2013

Risk and Uncertainty Program

RSMG Working Paper Series

TITLE:

Ordinal, Nonlinear Context Dependence

AUTHORS:

Patrick O'Callaghan

Schools of Economics and Political Science

The University of Queensland

St Lucia

Brisbane

Australia 4072

Web: www.uq.edu.au

Working Paper: R13_4

Ordinal, nonlinear context-dependence*

Patrick O'Callaghan

University of Queensland

May 18, 2013

Abstract

Prospect theory [KT79] and its more recent formalizations [KR06, KR07] prescribe "nonlinear" reference-dependence. The same may be said of other forms of *context-dependence* such as status quo bias [MO05]. Even in settings where there is a strong case for linear context-dependence such as Gilboa and Schmeidler's theory of case-based decisions, nonlinearity is typical in the absence of "diversity of preference". Furthermore, the hope of providing an axiomatic foundation for neuroscientific models of decision making, where context is interpreted as a physical state or "connectome" suggest a general, ordinal axiomatization of nonlinear context-dependence is called for. As with traditional, "context-free" models of ordinal utility (eg. Debreu [Deb54]), the issue of continuity is central: precise, yet simple and intuitive, conditions on the set of contexts are needed if preferences have a representation that is continuous across contexts. The continuity condition I employ is the obvious choice and is a generalisation of [GS03a]. There are interesting connections with literature on jointly continuous utility [Lev83, CCM09]. Finally, a promising feature of the present approach is it that may be used to axiomatise payoffs associated with discontinuous games (such as Bertrand oligopoly).

*I thank, without implication, Simon French, Simon Grant, Peter Hammond, Jeff Kline, Robert MacKay, John Quiggin and Horst Zank for their helpful suggestions concerning earlier versions of this paper. The main result of this paper first appeared in the third chapter of my PhD thesis at the University of Warwick in 2011.

Contents

1 Introduction

In psychology a context effect is understood to be the influence of environmental factors on the perception of a stimulus. Prospect theory [KT79], where the reference point identifies the context, provides a leading example of a static model where preferences vary with context. As the recent formalisation of Kőszegy and Rabin [KR06, KR07] highlights, prospect theory also exemplifies the case where preferences may vary with context in a continuous but "*non-linear*" way. The first goal of the present paper is to axiomatise this form of *context-dependence*.

More generally, in decision theory context may also refer to data on socio-economic variables or 'choice situations' [McF04]; conjectures about other players in a game [GS03a]; memories/databases in case-based decision theory (CBDT) [GS95, GS01]; or the composition of the set of alternatives available to the decision maker itself [Bar10, PX12].[1]

Note that all these forms of context are in some sense encoded in the brain of a decision maker. As the field of neuroscience strives with increasing success to show, the set of connections and strength of connections in the brain (recently coined the connectome) are physical representations of the information agents perceive [Seu12].[2] Yet there may also be influential factors that may have no counterpart in the sphere of consciousness of the decision maker, which the decision maker has no language to describe, but since they have an impact on the brain's output (decisions), they too must be encoded in the brain.

This latter point suggests that the traditional, state-dependent utility (see [DR04] for a recent survey), where the primitive is a single, context-free preference relation will not do: it seems absurd to include contexts of which the decision-maker may be entirely unaware. By contrast, there may be good reason to think that the decision-maker will be able to decide which of a fixed set of alternatives she prefers once the context is fixed (regardless of whether the decision-maker is aware of precisely which context she is in). In this case, a decision-maker is best characterised in terms of a family of preference relations as opposed to a single one. Such a decision-maker is said to exhibit context-dependence and a family of preference relations that is indexed by

[1]It is thanks to a lively debate between Ken Binmore and Prasanta Pattanaik at a talk by the latter at CRETA/Marie Curie Conference in honour of Peter Hammond that I became interested in context-dependence. There, Prof. Pattanaik made a convincing case in favour of the need for an axiomatic foundation of context-dependence in decision theory.

[2]Seung [Seu12] provides a very accessible introduction to this promising area of research.

contexts is referred to as *context preferences*.

Similarly, in prospect theory, the idea of extending preferences to alternative-context pairs (or functions from contexts to consequences) seems to defeat the original goal of focusing preferences on changes in wealth from a given reference level, that Kahnemann and Tversky so convincingly promoted. Indeed, Bleichrodt [Ble07] describes this principle in terms of incompleteness of the preference ordering over pairs of prospects and reference-points. Likewise, although [KR07] models reference dependence in terms of a single real-valued function over pairs of wealth and reference-wealth levels, it would be unwise to interpret this as meaning that the preferences that underly their utility function are complete over such pairs. This is because a casual inspection reveals that their utility function is actually strictly decreasing in the reference-wealth level. Surely the ranking of reference-wealth levels implied by such a utility function is arbitrary.

The present approach shares this property in that statements such as such as "I prefer context x to context y.", whilst not necessarily devoid of meaning, are excluded from the utility representation. A given utility representation implies a ranking of contexts for every fixed alternative, as well as a ranking of the set of alternative-context pairs, but it is understood to be arbitrary in the sense that there may well exist distinct representations of the same preferences that give rise to other orderings of alternative-context pairs. The only preference statements that are represented by the utility function are of the form: "*Given context x, I prefer alternative a to alternative b.*"; or in the case where the decision maker is unaware of the context relative to the observer/analyst, simply "*At present I prefer a to b.*"

In light of the present, broad interpretation of the term "context" it seems natural to ask: when is the space of contexts inappropriate for the purposes of being able to precisely characterise the preferences of a decision maker using a real-valued function? For instance, does a particular graph of all neurons and connections between them, present in the brain at decision time, constitute a plausible form of context? The answer will of course depend on what is required of the representation.

The present focus is on nonlinear context-dependence, and whilst this will be precisely defined below, the following quote serves to identify a basic property that a representation of context-dependence should preserve.

> When processing sensory input, it is of vital importance for the neural systems to be able to discriminate a novel stimulus from the background of redundant, unimportant signals. [MMB$^+$12]

This may be translated to mean that small perturbations in the context does not change strict preference over alternatives. It is clearly analogous to the concept of "*continuity of preference across contexts*"–at least in some regions of the space of contexts.

Each of the aforementioned models appeal to this kind of continuity. This common ground demands a common framework, and a foundation with results that parallel the classical (context-free) theorems of Wold, Eilenberg and Debreu (respectively [Wol44, Eil41, Deb54]) is required.

Beginning with Levin [Lev83], the literature on "*jointly continuous*" utility representations, [BM95, Meh98, CCM09], is the nearest thing to such a foundation. The reason it is not is due to the primary application being existence proofs in general equilibrium theory. That is, the approach imposes conditions that directly give rise to joint continuity across alternative-context pairs of the utility representation. These conditions are difficult to motivate in the decision theory setting, and are substantially stronger than the continuity across contexts informally motivated above.

1.1 The present contribution

The purpose of the present paper is to provide an axiomatic model for representing context preferences over a set of alternatives A where "non-linear" dependence on contexts x in X is plausible.

A representation of context preferences refers to a function of the form $U : A \times X \to \mathcal{R}$, such that for any context x an alternative a is strictly preferred to alternative b if and only if

$$U(a, x) > U(b, x).$$

That is, at each context x, there exists a utility function that represents a preferences in the classical sense.

Representations of linear context-dependence are also of this form. Gilboa and Schmeidler's model of context dependence, applied to the setting of case-based decision theory (CBDT) [GS95, GS01, GS03b] and choice under uncertainty [GS03a], is *linear* in the sense that the preferences they study may be represented using an operator that is linear on the space where contexts are embedded. As the following example highlights, the actual operation over the space of contexts that is preserved by the representation is not linearity but additivity across vectors of natural numbers for the case of CBDT.

Example 1.1. *In CBDT,* cases*, are represented by a dimension* $1, \ldots, n$*, and the context* x *is understood to be a database (or memory) of cases. Thus the context space* $X := \mathbb{N}^n$ *is the set of possible* databases*. A database* x *is a vector in* $\mathbb{N}^n$ *with* k^{th} *entry equal to the frequency of case* k *in* x *. Each* *alternative a gives rise to a vector $v(a)$ in $\mathcal{R}^n$ such that for each x in X*

$$\mathcal{U}(a, x) = v(a) \cdot x\,.$$

Formally, context preferences exhibit *linear context-dependence* when both of the following are true:

1 X can be embedded in a linear space Y;

2 there is a function $A \times Y \to \mathcal{R}$, that is linear across Y, and whose restriction to $A \times X$ represents context preferences.

Context preferences exhibit *nonlinear context-dependence*, when either of (1) and (2) is false.

A more detailed analysis of some of the difficulties associated with the linear approach is provided in section 3. As usual, it is the very global and simple nature of linearity that is both a strength and a weakness. For the purposes of explaining decision making in: prospect theory; discrete choice in econometrics; game theory; the presence of framing effects, or with a physical model of the brain, a more local, nonlinear form context-dependence seems appropriate.

The most important class of nonlinear representations are those that preserve continuity in the way preferences vary across contexts. The concept of continuity across contexts is only meaningful once we have identified a suitable collection τ of subsets of X that is closed under finite intersections and arbitrary unions. In this case, X is a topological space, τ is a topology, and if a set O lies in τ, it is said to be open. Whilst in certain settings it is obvious what sets should be called open, as examples below show, this need not always be the case.

Preferences are (Cac) at x if for any pair a, b of alternatives, such that a is strictly preferred to b at x, there is an open set O that

1 contains x;

2 for any y in O, a is strictly preferred to b at y.

Preferences satisfy (Cac) on Z if they are are (Cac) at x for all x in Z.

A utility representation preserves continuity across contexts if, for any alternative a, the function $U(a, \cdot)$ is continuous at x if and only if preferences are (Cac) at x. In ch.3 of [O'C12], I argue in favour of additional axioms on the way preferences vary for the case where uncertainty is context. Nonetheless, that model, which also allows for non-linearity across contexts, builds upon the present foundations.

This is where the main distinction between classical theory and the present problem lies. As will be shown, even if preferences are continuous across contexts, there is in general no guarantee that the representation will be. It is therefore *a fortiriori* necessary to check that the chosen context space is nice enough to ensure a continuity preserving representation exists. This raises the question of what exactly is meant by 'nice enough', and this is the first question that is definitively answered in this paper: the space of contexts must be a "*perfectly normal* topological space".

Whilst this result does not answer the question of whether it is appropriate to model the human brain, or any other context space, as being "perfectly normal", it does dictate a clear upper bound on the degree of generality modellers may assume when specifying their framework. Perfectly normal spaces are a minimal generalisation of the idea of a 'metrizable space', and as such most of the useful of settings that may arise in applications are covered. A prominent example of a space that is not perfectly normal is the set of probability measures on an uncountable state space. This tells us that generalisations of Savage's [Sav72] subjective expected utility model that allow for nonlinear dependence on beliefs, may have no representation that preserves continuity. A prominent example of a space that is not metrizable, but is perfectly normal is the unit-square with topology generated by the lexicographic order.

On the other hand, it turns out that *each* subset of a space that is perfectly normal is also perfectly normal, and so it is unnecessary to impose the condition that preferences are continuous across *all* contexts. Moreover, since many of the spaces encountered in economics are metric spaces, they are perfectly normal, and although their subsets may not be metric spaces, they too are perfectly normal. Thus the present approach takes a step towards providing a decision theoretic foundation for the theory of discontinuous games [BS12], where a player's payoff depends on the opponent's action in a way that may be discontinuous: consider the usual Bertrand oligopoly settings for instance.

Another reason why the (nonlinear) context preferences approach is well suited to game theory is that the utility representation that is obtained is the most general that preserves the set of Nash equilibria. This is easily seen by considering an example of a game where both players have dominant strategies. Any utility function that assigns a higher value to the dominant strategy, *for each of the opponent's strategies*, preserves the set of Nash equilibria. But interpreting the opponent's strategies as contexts, this is precisely the form of representation obtained here.

Lastly, an important question that arises in relation to the literature on jointly continuous utility is whether the present approach offers an alternative route obtaining a jointly continuous representation. Given that the present representation theorem is separately continuous in alternatives and contexts: can one remove joint discontinuities by exploiting the ordinality of the representation at a given context and the arbitrariness of the values the representation assigns to alternative-context pairs? Whilst the answer to this difficult question is only partially addressed at this stage, the results look promising.

2 Model and Results

2.1 Preferences

The primitives of the model consist of two nonempty sets A and X . Let X denote a set of possible contexts or situations at which the protagonist, Val, might face the problem of deciding amongst certain alternatives. To simplify the exposition, the set A of alternatives is assumed to be the same at each context. In the present, general setting, the question of whether Val is aware of X or her whereabouts in X is left unspecified, but it is assumed that once a context is fixed, some form of ranking of the alternatives according to what she prefers is feasible. That is, for a given context x, and alternatives a and b, she will be able to state whether or not she "strictly prefers" a to b .

Thus for each x in X, Val's preferences are described by a *(context) preference relation* $\succ_x$ which formally is a subset of $A \times A$.[3] This gives rise to a collection of preference relations $\{\succ_x : x \in X\}$, so that the variation of preference, for one alternative *over* another, *across* contexts is explicitly modelled. Where necessary, the more expressive notation $\{(A, \succ_x) : x \in X\}$ is used instead, and brevity favours $\{\succ_x\}_X$ when its status is unambiguous. The term *context preferences* will also refer to this collection of individual context preference relations.[4]

The situation where, for a given context x and pair of alternatives a and b, Val's preferences are such that neither $a \succ_x b$, nor $b \succ_x a$, is denoted by $a \sim_x b$. This situation could just as well be described by $b \sim_x a$. Thus the relation $\sim_x$ is symmetric, and given standard conditions, which are stated below, it is an 'equivalence' relation that characterises indifference between alternatives.

[3] I choose this approach, where strict as opposed to a weak preference relation $\succ_x$ is primitive, because strict preference is unambiguous in its meaning. It is adopted in standard texts such as Fishburn [Fis79] and Kreps [Kre88], and convincingly motivated by Adams [Ada65].

[4] The term context preferences is chosen due to similar terminology being used when state or time preferences are modelled.

As discussed in the introduction, context preferences exclude any preference statements that Val may in fact be in a position to make regarding pairs of contexts, or indeed between one alternative-context pair and another. This information is intentionally ignored so that no assumption need be made concerning preferences over such objects. In the language of measurement theory (see [dG02] for a recent survey) context preferences are low in the information hierarchy. Perhaps the best way to understand the approach is in terms of multiple epistemological viewpoints, each pertaining to a context: there are no hidden independence assumptions.

2.2 Axioms and the Context space

Definition (Asymmetry (Asy.)).
For all $a, b \in A$, $x \in X$: *if* $a \succ_x b$ *then* $\neg(b \succ_x a)$.

Definition (Continuity across contexts (Cac)).
For all $a, b \in A$, $x \in X$: *if* $a \succ_x b$, *then there exists an open neighbourhood* O *of* x *in* X *such that for every* $x \in O$ *we have* $a \succ_q b$.

Note that continuity has the intuitive appeal that it characterizes the stability of strict preferences. That is stability with respect to perturbations across contexts.

Definition (Negative transitivity (NT)).
For all $a, b, c \in A$ *and* $x \in X$: *if* $a \succ_x c$, *then either* $b \succ_x a$ *or* $c \succ_x b$; *equivalently* $\neg(a \succ_x b)$ *and* $\neg(b \succ_x c)$, *then* $\neg(a \succ_x c)$.

It is well known that conditions (NT) and (Asy.) on a binary relation $\succ_.$ are together equivalent to assuming that the union $\succsim_.$ of $\succ_.$ and $\sim_.$ is both complete and transitive–see [Fis79] ch.2 for instance.

The main theorem holds for spaces that are *perfectly normal*. A topological space X is said to be *normal* if for every pair of disjoint closed subsets of X there exist disjoint open sets containing A and B respectively. Then a topological space X is perfectly normal if for every set C that is closed in X, there exists a real-valued function f such that $C = f^{-1}(0)$. An equivalent definition is that every such C can be written as a countable intersection of sets that are open in X ([Mun00] p.229). The following theorem is a recent, intuitive restatement of Michael's [Mic56] selection theorem due to [GS00]. It provides a useful characterisation of perfectly normal topogical spaces.

Theorem 2.1 (Michael's selection theorem). *The following two statements are equivalent.*

1) X is a perfectly normal topological space.

2) If $g, h : X \to \mathcal{R}$ are upper and lower semi-continuous respectively and $g \leqslant h$, then there is a continuous $f : X \to \mathcal{R}$ such that $g \leqslant f \leqslant h$ and $g(x) < f(x) < h(x)$ whenever $g(x) < h(x)$.

This equivalence is relevant for preferences that are indexed by elements of a context space because if the context space is not perfectly normal then *there exist $g, h : X \to \mathcal{R}$* that are upper (resp. lower) semi-continuous and $g \leqslant h$, such that *for no* continuous $f : X \to \mathcal{R}$ do we have $g \leqslant f \leqslant h$ and $g(x) < f(x) < h(x)$ whenever $g(x) < h(x)$. This, as we will see in the proof of our theorem and subsequent discussion, would imply that there exists context preferences $\{(A, \succ_x) : x \in X\}$ that are continuous across contexts such that there is no representation $U : A \times X \to \mathcal{R}$ satisfying continuity across contexts.

Note that for countable S the simplex of probability measures Δ on S is a subspace of $\mathcal{R}^S$, which is a 'Hilbert space' under the standard Euclidean metric, so that by [SS70] p.65, Δ is a metric space. Then by [Mun00] p. 229 every metrizable space is perfectly normal. Thus for countable S, Δ is a suitable context space. This is not true for uncountable S, by counterexample 105, on p125 of [SS70].

Another example of a context space in the literature is any countable product of the discrete space of non-negative integers with the Cartesian product topology as is found in the case-based decision theory of Gilboa and Schmeidler [GS01, GS03b]. By [SS70] p.121, this is a complete metric space, and so this too is a suitable space of contexts for a nonlinear representation. On the other hand, uncountable products of the nonnegative integers with the product topology are, by counterexample 103 of [SS70], not normal spaces, and so they may be unsuitable depending on preferences.

2.3 Result

The representation we are seeking is of the following form.

Definition 2.1. *$U : A \times X \to \mathcal{R}$ is said to be a* context utility representation *of preferences $\{(A, \succ_x) : x \in X\}$ if for all $a, b \in A$ and $x \in X$,*

$$a \succ_x b \quad \Leftrightarrow \quad U(a, x) > U(b, x).$$

It is an ordinal if, for any other representation V of preferences, there exists a family of strictly increasing functions $\{f_x : \mathcal{R} \to \mathcal{R}\}_{x \in X}$ such that for each x, $V(\cdot, x) = f_x \circ U(\cdot, x)$.

Remark 2.1. *Note that in definition (2.1), the continuity of $f_{.}$ across contexts is implied by the same form of continuity of the functions U and V .*

I now state and prove our representation theorem for context preferences that satisfy the aforementioned conditions.

Theorem 2.2. *Let A be discrete and countable and X be a perfectly normal topological space of contexts. The following two statements are equivalent:*

1) *context preferences $\{(A, \succ_x) : x \in X\}$ satisfy (Asy.), (NT) and are (Cac) on $Z \subset X$;*

2) *context preferences have a ordinal context utility representation that is separately continuous on $A \times Z$.*

Proof. Let $\{1, 2, 3 \ldots\}$ be an arbitrary enumeration of A, and by $[j]$ we will denote the subset of A that contains the first j elements of the enumeration. By $U^j : [j] \times X \to \mathcal{R}$ we will denote the utility representation of the projection of preferences $\{(A, \succ_x) : x \in X\}$ onto the first j elements of the enumeration. That is, if we recall that for each $x \in X$, $\succ_x$ is a subset of $A \times A$, then we see that $\{\succ_x : x \in X\} \subset (A \times A)^X$. Hence by the projection of preferences onto $[j]$ we mean

$$\{\succ_x : x \in X\} \cap \big([j] \times [j]\big)^X$$

which is a well defined intersection since

$$\begin{aligned}\big([j] \times [j]\big)^X \cap (A \times A)^X &= \prod_{x \in X} \big([j] \times [j]\big) \cap (A \times A) \\ &= \prod_{x \in X} ([j] \cap A) \times ([j] \cap A) \\ &= \big([j] \times [j]\big)^X.\end{aligned}$$

We will use this projection to proceed by induction on A. Thus we first show that a continuous representation for the basic case: which we take to be $j = 1$.

Let $U^1(1, x) \equiv 0$ for all $x \in X$. By condition (Asy.), U^1 is a representation for the projection of preferences onto $[1]$ and it is clearly continuous. This completes the proof for the basic case. Now suppose that for some $j \in \mathbb{N}$ greater than 1 there exists a representation U^{j-1} of the projection onto $[j - 1]$. We need to show that the conditions imply the existence of a representation of the projection onto $[j]$.

For $a \in [j-1]$ we set $U^j(a,\cdot) = U^{j-1}(a,\cdot)$. Then by the induction hypothesis, for all $a, b \in [j-1]$ and $x \in X$ we have,

$$a \succ_x b \quad \Leftrightarrow \quad U^j(a,x) > U^j(b,x),$$

and on this subset of $[j]$, U^j is continuous.

We now need to select a continuous function $U^j(j,\cdot)$ on X such that for each x, $U^j(\cdot,x) : [j] \to \mathcal{R}$ represents $\succ_x \cap [j] \times [j]$. We will do this by first defining the lower and upper envelopes, $\underline{U}$ and $\overline{U}$ respectively of $U^j([j-1],X)$ relative to alternative j. That is, informally speaking, for the lower envelope relative to j we seek the function whose graph is the set of points $\{(x, U^j(\cdot,x)) : x \in X\}$ that can be seen by looking down from the position of j in the preference order at context x. (Similarly, the upper envelope relative to j, it is the set of points that can be seen by looking up.)

We then show that these two functions are respectively upper and lower semi-continuous and that the latter dominates the former pointwise. This, together with the fact that X is perfectly normal implies, via Michael's selection theorem that the required function $U^j(j,\cdot)$ exists.

First some useful notation is introduced.[5] Let $\mathcal{B}_{jk}$ refer to the set of x such that alternative j is better than alternative k. Similarly, $\mathcal{W}_{jk}$ refers to the set of context where k is strictly better than j. The set $\mathcal{N}_{jk}$ is the set of contexts where indifference holds between j and k. Analogously, for any subset D of A, let $\mathcal{B}_{jD}$ denote the set of contexts such that j is strictly better than every alternative in D, and $\mathcal{W}_{jD}$ the set of contexts such that j is strictly worse than each of the alternatives in the set D.

Let us now define the lower and upper envelopes of U^j relative to j.

$$\underline{U}(x) := \begin{cases} \min_{k\in[j-1]} \left\{U^j(k,x)\right\} - 1 & \text{if } x \in \mathcal{W}_{j[j-1]} \\ \max_{k\in[j-1]} \left\{U^j(k,x) : j \succsim_x k\right\} & \text{otherwise.} \end{cases}$$

This function is well defined for the following reasons: firstly, $[j-1]$ is compact; secondly, for each x in

$$X \backslash \mathcal{W}_{j[j-1]} \equiv \bigcup_{k=1}^{j-1} \left(\mathcal{B}_{jk} \cup N_{jk}\right)$$

[5] Similar notation is used elsewhere in the literature on context preferences. See [GS03a] for instance.

there exists $k \in [j-1]$ such that $j \gtrsim_x k$; and thirdly, the fact that $\gtrsim_x$ is complete and transitive means that if D_x is the subset of alternatives in $[j-1]$ that j weakly dominates at x, there is at least one element of D_x that is maximal in D_x under $\gtrsim_x$. Similarly, the upper envelope of $U^j([j-1], X)$ relative to j is also well defined as follows:

$$\overline{U}(x) := \begin{cases} \max_{k \in [j-1]} \left\{U^j(k,x)\right\} + 1 & \text{if } x \in \mathcal{B}_{j[j-1]} \\ \min_{k \in [j-1]} \left\{U^j(k,x) : k \gtrsim_x j\right\} & \text{otherwise.} \end{cases}$$

Claim 2.1. *For all $x \in X$, $\underline{U}(x) \leqslant \overline{U}(x)$.*

Proof. With a view to obtaining a contradiction, suppose that for some $x \in X$, $\overline{U}(x) < \underline{U}(x)$. Then by definition x cannot be in the union of $\mathcal{W}_{j[j-1]}$ and $\mathcal{B}_{j[j-1]}$. Thus for some $k, l \in [j-1]$, $\underline{U}(x) = U^j(k,x)$ and $\overline{U}(x) = U(l,x)$. Now once more by definition,

$$l \gtrsim_x j \gtrsim_x k,$$

so that by condition (NT) it follows that: $l \gtrsim_x k$. Then since U^j is equal to U^{j-1} on $[j-1] \times X$, which, by the induction hypothesis, represents the projection of preferences onto $[j-1]$, we have $U(l,x) \geqslant U(k,x)$. This is the required contradiction, and so we see that $\underline{U}$ is pointwise weakly dominated by $\overline{U}$. □

Purely for expositional purposes, we introduce two fictional alternatives $\underline{a}$ and $\overline{a}$, such that for all $x \in X$ and $k \in [j]$, we have $\overline{a} \succ_x k \succ_x \underline{a}$. Accordingly, we define $[j-1]^* := [j-1] \cup \{\underline{a}, \overline{a}\}$, and for each $x \in X$, let

$$U^j(\underline{a}, x) \equiv \min_{k \in [j-1]} \{U^j(k,x)\} - 1$$

and

$$U^j(\overline{a}, x) \equiv \max_{k \in [j-1]} \{U^j(k,x)\} + 1.$$

Now define the lower envelope of $U^j([j-1]^*, X)$ relative to j, and note that for all $x \in X$ there exists $k \in [j-1]^*$ such that $j \gtrsim_x k$, so that

$$\underline{U}(x) := \max_{k \in [j-1]^*} \left\{U^j(k,x) : j \gtrsim_x k\right\}$$

is well defined and similarly, so is

$$\overline{U}(x) := \min_{k \in [j-1]^*} \left\{U^j(k,x) : k \gtrsim_x j\right\}.$$

Now, by construction, the lower and upper envelopes of $U^j([j-1], X)$ relative to j are respectively identical to the lower and upper envelopes of $U^j([j-$

$1]^*, X)$ relative to j. Moreover, we can readily see that for all $x \in X$ if for some $k \in [j-1]^*$, $k \sim_x j$, then $\underline{U}(x) = U^j(k,x) = \overline{U}(x)$; on the otherhand, if $\underline{U}(x) = \overline{U}(x)$ for some x, then there exists $k \in [j-1]$, for it cannot be $\underline{a}$ or $\overline{a}$, such that $k \sim_x j$. So we see that there exists $\underline{U}(x) = \overline{U}(x)$ if and only if there exists $k \in [j-1]$ such that $k \sim_x j$. Equivalently, for each $x \in X$, if $\underline{U}(x) = U^j(g,x)$ and $\overline{U}(x) = U^j(l,x)$ for some $g, l \in [j-1]^*$, then

$$\overline{U}(x) > \underline{U}(x) \quad \Leftrightarrow \quad l \succ_x g,$$

and for all other $k \in [j-1]^*$, either $k \succ_x l$ or $g \succ_x k$. This, together with claim (2.1), shows that provided $\underline{U}$ is usc and $\overline{U}$ is lsc, Michael's selection theorem tells us that the desired continuous function exists.

We now show that $\underline{U}$ is upper semi continuous (usc). We do so by showing that $\underline{U}$ is infact continuous everywhere except the contexts where alternative j changes from being strictly worse to being indifferent to some other alternative(s) and the set of alternatives that j dominates is unchanged. At such points, $\underline{U}$, should, intuitively speaking, increase because the set of elements of $[j-1]^*$ that are weakly dominated by j at x will have increased in cardinality by at least one, and we recall that $\underline{U}$ is defined in terms of the maximum over D_x. The fact that the alternative space is discrete and strict preference is continuous then implies that there is a jump up in the value of $\underline{U}$ in such contexts, a property that is satisfied by usc functions.

We now provide a formal proof of this argument.

Claim 2.2. *On X, $\underline{U}$ is upper semicontinuous and $\overline{U}$ is lower semicontinuous.*

Proof. For each $x \in X$, let D^*_x be the set of elements of $[j-1]^*$ that are weakly dominated by j at x. Similarly, let $E^*_x := \{k \in [j-1]^* : k \succsim_x j\}$. Condition (Asy.) and the definition of $\sim_\cdot$ imply that for all $x \in X$ and $k \in [j-1]$, exactly one of following relationships must hold:

$$j \succ_x k, \quad k \succ_x j \quad \text{or} \quad k \sim_x j\,.$$

So that for each x, we have $D^*_x \cup E^*_x = [j-1]^*$. Furthermore, $l \in D^*_x$, $k \in E^*_x$ implies that $l \succsim_x k$ by transitivity of $\succsim_x$, and $l \sim_x k$ holds if and only if $D^*_x \cap E^*_x \neq \varnothing$, which in turn is true, if and only if $x \in \mathcal{N}_{jk}$ for some $k \in [j-1]$. When $D^*_x \cap E^*_x$ is empty, therefore, the two sets form a partition of $[j-1]^*$.

We now prove the claim that $\underline{U}$ is upper semicontinuous.

Let $x \in X$ be such that for any open set O that contains x and is sufficiently small, we have $D_z^* = D_x^*$ for all $z \in O$. Thus for each $z \in O$, $\underline{U}(z) = U^j(k,z)$ for some $k \in D_x^*$ and so

$$O = \bigcup_{k \in D_x^*} \bigcap_{l \in D_x^*} \{x \in O : k \succsim l\} = O \cap \big(\bigcup_{k \in D_x^*} \Delta \backslash \mathcal{W}_{kD_x^*} \big) .$$

By condition (C'ty) and the fact that the finite intersection of open sets is open, we see that $\mathcal{W}_{kD_x^*}$ is open in X and so its complement is closed. Therefore in the subspace topology of O, $O \cap X \backslash \mathcal{W}_{kD_x^*}$ is also closed. Now on each of the sets $O \cap \Delta \backslash \mathcal{W}_{kD_x^*}$, $\underline{U} = U^j(k, \cdot)$, and by the induction hypothesis, $U^j(k, \cdot)$ is continuous on X and hence continuous on each of its subsets, so in particular the restriction of $U^j(k, \cdot)$ to $O \cap \Delta \backslash \mathcal{W}_{kD_x^*}$ is a continuous function.

In the intersection of any pair of sets $\Delta \backslash \mathcal{W}_{kD_x^*}$, $\Delta \backslash \mathcal{W}_{lD_x^*}$, $k, l \in D_x^*$ it is clear that by condition (Asy.) we cannot have strict preference in either direction between k and l. Thus for any $z \in O$ in such an intersection we have:

$$U^j(k, \cdot) = \underline{U}(\cdot) = U^j(l, \cdot) .$$

In the same way, for any $z \in O$ and $D \subset D_x^*$ such $\bigcap_{k \in D} \Delta \backslash \mathcal{W}_{kD_x^*}$ we have $U^j(k,z) = \underline{U}(z)$ for all $k \in D$.

Now let C be any closed subset of $\mathcal{R}$. Then we have

$$O \cap \underline{U}^{-1}(C) = O \cap \big(\bigcup_{k \in D_x} U^j(k, \cdot)^{-1}(C) \big) .$$

Continuity of $U^j(k, \cdot)$ for each $k \in D_x$ implies that $U^j(k, \cdot)^{-1}(C)$ are closed in X, and the fact that D_x is finite implies that the union on the right-hand-side of this equation is closed in X. Hence, in the subspace topology, both sides are closed in O and therefore $\underline{U}$ is continuous at x.[6]

By condition (C'ty), the above argument accounts for $x \in \mathcal{B}_{kD} \cap \mathcal{W}_{kE}$, for each partition D, E of $[j-1]^*$. That is, all points x such that $j \not\sim_x k$ for any $k \in [j-1]$. On the interior of the set $\bigcup_{k=1}^{j-1} \mathcal{N}_{jk}$ relative to X, there exists $k \in [j-1]$ and an open neighborhood O of x that is contained in $\mathrm{Int}_X \mathcal{N}_{jk}$. As such, on O, $\underline{U}$ is equal to the continuous function $U^j(k, \cdot)$ and is thereby continuous.

[6]This argument is based on what is called "The pasting lemma" (Munkres p.124).

It remains for me to consider $x \in \mathrm{bd}_X \mathcal{N}_{jk}$ for arbitrary $k \in [j-1]$. Here, in every open neighborhood O of x, there exists $y \in O$ such that $y \in \mathcal{B}_{jk} \cup \mathcal{W}_{jk}$. Let $\{O_n : n \in \mathbb{N}\}$ be a sequence of open sets that contain x which satisfies $\bigcap_{n=1}^{\infty} O_n = \{x\}$. (Such a sequence exists precisely because X is perfectly normal.)

We will consider sequences of contexts $\{y_n\}$ such that $y_n \in O_n$ for each n, so that $\lim_n y_n = x$. We do so by first partitioning each set O_n into the three sets

$$O_n \cap \mathcal{N}_{jk}, \quad O_n \cap \mathcal{W}_{jk} \quad \text{and} \quad O_n \cap \mathcal{B}_{jk}.$$

Any such sequence of contexts $\{y_n\}$ that has infinitely many elements that lie in more than one of the three sets that determine the partition, say $\mathcal{B}_{jk}$ and $\mathcal{W}_{jk}$, has convergent subsequences $\{y'_{n_m}\}$ and $\{y''_{n_m}\}$ that lie wholly in $\mathcal{B}_{jk}$ and $\mathcal{W}_{jk}$ respectively. So it suffices to consider sequences in each of the partitions separately.

For each sequence $\{y_n \in O_n \cap \mathcal{N}_{jk}\}$ we have $\underline{U}(y_n) = U^j(k, y_n)$ for each $n \in \mathbb{N}$, and by the induction hypothesis $U^j(k, \cdot)$ is continuous, so $\underline{U}(y_n)$ converges to $\underline{U}(x)$.

For the sequence of sets $\{O_n \cap \mathcal{W}_{jk}\}$, we note that $z \in O_n \cap \mathcal{W}_{jk}$ implies that $k \succ_z j$, and, by condition (NT), for all $l \in D_z$ we have $k \succ_z l$, so that $U^j(k, z) > \underline{U}(z)$. Let us consider the following two exhaustive cases.

Case 1. There exists $n \in \mathbb{N}$ such that for all $z \in P_n \equiv \bigcup_{m \geqslant n} (O_m \cap \mathcal{W}_{jk}) \cup \{x\}$: if $l \in D_z$, then $k \succ_x l$. That is, there exists an open neighborhood Q of x such that for all $y \in (Q \cap \mathcal{W}_{jk}) \cup \{x\}$, $U^j(k, y) > \underline{U}(y)$. Thus, for any sequence of contexts $\{y_m \in O_m \cap \mathcal{W}_{jk} : m \in \mathbb{N}\}$, $\{y_m\}$ converges to $x \in \mathrm{bd}_X \mathcal{W}_{jk}$ and moreover,

$$\limsup_m \underline{U}(y_m) < U(x) = U^j(k, x)$$

Thus in this case, $\underline{U}$ is x is upper semicontinuous.

Case 2. (This is the negation of Case 1.) That is for all open neighborhoods Q of x, there exists $z \in (Q \cap \mathcal{W}_{jk}) \cup \{x\}$ with $U^j(k, z) \leqslant \underline{U}(z)$. Now if $z \neq x$, then $z \in \mathcal{W}_{jk}$ and $U^j(k, z) \leqslant \underline{U}(z)$ imply $j \succsim_z l \succsim_z k \succ_z j$, which contradicts condition (NT). Thus $z = x$ is the only context such that for all $n \in \mathbb{N}$ there exists $z \in P_n$ and $l \in D_z$ such that $l \succsim_x k$. Once more by condition (NT) we see that for such l, $l \sim_x k$.

Now note that for each $y_n \in O_n \cap \mathcal{W}_{jk}$ we have $U^j(l, y_n) \leqslant \underline{U}(y_n) < U^j(k, y_n)$;

this is equivalent to

$$0 \leqslant \underline{U}(y_n) - U^j(l, y_n) < U^j(k, y_n) - U^j(l, y_n)\,.$$

Then using the triangle inequality on $\mathcal{R}$, the above bound, and the fact that $U^j(k, x) = \underline{U}(x) = U^j(l, x)$ we have:

$$\begin{aligned}
0 &\leqslant |\underline{U}(y_n) - \underline{U}(x)| \\
&\leqslant |\underline{U}(y_n) - U^j(l, y_n)| + |U^j(l, y_n) - \underline{U}(x)| \\
&\leqslant |U^j(k, y_n) - U^j(l, y_n)| + |U^j(l, y_n) - \underline{U}(x)| \\
&\leqslant |U^j(k, y_n) - \underline{U}(x)| + 2|U^j(l, y_n) - \underline{U}(x)|\,.
\end{aligned}$$

Thus $|\underline{U}(y_n) - \underline{U}(x)|$ converges to 0 by continuity of $U^j(k, \cdot)$ and $U^j(l, \cdot)$ via the induction hypothesis. This completes the proof of Case (2). Indeed, in this case, if for instance $\mathcal{B}_{jk}$ is empty, $\underline{U}$ is in fact continuous at the boundary of $\mathcal{W}_{jk}$.

We now turn to the remaining sequence of sets $\{O_n \cap \mathcal{B}_{jk}\}_{n\in\mathbb{N}}$. Now as above, for each n, let $P_n \equiv \bigcup_{m\geqslant n}(O_m \cap \mathcal{B}_{jk}) \cup \{x\}$. Now define the sequence of subsets of $[j-1]^*$

$$F_n := \bigcup_{y\in P_n} D_y,$$

then because $F_n = F_{n+1} \cup \bigcup_{y\in O_n\cap\mathcal{B}_{jk}} D_y$, F_n is a decreasing sequence of sets: each of which contains the alternative k . Now since $[j-1]^*$ is a finite set, there exists $n' \in \mathbb{N}$ such that for all $n \geqslant n'$, $F_n = F_{n'} \equiv F$.

Once more there are two cases. The first holds when there exists $n' \in \mathbb{N}$ such that $y, z \in P_{n'}$ implies $D_y = D_z$. By the construction of $\{P_n\}$, this implies that the same is true for all $n \geqslant n'$. In this case we may use the partition lemma argument on $P_{n'}$ to show that the restriction of $\underline{U}$ to P_n' is continuous. Although, $P_{n'}$ is not a neighborhood of x, it is clear that for any sequence $\{y_n\} \subset P_{n'}$ that is converging to x, we have $\underline{U}(y_n)$ converges to $\underline{U}(x)$.

The remaining case is where for all $n \in \mathbb{N}$, there exists $y, z \in P_n$ such that $D_y \neq D_z$. Let us define the set

$$L := \bigcap_n \{l \in F : l \in D_y \backslash D_z, \text{for some } y, z \in P_n\}.$$

The fact that $P_n \subset \mathcal{B}_{jk} \cup \{x\}$ implies that $k \in D_y \cap D_z$ for all $y, z \in P_n$, thus $F \neq L_n$ for all n .

Suppose that $l \in L$, and $l \succ_x j$. Then by condition (C'ty) there exists an open neighborhood Q of x such that $l \notin D_y$ for all $y \in Q$. Now for n sufficiently large, the fact that $\bigcap_n O_n = \{x\}$ implies that $P_n \subset Q$. Thus, either $P_n \cap Q = \{x\}$ (if we are considering the discrete topology on X), or, by definition of L, there exist $y, z \in Q$ such that $l \in D_y \backslash D_z$. In either case we obtain a contradiction: in the former it is the fact that we have found $n \geqslant n'$ such that $D_y = D_z$ for all $y, z \in P_n$; and in the second, it is the definition of Q.

Now suppose that $j \succ_x l$. In this case there exists an open neighborhood Q of x such that $l \in D_y$ for all $y \in Q$. In this case, we see that like k, $l \in F \backslash L$.

Thus, by default, we see that $l \sim_x j$. Now let $\{z_n\}$, $z_n \in P_n$ be such that $l \in D_{y_n} \backslash D_{z_n}$ for some $y_n \in P_n$. Thus, for all z_n we have $l \succ_{z_n} j$. Now since z_n converges to x, we know that for every open neighborhood Q of x there exists n such that $z_n \in Q$. Thus, $x \in \mathrm{bd}_X \mathcal{W}_{jl}$. This case is therefore identical to Case (2) above with the roles of k and l reversed.

This completes the proof of the fact that $\underline{U}$ is upper semicontinuous on X. The proof that $\overline{U}$ is lower semicontinuous is identical except for the fact that the points of discontinuity lie on the boundary of $\mathcal{B}_{j\cdot}$ rather than $\mathrm{bd}_X \mathcal{W}_{j\cdot}$ as we have shown to be the case for $\underline{U}$. □

Finally, to extend the result to the countably infinite case we appeal to the axiom of (countable) choice and let $U(j, \cdot) = U^j(j, \cdot)$ for each $j \in \mathbb{N}$. □

3 Discussion

Preserving continuity across contexts is clearly a desirable property of any utility representation of preferences. The fact that the main result of this paper does not impose continuity upon preferences at all contexts means that the present representation should not revive the discussion concerning the empirical relevance of continuity. Whenever it is natural for continuity across contexts to hold, the present result tells us that provided the context space is perfectly normal, there exists a utility representation.

On the other hand, if there are subsets of the context space where continuity across contexts fails to hold, such preferences are still within the scope of the present result. At first this may seem to be an unnecessary level of generality, but the existing literature on discontinuous games (of which

Bertrand Oligopoly is a prominent example) and the absence of related results in decision theory suggests otherwise.

The fact that the result is not confined to continuous context dependence and the fact that the mathematics involved draw from the field of nonlinear analysis and topology justifies the title of nonlinear context dependence. This title facilitates the comparison with the existing theory of context dependence that has been developed by Gilboa and Schmeidler–principally through their axiomatization of case-based decision theory.

Whilst generality that is both the strength and weakness of the nonlinear approach, there are some irrefutable reasons to prefer it over the linear approach. There are two issues that are troublesome with the linear approach to context dependence. The first is that preferences must vary a lot across the context space in order to guarantee a representation that is simultaneously separable across A and linear across X. It will not do for instance for an alternative to be the best in all contexts, or for just four of the possible six complete and transitive strict orderings of three alternatives to be present in the collection of preference relations that form the decision maker's context preferences as is the case in figure (3.1).

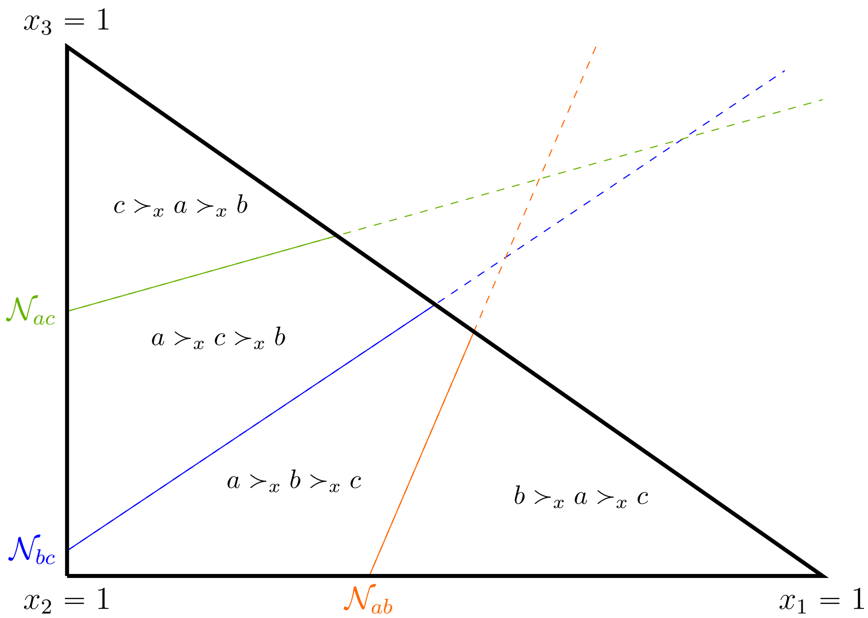

Figure 3.1: Here preferences that satisfy (Asy.), (NT), (Cac), a strong convexity across contexts condition (the combination axiom of [GS03a]), but which are not 3-diverse. Since the three lines are not congruent (to be congruent they should meet at a point), there is no linear representation. Such preferences are nonlinear.

However, even diversity over triples of alternatives is not enough. Figure (3.2) depicts the case where for each set of three distinct alternatives there are $3! = 6$ contexts each corresponding to one of the six possible orderings of the three alternatives. Nonetheless, there is no utility representation that is linear across contexts because there is not enough diversity over quadruples of alternatives. A minimal condition that rules out such preferences and is sufficient for linearity to hold is independence of the vectors that are normal to the hyperplanes defined by indifference and is discussed in [O'C12]. However, even this condition is so strong as to imply that for a context space of dimension two and a set of three or more alternatives, there are no context preferences satisfying the conditions for a linear representation.[7]

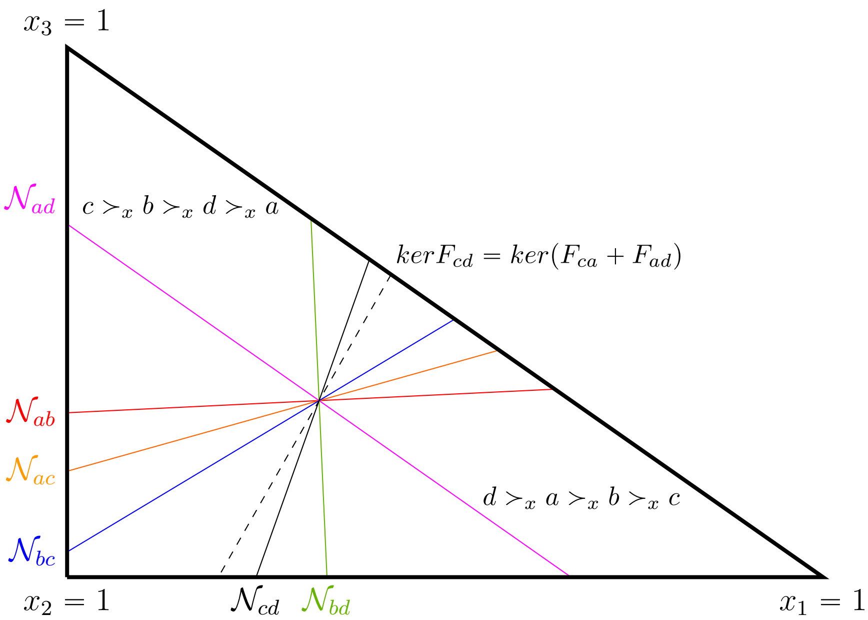

Figure 3.2: Here, preferences satisfy (Asy.), (NT), (Cac), a strong convexity across contexts condition (the combination axiom of [GS03a]), and 3-diversity, but not 4-diversity. For preferences to have a linear representation, the solid black line would have to lie on the dashed black line. There is no behavioural reason why this should be so. Such preferences are nonlinear.

It worth recalling that there is no counterpart to such conditions in the context-free [vNM44] setting. Indeed in the uncertainty setting, ignoring the issue of uniqueness of the representation: for each family of context preferences that Gilboa and Schmeidler [GS03a] represent, if consequences

[7] Ashkenazy and Lehrer [AL01] are to my knowledge the first to discuss this matter and [O'C12] discusses this in some detail.

are identified with the set $A \times X$, there is a suitably defined vNM expected utility representation [O'C12]. Due to the diversity condition, it is clear that the converse is not true. Moreover, whilst one would not wish to rule such diversity of preference out, there is no behavioural reason to insist upon it unless we choose to allow the context space to vary instead of treating it as a primitive of the model.

It is the global nature of linearity that forces diversity upon the modeller. Typically a local approach will be more appropriate, and this seems particularly true, given the largely distinct regions of the brain, when the connectome is context [Seu12].

Another issue with linearity is that it implies a "thinness" of the set of contexts where strict preference does not hold in either direction between two alternatives. An informal way to understand this point is to note that in finite dimensions a hyperplane is of dimension one lower than the ambient space, and as such it has empty interior. Even in the general case, linear context dependence implies that if for some context x and alternatives a and b we have no strict preference in either direction, then, informally speaking, almost all the contexts y near x satisfy $a \succ_y b$ or $b \succ_y a$.[8] The work of Deco and Rolls [RD10] suggests that this knife-edged kind of "instability" of no-strict preference *across contexts* is a characteristic of schizophrenia. Healthy decision-makers instead have a range where they are indecisive and for decision-makers with obsessive-compulsive disorders, the range is large.

Whilst the nonlinear approach presented here offers an alternative to those who wish to avoid ruling out reasonable preferences, it does come at a cost. The main ones being the absence of a cardinal scale and the related cost of eliciting a preferences. That is to say, if the decision maker does find the conditions for linearity agreeable, then there is a significant reduction in the number of binary choice questions and answers at different contexts that are needed to identify a linear utility function.

4 Future work

At this stage, it remains to shown that the present approach can be extended to uncountably many alternatives. This would extend the result to apply to a much wider variety of settings including the Kőszegi–Rabin model of reference-dependence. It would also facilitate a more detailed study of the interesting fact that the unit square endowed with the topology generated by a lexicographic ordering is a perfectly normal space.

[8] Thinness is implied by the combination axiom of Gilboa and Schmeidler.

It also remains to shown that joint discontinuities are not "essential" in the sense that they may be removed by exploiting the ordinal nature of the representation together with the fact that the ordering over alternative-context pairs is arbitrary in so far as preferences are concerned. Simple examples, and sketch proofs indicate that these extensions may well be possible provided the context space is metrizable. There is a large mathematical literature on the relationship between separate and joint continuity, and applications to the present problem look promising.

Another interesting question is whether or not the literature on jointly continuous utility representations can be intuitively applied to settings with context-dependence in decision theory.

A less technical, yet equally interesting, question concerns the appropriate definition of an open set in the setting of case-based decision theory. In the linear setting this issue does not arise, but to borrow from [MMB+12]: it is important to know what constitutes a "novel stimulus" and what constitutes "background, unimportant noise" in case-based decision theory.

Finally, there appear to be sound reasons for using the context preferences, as opposed to a single preference relation, to model unawareness. For similar reasons, applications to the field of neuroscience are also promising.

References

[Ada65] E. W. Adams. Elements of a theory of inexact measurement. *Philosophy of Science*, 32(3/4):205–228, 1965.

[AL01] G. Ashkenazi and E. Lehrer. Relative utility. Preprint available at `www.math.tau.ac.il/ lehrer/Papers/relativeutil8.pdf`, 2001.

[Bar10] Andrei Barbos. Context effects: A representation of choices from categories. *Journal of Economic Theory*, 145:1224–1243, 2010.

[Ble07] H. Bleichrodt. Reference-dependent utility with shifting reference points and incomplete preferences. *Journal of mathematical psychology*, 51:266–276, 2007.

[BM95] D. S. Bridges and G. B. Mehta. *Representations of Preference Orderings*. Springer-Verlag, Berlin–Heidelberg, 1995.

[BS12] Paulo Barelli and Idione Soza. Existence of nash equilibria in relative and discontinuous games. Preprint available at http://www.uq.edu.au/economics/..., 2012.

[CCM09] Alessandro Caterino, Rita Ceppitelli, and Francesca Maccarino. The euclidean distance construction of order homomorphisms. *Applied General Topology*, 10:187–195, 2009.

[Deb54] G. Debreu. Representation of a preference ordering by a numerical function. In R. M. Thrall, C. H. Coombs, and R. L. Davis, editors, *Decision processes*, pages 159–165, 1954.

[dG02] C. d'Aspremont and L. Gevers. Social welfare functionals and interpersonal comparability. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of social choice and welfare*, volume 1, pages 459–541, 2002.

[DR04] Jacques H. Drèze and Aldo Rustichini. State-dependent utility and decision theory. In S. Barberá, P. J. Hammond, and C. Seidl, editors, *Handbook of Utility Theory*, pages 839–892, 2004.

[Eil41] S. Eilenberg. Ordered topological spaces. *American Journal of Mathematics*, 63:39–45, 1941.

[Fis79] P. C. Fishburn. *Utility theory for decision making*. Publications in operations research. R. E. Krieger Pub. Co., 1979.

[GS95] I. Gilboa and D. Schmeidler. Case-based decision theory. *The Quarterly Journal of Economics*, 110(3):605–639, 1995.

[GS00] Chris Good and Ian Stares. New proofs of classical insertion theorems. *Comment. Math. Univ. Carolinae*, 41(1):139 – 142, 2000.

[GS01] I. Gilboa and D. Schmeidler. *A theory of case-based decisions*. Cambridge University Press, 2001.

[GS03a] I. Gilboa and D. Schmeidler. A derivation of expected utility maximization in the context of a game. *Games and Economic Behavior*, 44(1):172 – 182, 2003.

[GS03b] Itzhak Gilboa and David Schmeidler. Inductive inference: An axiomatic approach. *Econometrica*, 71(1):pp. 1–26, 2003.

[KR06] B. Kőszegi and M. Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006.

[KR07] B. Kőszegi and M. Rabin. Reference-dependent risk attitudes. *American Economic Review*, 97(4):1047–1073, 2007.

[Kre88] D. M. Kreps. *Notes on the theory of choice*. Underground classics in economics. Westview Press, 1988.

[KT79] D. Kahneman and A. Tversky. Prospect theory: and analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

[Lev83] V. Levin. A continuous utility theorem for closed preorders on a σ-compact metrizable space. *Soviet Math. Doklady*, 28:715–718, 1983.

[McF04] Daniel L. McFadden. Revealed stochastic preference: A synthesis. Preprint available at http://hdl.handle.net/10.1007/s00199-004-0495-3, 2004.

[Meh98] Ghanshyam B. Mehta. Preference and utility. In S. Barberá, P. J. Hammond, and C. Seidl, editors, *Handbook of Utility Theory*, volume 1, pages 1–47, 1998.

[Mic56] E. Michael. Continuous selections i. *The Annals of Mathematics*, 63(2):361–382, 1956.

[MMB+12] Jorge F Mejias, Gary Marsat, Kieran Bol, Erik Harvey-Girard, Leonard Maler, and Andre Longtin. Signal cancellation and contrast invariance in electrosensory systems. *BMC Neuroscience*, 13(Suppl 1):F2, 2012.

[MO05] Yusufcan Masatlioglu and Efe A Ok. Rational choice with status quo bias. *Journal of Economic Theory*, 121(1):1–29, 2005.

[Mun00] J. R. Munkres. *Topology*. Prentice Hall, 2nd edition, 2000.

[O'C12] Patrick H.D. O'Callaghan. Context and decision: three approaches to representing preferences. *University of Warwick PhD Thesis*, 2012.

[PX12] Prasanta K. Pattanaik and Yongsheng Xu. On dominance and context-dependence in decisions involving multiple attributes. *Economics and Philosophy*, 28:117–132, 2012.

[RD10] E. T. Rolls and G. Deco. *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function.* Oxford University Press, Oxford, 2010.

[Sav72] L. J. Savage. *The foundations of statistics.* Dover Publications, 1972.

[Seu12] S. Seung. *Connectome: How the Brain's Wiring Makes Us Who We Are.* Houghton Mifflin Harcourt, 2012.

[SS70] L. A. Steen and J. A. Seebach. *Counterexamples in Topology.* Hot, Rinehart and Winston, Inc., USA, 1970.

[vNM44] J. von Neumann and O. Morgenstern. *Theory of games and economic behavior.* Princeton University Press, Princeton and Oxford, Sixtieth anniversary edition, 1944.

[Wol44] H. Wold. A synthesis of pure demand analysis i, ii and iii. *Skandinavisk Actuarietidskrift*, 26:85–118, 220–263, 69–120, 1943–44.

PREVIOUS WORKING PAPERS IN THE SERIES

RISK AND UNCERTAINTY PROGRAM

R03_1 Bargaining power and efficiency in principal-agent relationships, Robert G Chambers and John Quiggin (August 2003).

R03_2 Information and the risk-averse firm, Robert G. Chambers and John Quiggin (2003).

R03_3 Narrowing the No-Arbitrage Bounds, Robert G. Chambers and John Quiggin (2003).

R03_4 Separability of stochastic production decisions from producer risk preferences in the presence of financial markets, Robert G Chambers and John Quiggin (2003).

R03_5 Comparative statics for state-contingent technologies, Robert G Chambers and John Quiggin (2003).

R03_6 Dual structures for the sole-proprietorship firm, Robert G Chambers and John Quiggin (December, 2003).

R04_1 Output price subsidies in a stochastic world, Robert G Chambers and John Quiggin (2004).

R04_2 Supermodularity and risk aversion, John Quiggin and Robert G Chambers (March 2004).

R04_3 Linear-risk-tolerant, invariant risk preferences, Robert G Chambers and John Quiggin (April, 2004).

R04_4 Increasing uncertainty: a definition, Simon Grant and John Quiggin (May 2004)

R04_5 Supermodularity and the comparative statics of risk, John Quiggin and Robert G Chambers (June, 2004).

R04_6 Fixed wages and bonuses in agency contracts: the case of a continuous state space, Maria Racionero and John Quiggin (June, 2004).

R04_7 Games without rules, Flavio Menezes and John Quiggin (July, 2004).

R04_8 The risk premium for equity: implications for resource allocation, welfare and policy, Simon Grant and John Quiggin (August, 2004).

R04_9 Capital market equilibrium with moral hazard and flexible technology, John Quiggin and Robert G Chambers (September, 2004).

R05_1 The state-contingent approach to production under uncertainty, John Quiggin and Robert G Chambers (January, 2005).

R05_2 Economists and uncertainty, John Quiggin (April, 2005).

R05_3 Cost minimization and asset pricing, Robert G Chambers and John Quiggin (2005).

R05_4 Consistent Bayesian updating with unconsidered propositions, Simon Grant and John Quiggin (February, 2005).

R05_5 Comparative risk aversion for state-dependent preferences, John Quiggin and Robert G Chambers (May, 2005).

R05_6 Outcomes and strategy choices in Tullock contests, Flavio Menezes and John Quiggin (March, 2005).

R05_7 Learning and Discovery, Simon Grant and John Quiggin (July 2005).

R06_1 Dual approaches to the analysis of risk aversion, Robert G Chambers and John Quiggin (June, 2006).

R06_2 Efficiency analysis in the presence of uncertainty, Chris O'Donnell, Robert G Chambers and John Quiggin (May, 2006).

R06_3 Lost in translation: honest misunderstandings and ex post disputes, Simon Grant, Jeff Kline and John Quiggin (August, 2006).

R07_1 Estimating complex production functions: The importance of starting values, Mark Neal (January, 2007).

R07_2 Risk and derivative price, Yusuke Osaki (2007).

R07_3 A risk-neutral characterization of optimism and pessimism, and its applications, Yusuke Osaki and John Quiggin (2007).

R07_4 Can game theory be saved? Flavio Menezes and John Quiggin (August, 2007).

R07_5 Bargaining power and efficiency in insurance contracts, John Quiggin and Robert G Chambers (April 2007).

R07_6 Sharp and diffuse incentives in contracting, Flavio Menezes and John Quiggin (2007).

R07_7 Markets for influence, Flavio Menezes and John Quiggin (July, 2007).

R07_8 Event-specific data envelopment models and efficiency analysis, Robert G. Chambers, Atakelty Hailu and John Quiggin (2007).

R08_1 Generalized invariant preferences: two-parameter representations of preferences by Robert G Chambers and John Quiggin (February, 2008).

R08_2 Bounded rationality and small worlds, Simon Grant and John Quiggin (June, 2008).

R09_1 Inductive reasoning about unawareness, Simon Grant and John Quiggin (June 2009).

R09_2 Markets for influence, Flavio M. Menezes and John Quiggin (September, 2009).

R09_3 A matter of interpretation: bargaining over ambiguous contracts, Simon Grant, Jeff Kline and John Quiggin (November, 2009).

R10_1 The computation of perfect and proper equilibrium for finite games via simulated annealing, Stuart MacDonald and Liam Wagner (April, 2010).

R10_2 Uncertainty and technical efficiency in Finnish agriculture: a state-contingent approach, Céline Nauges, Christopher O'Donnell and John Quiggin (2010).

R10_3 Production under uncertainty: a simulation study, Sriram Shankar, Chris O'Donnell and John Quiggin, (2010).

R10_4 Economics as a social science: financial regulation after the crisis, John Quiggin (2010).

R11_1 More competitors or more competition? Market concentration and the intensity of competition, Flavio M. Menezes and John Quiggin (August 2011).

R11_2 A two-parameter model of dispersion aversion, Robert G Chambers, Simon Grant, Ben Polak and John Quiggin (August 2011).

R11_3 Intensity of competition and the number of competitors, Flavio M Menezes and John Quiggin (June, 2011).

R11_4 A matter of interpretation: ambiguous contracts and liquidated damages, Simon Grant, Jeff Kline and John Quiggin (September, 2011).

R11_5 Capabilities as menus: A non-welfarist basis for QALY evaluation, Han Bleichrodt and John Quiggin (October, 2011).

R12_1 Differential awareness ambiguity, and incomplete contracts: a model of contractual disputes, Simon Grant, J Jude Kline and John Quiggin (December, 2012).

R12_2 Optimal access regulation with downstream competition, Tina Kao, Flavio Menezes and John Quiggin (March, 2012).

R12_3 Inductive reasoning about unawareness, Simon Grant and John Quiggin (April, 2012).

R12_4 Bounded awareness, heuristics and the precautionary principle, Simon Grant, John Quiggin (April, 2012).

R12_5 A matter of interpretation: ambiguous contracts and liquidated damages, Simon Grant, Jeff Kline and John Quiggin (May, 2012).

R12_6 Meeting the competition: commitment and competitive behaviour, Tina Kao, Flavio Menezes and John Quiggin (December 2012).

R13_1 Inferring the strategy space from market outcomes, Flavio Menezes and John Quiggin (January, 2013).

R13_2 The value of information and the value of awareness, John Quiggin (January, 2013).

R13_3 A theory of strategic interaction with purely subjective uncertainty, Simon Grant, Idione Meneghel and Rabee Tourky (June, 2013).

R13_4 Ordinal, nonlinear context-dependence, Patrick O’Callaghan (May, 2013).