that $(a, 1) \succsim (a, b)$. 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. □

Example 3: Differentiability with single-peaked orderings. Let $X = \mathbb{R}$. A preference relation $\trianglerighteq$ is *single peaked* if there exists a $peak \in \mathbb{R}$ such that $a < b < peak < c < d$ implies $a \lhd b \lhd peak \rhd c \rhd d$. We take Λ to be a set of single-peaked orderings. In the political economy literature, single-peaked preferences are used to model a voter's preferences over social policies with $peak$ representing the voter's ideal policy. The Λ-CD property can be thought of as a requirement on the aggregation of different preferences into a social preference.

We restrict the members of Λ to be *balanced* in the sense that for every $x \neq peak$ there is another unique point $m(x)$ on the other side of the peak such that x and $m(x)$ are indifferent. Define $m(peak) = peak$.

We say that $\trianglerighteq_1, \trianglerighteq_2, \ldots, \trianglerighteq_K$ are *systematically ordered* if for every x we have $m_1(x) < m_2(x) < \cdots < m_K(x)$ and in particular $peak_1 < peak_2 < \cdots < peak_K$.

In the following claim, attention is restricted to preferences that are continuous and have at most one maximal point (in order to rule out trivialities such as preferences with total indifference).

Claim 4. *Let $\Lambda = \{\trianglerighteq_1, \trianglerighteq_2, \ldots, \trianglerighteq_K\}$ be a finite set of single-peaked orderings on $\mathbb{R}$ which are balanced and systematically ordered.*
(a) Assume $K = 2$. If $\succsim$ is a continuous preference relation on $\mathbb{R}$ that has at most one maximal point, then $\succsim$ is Λ-CD if and only if
(i) $\succsim$ is single peaked with $peak \in [peak_1, peak_2]$; and
(ii) $\succsim$ is balanced; if $x < peak$, then $m_1(x) < m(x) \leq m_2(x)$; and if $x > peak$, then $m_1(x) \leq m(x) < m_2(x)$.
(b) If $K \geq 3$, then no continuous Λ-CD preference with a unique maximal point exists.

Proof. (a) Suppose that $\succsim$ satisfies conditions (i) and (ii) and denote its unique maximum point by $peak$. Without loss of generality, it is sufficient to deal with $x < peak$. By condition (ii), $m(x) \leq m_2(x)$, and thus $U_{\triangleright_2}(x) = (x, m_2(x)) \supseteq (x, m(x)) = U_{\succ}(x)$ and $L_{\triangleright_2}(x) = (-\infty, x) \cup (m_2(x), \infty) \subseteq (-\infty, x) \cup (m(x), \infty) = L_{\succ}(x)$, so $\trianglerighteq_2$ is a gradient at x. To see that $\trianglerighteq_1$ is not a gradient at x, note that for $x \in [peak_1, peak)$, $peak \succ x$ but $x \triangleright_1 peak$. For $x < peak_1$, we have $m_1(x) < m(x)$ by condition (ii), and so for any point y such that $m_1(x) < y < m(x)$ it holds that $y \succ x$ and $x \triangleright_1 y$.

Now suppose that $\succsim$ is Λ-CD, continuous, and has at most one peak. It must have a maximum point (and is thus unique) since by the strict Pareto property the preferences are strictly increasing to the left of $peak_1$ and are strictly decreasing to the right of $peak_2$ and by the continuity of the preferences they have a maximum in the closed interval $[peak_1, peak_2]$. Denote the maximal point by $peak$.

For every $x < peak$ the gradient must be $\trianglerighteq_2$. This is obvious for $x \in [peak_1, peak)$ while for $x < peak_1$, it holds that $U_{\triangleright_1}(x) \subset U_{\triangleright_2}(x)$ and $L_{\triangleright_1}(x) \supset L_{\triangleright_2}(x)$ and thus by the uniqueness of the gradient, it must be $\trianglerighteq_2$. Therefore, for every $x < peak$ the preferences are increasing. Analogously, the preferences are decreasing for $x > peak$ and thus, the preferences are single peaked. Furthermore, $U_{\succ}(x) \subseteq (x, m_2(x))$ for every $x < peak$ and by the continuity of the preferences $m(x)$ exists and $m(x) \leq m_2(x)$. In order for $\trianglerighteq_1$ not to be a gradient at x, it must be that $m(x) > m_1(x)$. An analogous argument applies to every point to the right of $peak$.

(b) Suppose that $\Lambda = \{\trianglerighteq_1, \trianglerighteq_2, \ldots, \trianglerighteq_K\}$ are systematically ordered single-peaked orderings. We now show that there are no Λ-CD preferences. Let x, y be ordered $x < peak_1 < peak_K < y$ such that $x \sim_2 y$. (Such a pair exists since take $x < peak_1$. If $m_2(x) \leq peak_K$, then choose y to the right of $peak_K$ and its mirror must be to the left of $peak_1$.)

Suppose that $x \succsim y$ and x is not $\succsim$-maximal (since there is at most one $\succsim$-maximal alternative, if x is $\succsim$-maximal, start with a different pair $x \sim_2 y$ to begin with). By the strict Pareto property, every point to the left of x is strictly $\succsim$-inferior to x and every point to the right of y is strictly $\succsim$-inferior to y. Thus, $U_{\succ}(x) \subseteq (x, y) = U_{\triangleright_2}(x)$ and $L_{\succ}(x) \supseteq (-\infty, x) \cup (y, \infty) = L_{\triangleright_2}(x)$. Thus, $\trianglerighteq_2$ is a gradient at x. Since the primitives are systematically ordered, it holds that $U_{\triangleright_2}(x) \subseteq U_{\triangleright_3}(x)$ and $L_{\triangleright_2}(x) \supseteq L_{\triangleright_3}(x)$, and thus $\trianglerighteq_3$ is also a gradient at x. Similarly, if $x \precsim y$ then both $\trianglerighteq_2$ and $\trianglerighteq_1$ are gradients at y, a contradiction. □

7. A Modified Definition (Λ-CMD)

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 one which is tighter than all the others in the sense that it has a smallest upper contour set. We will call such a preference relation *convex minimal differentiable*, denoted CMD, and will refer to such a gradient as a *minimal gradient.*

Formally, given a preference relation $\succsim$, a gradient $\trianglerighteq \in \Lambda$ at x is a *minimal gradient* of $\succsim$ at x if for every gradient $\trianglerighteq'$ at x we have that $y \trianglerighteq x$ implies $y \trianglerighteq' x$. A Λ-*convex-minimal-differentiable preference* (Λ-CMD preference) is a preference satisfying that for any non-maximal x there is a minimal gradient. Obviously, every Λ-CD preference is Λ-CMD. Also, notice that any $\trianglerighteq \in \Lambda$ is Λ-CMD where at any point, the minimal gradient of $\trianglerighteq$ is $\trianglerighteq$ 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, these preferences are not Ψ-CD since at the point $(0,1)$ any ordering $\geq_{(x,y)}$ where $0 < y \leq x$ is a gradient. However, the preferences are Ψ-CMD since $\geq_{(1,1)}$ is a minimal gradient at any point.

We will now see that continuous Ψ-CMD preferences might not be representable by a standard differentiable function.

Claim 5. *A continuous preference $\succsim$ on $\mathbb{R}^2_+$ may be Ψ-CMD and yet not representable by a differentiable u which has the property that $b \succ a$ implies $\nabla_u(a)\cdot(b-a) > 0$.*

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

The preferences are continuous and Ψ-CMD. Suppose $\succsim$ is representable by a differentiable u and denote by u_i the partial derivative according to dimension i. Since $u(x,0) = u(0,g(x))$, $u_1(x,0) = u_2(0,g(x))g'(x)$. Thus $u_1(x,0) = 0$ almost everywhere in $[1,2]$. But $(x',0) \succ (x,0)$ for all $x' > x$ and yet $\nabla_u(x,0)\cdot(x'-x,0) = 0$. That is, the gradient of u does not point in the direction of improvement. □

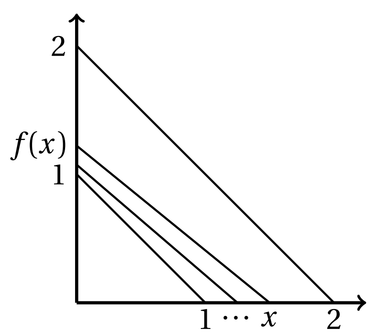

FIGURE 3: A continuous Ψ-CMD preference relation need not be representable by a differentiable utility function with a gradient which points in the direction of improvement.

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 will show the validity of the second welfare theorem for any two-agent economy with a finite set of distinct and indivisible goods under the assumption that the agents' preference relations are Ψ-minimal differentiable, which need not be linear.

Consider a two-agent exchange economy where there is a finite set of 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, that is, $\succ_1$ and $\succ_2$. A pricing function $p : Z \to \mathbb{R}_{++}$ gives rise to a price for any bundle $A \in X$ by $p(A) = \sum_{a \in A} p(a)$. The ordering $\trianglerighteq_p$ on X is the "more expensive" relation according to p, that is, $A \trianglerighteq_p B$ if $p(A) \geq p(B)$. Let Λ 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 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 Ψ-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.

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

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

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 Ψ-CMD preferences includes Ψ (all linear preferences), 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 Ψ 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 Ψ-CMD: For any doubleton, the pricing function $p(a) = 6$, $p(b) = 5$, $p(c) = 4$ is a minimal gradient. For any singleton, the pricing function $p(a) = 4$, $p(b) = 5$, $p(c) = 6$ is a minimal gradient. Every gradient is minimal at $\emptyset$.

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

Claim 6. *Let $\succ_1, \succ_2$ be strict Ψ-CMD 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 minimal gradient of $\succ_i$ at A_i by $\rhd_{p_i}$. We will show that each $\rhd_{p_i}$ 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 $\rhd_{p_j}$ is an equilibrium ordering (agent j does not prefer any strictly p_j-lower bundle because $\rhd_{p_j}$ 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 $SD(A_i, V) > SD(A_j, V')$.

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 $\rhd_q$ 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 Ψ-CMD and $\rhd_{p_2}$ is a minimal gradient at A_2, there is no gradient $\rhd_q$ at A_2 with $q(A_2) > q(V^c)$. That is, for every positive vector q which satisfies that $q \cdot 1_S > q \cdot 1_{A_2}$ for all $S \in U_{\succ_2}(A_2)$ it must hold that $q \cdot 1_{V^c} \geq q \cdot 1_{A_2}$. By Farkas' lemma the following holds:

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

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 the coefficients are positive.

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 $SD(W, A_2) < SD(V^c, A_2)$.

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 $\rhd_{p_1}$ is the gradient of $\succ_1$ at A_1. Therefore, $p_1 \cdot \left(\sum_{W \in U^+(A_2)} \alpha_W (1_W - 1_{A_2})\right) < 0$ 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. □

Remark: Claim 6 also proves a stronger result: in a two-agent economy, if both preferences are Ψ-differentiable at a Pareto-efficient allocation (but not necessarily elsewhere) and monotonic, the second welfare theorem holds for that allocation.

9. 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.

Baldwin, Elizabeth, Ravi Jagadeesan, Paul Klemperer, and Alexander Teytelboym (2023). "The Equilibrium Existence Duality." *Journal of Political Economy*, 131 (6), 1440–1476.

Baldwin, Elizabeth and Paul Klemperer (2019). "Understanding Preferences: 'Demand Types', and the Existence of Equilibrium with Indivisibilities." *Econometrica*, 87 (3), 867–932.

Billingsley, Patrick (1986). *Probability and Measure.* Second Edition, John Wiley & Sons.

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

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

Neilson, William S. (1991). "Smooth Indifference Sets." *Journal of Mathematical Economics*, 20 (2), 181–197.

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

Richter, Michael and Ariel Rubinstein (2015). "Back to Fundamentals: Equilibrium in Abstract Economies." *The American Economic Review*, 105 (8), 2570–2594.

Richter, Michael and Ariel Rubinstein (2019). "Convex Preferences: A New Definition." *Theoretical Economics*, 14 (4), 1169–1183.

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

Rubinstein, Ariel (2012). *Lecture Notes in Microeconomic Theory: The Economic Agent*. Second Edition, Princeton University Press.

Shapley, Lloyd and Herbert Scarf (1974). "On Cores and Indivisibility." *Journal of Mathematical Economics*, 1 (1), 23–37.