

Geometry on the Utility Space

François Durand, Benoît Kloeckner, Fabien Mathieu, Ludovic Noirie

► **To cite this version:**

François Durand, Benoît Kloeckner, Fabien Mathieu, Ludovic Noirie. Geometry on the Utility Space. Fourth International Conference on Algorithmic Decision Theory, Sep 2015, Lexington, United States. pp.16, 2015, Fourth International Conference on Algorithmic Decision Theory. <10.1007/978-3-319-23114-3_12>. <hal-01222871>

HAL Id: hal-01222871

<https://hal.inria.fr/hal-01222871>

Submitted on 30 Oct 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Geometry on the Utility Space

François Durand¹, Benoît Kloeckner², Fabien Mathieu³, and Ludovic Noirie³

¹ Inria, Paris, France, francois.durand@inria.fr

² Université Paris-Est, Laboratoire d'Analyse et de Mathématiques Appliquées (UMR 8050), UPEM, UPEC, CNRS, F-94010, Créteil, France,

benoit.kloeckner@u-pec.fr

³ Alcatel-Lucent Bell Labs France, Nozay, France,
{fabien.mathieu,ludovic.noirie}@alcatel-lucent.com

Abstract. We study the geometrical properties of the utility space (the space of expected utilities over a finite set of options), which is commonly used to model the preferences of an agent in a situation of uncertainty. We focus on the case where the model is neutral with respect to the available options, i.e. treats them, *a priori*, as being symmetrical from one another. Specifically, we prove that the only Riemannian metric that respects the geometrical properties and the natural symmetries of the utility space is the round metric. This canonical metric allows to define a uniform probability over the utility space and to naturally generalize the Impartial Culture to a model with expected utilities.

Keywords: Utility theory, Geometry, Impartial culture, Voting

1 Introduction

Motivation. In the traditional literature of Arrowian social choice [2], especially voting theory, the preferences of an agent are often represented by ordinal information only⁴: a strict total order over the available options, or sometimes a binary relation of preference that may not be a strict total order (for example if indifference is allowed). However, it can be interesting for voting systems to consider cardinal preferences, for at least two reasons.

Firstly, some voting systems are not based on ordinal information only, like Approval voting or Range voting.

Secondly, voters can be in a situation of uncertainty, either because the rule of the voting system involves a part of randomness, or because each voter has incomplete information about other voters' preferences and the ballots they will choose. To express preferences in a situation of uncertainty, a classical and elegant model is the one of expected utilities [17,5,10,12]. The utility vector \vec{u} representing the preferences of an agent is an element of \mathbb{R}^m , where m is the number of available options or *candidates*; the utility of a lottery over the options is computed as an expected value.

⁴ This is not always the case: for example, Gibbard [7] considers voters with expected utilities over the candidates.

For a broad set of applications in economics, the options under consideration are financial rewards or quantities of one or several economic goods, which leads to an important consequence: there is a natural structure over the space of options. For example, if options are financial rewards, there is a natural order over the possible amounts and it is defined prior to the agents' preferences.

We consider here the opposite scenario where options are symmetrical *a priori*. This symmetry condition is especially relevant in voting theory, by a normative principle of neutrality, but it can apply to other contexts of choice under uncertainty when there is no natural preexistent structure on the space of available options.

The motivation for this paper came from the possible generalizations of the *Impartial Culture* to agents with expected utilities. The *Impartial Culture* is a classical probabilistic model in voting theory where each agent draws independently her strict total order of preference with a uniform probability over the set of possible orders.

The difficulty is not to define a probability law over utilities such that its projection on ordinal information is the *Impartial Culture*. Indeed, it is sufficient to define a distribution where voters are independent and all candidates are treated symmetrically. The issue is to choose one in particular: an infinity of distributions meet these conditions and we can wonder whether one of these generalizations has canonical reasons to be chosen rather than the others.

To answer this question, we need to address an important technical point. As we will see, an agent's utility vector is defined up to two constants, and choosing a specific normalization is arbitrary. As a consequence, the utility space is a quotient space of \mathbb{R}^m , and *a priori*, there is no canonical way to push a metric from \mathbb{R}^m to this quotient space. Hence, at first sight, it seems that there is no natural definition of a uniform distribution of probability over that space.

More generally, searching a natural generalization of the *Impartial Culture* to the utility space naturally leads to investigate different topics about the geometry of this quotient space and to get a better understanding of its properties related to algebra, topology and measure theory.

We emphasize that our goal is not to propose a measure that represents real-life preferences. Such approach would follow from observation and experimental studies rather than theoretical work. Instead, we aim at identifying a measure that can play the role of a neutral, reference, measure. This will give a model for uniformness to which real distributions of utilities can be compared. For example, if an observed distribution has higher density than the reference measure in certain regions of the utility space, this could be interpreted as a non-uniform culture for the population under study and give an indication of an overall trend for these regions. With this aim in mind, we will assume a symmetry hypothesis over the different candidates.

Contributions and plan. The rest of the paper is organized as follows. In Section 2, we quickly present Von Neumann–Morgenstern framework and define the utility space. In Section 3, we show that the utility space may be seen as a quotient of the dual of the space of pairs of lotteries over the candidates.

In Section 4, we naturally define an inversion operation, that corresponds to reversing preferences while keeping their intensities, and a summation operation, that is characterized by the fact that it preserves unanimous preferences.

Since the utility space is a manifold, it is a natural wish to endow it with a Riemannian metric. In Section 5, we prove that the only Riemannian representation that preserves the natural projective properties and the *a priori* symmetry between the candidates is the round metric. Finally, in Section 6, we use this result to give a canonical generalization of the Impartial culture and to suggest the use of Von Mises–Fisher model to represent polarized cultures.

2 Von Neumann–Morgenstern Model

In this section, we define some notations and we quickly recall Von Neumann–Morgenstern framework in order to define the utility space.

Let $m \in \mathbb{N} \setminus \{0\}$. We will consider m mutually exclusive options called *candidates*, each one represented by an index in $\{1, \dots, m\}$. A *lottery* over the candidates is an m -tuple $(L_1, \dots, L_m) \in (\mathbb{R}_+)^m$ such that (s.t.) $\sum_{j=1}^m L_j = 1$. The set of lotteries is denoted \mathcal{L}_m .

The preferences of an agent over lotteries are represented by a binary relation \leq over \mathcal{L}_m .

Von Neumann–Morgenstern theorem states that, provided that relation \leq meets quite natural assumptions⁵, it can be represented by a utility vector $\vec{u} = (u_1, \dots, u_m) \in \mathbb{R}^m$, in the sense that following the relation \leq is equivalent to maximizing the expected utility. Formally, for any two lotteries L and M :

$$L \leq M \Leftrightarrow \sum_{j=1}^m L_j u_j \leq \sum_{j=1}^m M_j u_j.$$

Mathematical definitions, assumptions and proof of this theorem can be found in [16,12,10], and discussions about the experimental validity of the assumptions are available in [6,12].

For the purposes of this paper, a crucial point is that \vec{u} is defined up to an additive constant and a positive multiplicative constant. Formally, for any two vectors \vec{u} and \vec{v} , let us note $\vec{u} \sim \vec{v}$ iff $\exists a \in (0, +\infty), \exists b \in \mathbb{R}$ s.t. $\vec{v} = a\vec{u} + b\vec{1}$, where $\vec{1}$ denotes the vector whose m coordinates are equal to 1. With this notation, if $\vec{u} \in \mathbb{R}^m$ is a utility vector representing \leq , then a vector $\vec{v} \in \mathbb{R}^m$ is also a utility vector representing \leq iff $\vec{u} \sim \vec{v}$.

In order to define the utility space, all vectors representing the same preferences must be identified as only one point. The *utility space* over m candidates, denoted \mathcal{U}_m , is defined as the quotient space \mathbb{R}^m / \sim . We call *canonical projection* from \mathbb{R}^m to \mathcal{U}_m the function:

$$\tilde{\pi} : \begin{cases} \mathbb{R}^m \rightarrow \mathcal{U}_m \\ \vec{u} \rightarrow \tilde{u} = \{\vec{v} \in \mathbb{R}^m \text{ s.t. } \vec{v} \sim \vec{u}\}. \end{cases}$$

For any $\vec{u} \in \mathbb{R}^m$, we denote without ambiguity $\leq_{\tilde{u}}$ the binary relation over \mathcal{L}_m represented by \tilde{u} .

⁵ The necessary and sufficient condition is that relation \leq is complete, transitive, archimedean and independent of irrelevant alternatives.





Fig. 1. Space \mathbb{R}^3 of utility vectors for 3 candidates (without quotient).

Figure 1 represents the space \mathbb{R}^3 of utility vectors for 3 candidates, without projecting to the quotient space. The canonical base of \mathbb{R}^3 is denoted $(\vec{e}_1, \vec{e}_2, \vec{e}_3)$. Utility vectors $\vec{u}^{(1)}$ to $\vec{u}^{(4)}$ represent the same preferences as any other vector of the half-plane \tilde{u} , represented in gray. More generally, each non-trivial point \tilde{u} in the quotient utility space corresponds to a half-plane in \mathbb{R}^m , delimited by the line $\text{vect}(\vec{1})$, the linear span of $\vec{1}$. The only exception is the point of total indifference $\tilde{0}$. In \mathbb{R}^m , it does not correspond to a plane but to the line $\text{vect}(\vec{1})$ itself.



(a) Edges of the unit cube in \mathbb{R}^3 .

(b) Circle in \mathbb{R}^3 .

Fig. 2. Two representations of \mathcal{U}_3 .

To deal with utilities, conceptually and practically, it would be convenient to have a canonical representative \vec{u} for each equivalence class \tilde{u} . In Figure 2(a), for each non-indifferent \tilde{u} , we choose its representative satisfying $\min(u_i) = -1$ and $\max(u_i) = 1$. The utility space \mathcal{U}_3 (except the indifference point) is represented in \mathbb{R}^3 by six edges of the unit cube. In Figure 2(b), we choose the representative satisfying $\sum u_i = 0$ and $\sum u_i^2 = 1$. In that case, $\mathcal{U}_3 \setminus \{\tilde{0}\}$ is represented in \mathbb{R}^3 by the unit circle in the linear plane that is orthogonal to $\vec{1}$.

If we choose such a representation, the quotient space \mathcal{U}_m can inherit the Euclidean distance from \mathbb{R}^m ; for example, we can evaluate distances along the edges of the cube in Figure 2(a), or along the unit circle in Figure 2(b). But it is clear that the result will depend on the representation chosen. Hence, it is an interesting question whether one of these two representations, or yet another one, has canonical reasons to be used. But before answering this question, we need to explore in more generality the geometrical properties of the utility space.

3 Duality with the Tangent Hyperplane of Lotteries

In this section, we remark that the utility space is a dual of the space of pairs of lotteries. Not only does it give a different point of view on the utility space (which we think is interesting by itself), but it will also be helpful to prove Theorem 3, which characterizes the summation operator that we will define in Section 4.

In the example represented in Figure 3, we consider $m = 3$ candidates and $\vec{u} = (\frac{5}{3}; -\frac{1}{3}; -\frac{4}{3})$. The great triangle, or simplex, is the space of lotteries \mathcal{L}_3 . Hatchings are the agent's indifference lines: she is indifferent between any pair of lotteries on the same line (see [12], Section 6.B). The utility vector \vec{u} represented here is in the plane of the simplex, but it is not mandatory: indeed, \vec{u} can be arbitrarily chosen in its equivalence class \tilde{u} . Nevertheless, it is a quite natural choice, because the component of \vec{u} in the direction $\vec{1}$ (orthogonal to the simplex) has no meaning in terms of preferences.




Fig. 3. Space \mathcal{L}_3 of lotteries over 3 candidates.

The utility vector \vec{u} can be seen as a gradient of preference: at each point, it reveals in what directions the agent can find lotteries she likes better. However, only the direction of \vec{u} is important, whereas its norm has no specific meaning. As a consequence, the utility space is not exactly a dual space, but rather a quotient of a dual space, as we will see more formally.

For any two lotteries $L = (L_1, \dots, L_m)$ and $M = (M_1, \dots, M_m)$, we call *bipoint* from L to M the vector $\vec{LM} = (M_1 - L_1, \dots, M_m - L_m)$. We call *tangent polytope* of \mathcal{L}_m the set \mathcal{T} of bipoints of \mathcal{L}_m .

We call *tangent hyperplane* of \mathcal{L}_m :

$$\mathcal{H} = \{(\Delta_1, \dots, \Delta_m) \in \mathbb{R}^m \text{ s.t. } \sum_{j=1}^m \Delta_j = 0\}.$$

In Figure 3, the tangent polytope \mathcal{T} is the set of the bipoints of the great triangle, seen as a part of a vector space (whereas \mathcal{L}_m is a part of an affine space). The tangent hyperplane \mathcal{H} is the whole linear hyperplane containing \mathcal{T} .

Let us note $\langle \vec{u} \mid \vec{v} \rangle$ the canonical inner product of \vec{u} and \vec{v} . We call *positive half-hyperplane associated to \vec{u}* the set $\vec{u}^+ = \{\vec{\Delta} \in \mathcal{H} \text{ s.t. } \langle \vec{u} \mid \vec{\Delta} \rangle \geq 0\}$. By definition, a lottery M is preferred to a lottery L iff the bipoint \overrightarrow{LM} belongs to this positive half-hyperplane:

$$L \leq_{\vec{u}} M \Leftrightarrow \langle \vec{u} \mid \overrightarrow{LM} \rangle \geq 0 \Leftrightarrow \overrightarrow{LM} \in \vec{u}^+.$$

Let \mathcal{H}^* be the dual space of \mathcal{H} , that is, the set of linear forms on \mathcal{H} . For any $\vec{u} \in \mathbb{R}^m$, we call *linear form associated to \vec{u}* the following element of \mathcal{H}^* :

$$\langle \vec{u} \mid \cdot \rangle : \begin{cases} \mathcal{H} \rightarrow \mathbb{R} \\ \vec{\Delta} \rightarrow \langle \vec{u} \mid \vec{\Delta} \rangle. \end{cases}$$

We observed that the utility vector can be seen as a gradient, except that only its direction matters, not its norm. Let us formalize this idea. For any $(f, g) \in (\mathcal{H}^*)^2$, we denote $f \sim g$ iff these two linear forms are positive multiples of each other, that is, iff $\exists a \in (0, +\infty)$ s.t. $g = af$. We denote $\tilde{\pi}(f) = \{g \in \mathcal{H}^* \text{ s.t. } f \sim g\}$: it is the equivalence class of f , up to positive multiplication.

Proposition 1 *For any $(\vec{u}, \vec{v}) \in (\mathbb{R}^m)^2$, we have:*

$$\vec{u} \sim \vec{v} \Leftrightarrow \langle \vec{u} \mid \cdot \rangle \sim \langle \vec{v} \mid \cdot \rangle.$$

The following application is a bijection:

$$\Theta : \begin{cases} \mathcal{U}_m \rightarrow \mathcal{H}^* / \sim \\ \tilde{\pi}(\vec{u}) \rightarrow \tilde{\pi}(\langle \vec{u} \mid \cdot \rangle). \end{cases}$$

Proof. $\vec{u} \sim \vec{v}$

$$\begin{aligned} &\Leftrightarrow \exists a \in (0, +\infty), \exists b \in \mathbb{R} \text{ s.t. } \vec{v} - a\vec{u} = b\vec{1} \\ &\Leftrightarrow \exists a \in (0, +\infty) \text{ s.t. } \vec{v} - a\vec{u} \text{ is orthogonal to } \mathcal{H} \\ &\Leftrightarrow \exists a \in (0, +\infty) \text{ s.t. } \forall \vec{\Delta} \in \mathcal{H}, \langle \vec{v} \mid \vec{\Delta} \rangle = a \langle \vec{u} \mid \vec{\Delta} \rangle \\ &\Leftrightarrow \langle \vec{u} \mid \cdot \rangle \sim \langle \vec{v} \mid \cdot \rangle. \end{aligned}$$

The implication \Rightarrow proves that Θ is correctly defined: indeed, if $\tilde{\pi}(\vec{u}) = \tilde{\pi}(\vec{v})$, then $\tilde{\pi}(\langle \vec{u} \mid \cdot \rangle) = \tilde{\pi}(\langle \vec{v} \mid \cdot \rangle)$. The implication \Leftarrow ensures that Θ is injective. Finally, Θ is obviously surjective.

Hence, the utility space can be seen as a quotient of the dual \mathcal{H}^* of the tangent space \mathcal{H} of the lotteries \mathcal{L}_m . As noticed before, a utility vector may be seen, up to a positive constant, as a uniform gradient, i.e. as a linear form over \mathcal{H} that reveals, for any point in the space of lotteries, in what directions the agent can find lotteries that she likes better.

4 Inversion and Summation Operators

As a quotient of \mathbb{R}^m , the utility space inherits natural operations from \mathbb{R}^m : inversion and summation. We will see that both these quotient operators have an

intuitive meaning regarding preferences. The summation will also allow to define *lines* in Section 5, which will be a key notion for Theorem 5, characterizing the suitable Riemannian metrics.

We define the *inversion* operator of \mathcal{U}_m as:

$$- : \begin{cases} \mathcal{U}_m & \rightarrow \mathcal{U}_m \\ \tilde{\pi}(\vec{u}) & \rightarrow \tilde{\pi}(-\vec{u}). \end{cases}$$

This operator is correctly defined and it is a bijection; indeed, $\tilde{\pi}(\vec{u}) = \tilde{\pi}(\vec{v})$ iff $\tilde{\pi}(-\vec{u}) = \tilde{\pi}(-\vec{v})$. Considering this additive inverse amounts to reverting the agent's preferences, without modifying their relative intensities.

Now, we want to push the summation operator from \mathbb{R}^m to the quotient \mathcal{U}_m . We use a generic method to push an operator to a quotient space: considering \tilde{u} and \tilde{v} in \mathcal{U}_m , their antecedents are taken in \mathbb{R}^m thanks to $\tilde{\pi}^{-1}$, the sum is computed in \mathbb{R}^m , then the result is converted back into the quotient space \mathcal{U}_m , thanks to $\tilde{\pi}$.

However, the result is not unique. Indeed, let us take arbitrary representatives $\vec{u} \in \tilde{u}$ and $\vec{v} \in \tilde{v}$. In order to compute the sum, we can think of any representatives. So, possible sums are $a\vec{u} + a'\vec{v} + (b+b')\vec{1}$, where a and a' are positive and where $b+b'$ is any real number. Converting back to the quotient, we can get for example $\tilde{\pi}(2\vec{u} + \vec{v})$ and $\tilde{\pi}(\vec{u} + 3\vec{v})$, which are generally distinct. As a consequence, the output is not a point in the utility space \mathcal{U}_m , but rather a set of points, i.e. an element of $\mathcal{P}(\mathcal{U}_m)$.

This example shows how we could define the sum of two elements \tilde{u} and \tilde{v} . In order to be more general, we will define the sum of any number of elements of \mathcal{U}_m . Hence we also take $\mathcal{P}(\mathcal{U}_m)$ as the set of inputs.

We define the *summation* operator as:

$$\Sigma : \begin{cases} \mathcal{P}(\mathcal{U}_m) & \rightarrow \mathcal{P}(\mathcal{U}_m) \\ A & \rightarrow \{ \tilde{\pi}(\sum_{i=1}^n \vec{u}_i), n \in \mathbb{N}, (\vec{u}_1, \dots, \vec{u}_n) \in (\tilde{\pi}^{-1}(A))^n \}. \end{cases}$$




Fig. 4. Sum of two utility vectors in the utility space \mathcal{U}_4 .

Example 2 Let us consider \mathcal{U}_4 , the utility space for 4 candidates. In Figure 4, for the purpose of visualization, we represent its projection in \mathcal{H} , which is permitted by the choice of normalization constants b . Since \mathcal{H} is a 3-dimensional space, let $(\vec{h}_1, \vec{h}_2, \vec{h}_3)$ be an orthonormal base of it.

For two non-trivial utility vectors \tilde{u} and \tilde{v} , the choice of normalization multipliers a allow to choose representatives \vec{u} and \vec{v} whose Euclidean norm is 1.

In this representation, the sum $\sum\{\tilde{u}, \tilde{v}\}$ consists of utilities corresponding to vectors $a\vec{u} + a'\vec{v}$, where a and a' are nonnegative. Indeed, we took a representation in \mathcal{H} , so all normalization constants b vanish. Moreover, a , a' or both can be equal to zero because our definition allows to ignore \vec{u} , \vec{v} or both. Up to taking representatives of unitary norm for non-trivial utility vectors, let us note that the sum $\sum\{\tilde{u}, \tilde{v}\}$ can be represented by the dotted line and the point $\vec{0}$ of total indifference.

Geometrically, the sum is the quotient of the convex hull of the inputs. Note that this convex hull is actually a convex cone. We will see below the interpretation of the sum in terms of preferences.

Due to our definition of the sum operator, we consider the closed cone: for example, the inputs themselves (e.g. \tilde{u}) fit in our definition, and so is the total indifference $\tilde{\pi}(\vec{0})$. That would not be the case if we took only $\tilde{\pi}(a\vec{u} + a'\vec{v} + b\vec{1})$, where $a > 0$ and $a' > 0$. The purpose of this convention is to have a concise wording for Theorem 3 that follows.

We now prove that, if A is the set of the utility vectors of a population, then $\sum A$ is the set of utility vectors that respect the unanimous preferences of the population.

Theorem 3 (characterization of the sum) Let $A \in \mathcal{P}(\mathcal{U}_m)$ and $\tilde{v} \in \mathcal{U}_m$.

The following conditions are equivalent.

1. $\tilde{v} \in \sum A$.
2. $\forall (L, M) \in \mathcal{L}_m^2: [(\forall \tilde{u} \in A, L \leq_{\tilde{u}} M) \Rightarrow L \leq_{\tilde{v}} M]$.

Proof. First, let us remark that the tangent polytope \mathcal{T} generates the tangent hyperplane \mathcal{H} by positive multiplication. That is:

$$\forall \vec{\Delta} \in \mathcal{H}, \exists \vec{LM} \in \mathcal{T}, \exists \lambda \in (0, +\infty) \text{ s.t. } \vec{\Delta} = \lambda \vec{LM}.$$

Indeed, \mathcal{T} contains a neighborhood of the origin in vector space \mathcal{H} .

Let $\vec{v} \in \tilde{\pi}^{-1}(\tilde{v})$. We have the following equivalences.

- $\forall (L, M) \in \mathcal{L}_m^2, (\forall \tilde{u} \in A, L \leq_{\tilde{u}} M) \Rightarrow L \leq_{\tilde{v}} M,$
- $\forall \vec{LM} \in \mathcal{T}, \left(\forall \vec{u} \in \tilde{\pi}^{-1}(A), \langle \vec{u} | \vec{LM} \rangle \geq 0 \right) \Rightarrow \langle \vec{v} | \vec{LM} \rangle \geq 0,$
- $\forall \vec{\Delta} \in \mathcal{H}, \left(\forall \vec{u} \in \tilde{\pi}^{-1}(A), \langle \vec{u} | \vec{\Delta} \rangle \geq 0 \right) \Rightarrow \langle \vec{v} | \vec{\Delta} \rangle \geq 0$ (because \mathcal{T} generates \mathcal{H}),
- $\bigcap_{\vec{u} \in \tilde{\pi}^{-1}(A)} \vec{u}^+ \subset \vec{v}^+,$

- \vec{v} is in the convex cone of $\tilde{\pi}^{-1}(A)$ (because of the duality seen in Section 3),
- $\tilde{v} \in \sum A$.

Example 4 *Let us consider a non-indifferent \tilde{u} and let us examine the sum of \tilde{u} and its additive inverse $-\tilde{u}$. By a direct application of the definition, the sum consists of \tilde{u} , $-\tilde{u}$ and $\tilde{0}$.*

However, we have just proved that the sum is the subset of utility vectors preserving the unanimous preferences over lotteries. Intuitively, we could think that, since \tilde{u} and $-\tilde{u}$ seem to always disagree, any utility vector \tilde{v} respects the empty set of their common preferences; so, their sum should be the whole space. But this is not a correct intuition.

Indeed, let us examine the example of $\vec{u} = (1, 0, \dots, 0)$. About any two lotteries L and M , inverse opinions \tilde{u} and $-\tilde{u}$ agree iff $L_1 = M_1$: in that case, both \tilde{u} and $-\tilde{u}$ are indifferent between L and M . The only points of the utility space meeting this common property are \tilde{u} and $-\tilde{u}$ themselves, as well as the indifference point $\tilde{0}$.

5 Riemannian Representation of the Utility Space

Since the utility space is a manifold, it is a natural desire to endow it with a Riemannian metric. In this section, we prove that there is a limited choice of metrics that are coherent with the natural properties of the space and with the *a priori* symmetry between the candidates.

First, let us note that the indifference point $\tilde{0}$ must be excluded. Indeed, its unique open neighborhood is \mathcal{U}_m in whole, and no distance is consistent with this property⁶. In contrast, $\mathcal{U}_m \setminus \{\tilde{0}\}$ has the same topology as a sphere of dimension $m - 2$, so it can be endowed with a distance.

Now, let us define the round metric. The quotient $\mathbb{R}^m / \text{vect}(\vec{1})$ is identified to \mathcal{H} and endowed with the inner product inherited from the canonical one of \mathbb{R}^m . The utility space $\mathcal{U}_m \setminus \{\tilde{0}\}$ is identified to the unit sphere of \mathcal{H} and endowed with the induced Riemannian structure. We note ξ_0 this Riemannian metric on $\mathcal{U}_m \setminus \{\tilde{0}\}$.

In order to get an intuitive vision of this metric, one can represent any position \tilde{u} by a vector \vec{u} that verifies $\sum u_i = 0$ and $\sum u_i^2 = 1$. We obtain the $(m - 2)$ -dimensional unit sphere of \mathcal{H} and we consider the metric induced by the canonical Euclidean metric of \mathbb{R}^m . That is, distances are measured on the surface of the sphere, using the restriction of the canonical inner product on each tangent space. For $m = 3$, such a representation has already been illustrated in Figure 2(b).

With this representation in mind, we can give a formula to compute distances with ξ_0 . We denote J the $m \times m$ matrix whose all coefficients are equal to 1 and $P_{\mathcal{H}}$ the matrix of the orthogonal projection on \mathcal{H} :

$$P_{\mathcal{H}} = \text{Id} - \frac{1}{m}J.$$

⁶ Technically, this remark proves that \mathcal{U}_m (with its natural quotient topology) is not a T_1 space [8].

The canonical Euclidean norm of \vec{u} is denoted $\|\vec{u}\|$. For two non-indifferent utility vectors \vec{u} and \vec{v} , the distance between \tilde{u} and \tilde{v} in the sense of metric ξ_0 is:

$$d(\tilde{u}, \tilde{v}) = \arccos \left\langle \frac{P_{\mathcal{H}} \vec{u}}{\|P_{\mathcal{H}} \vec{u}\|} \mid \frac{P_{\mathcal{H}} \vec{v}}{\|P_{\mathcal{H}} \vec{v}\|} \right\rangle.$$

If \vec{u} and \vec{v} are already unit vectors of \mathcal{H} , i.e. canonical representatives of their equivalence classes \tilde{u} and \tilde{v} , then the formula amounts to $d(\tilde{u}, \tilde{v}) = \arccos \langle \vec{u} \mid \vec{v} \rangle$.

A natural property for the distance would be that its geodesics coincide with the unanimity segments defined by the sum. Indeed, imagine that Betty with utilities \tilde{u}_B is successfully trying to convince Alice with initial utilities \tilde{u}_A to change her mind. Assume that Alice evolves continuously, so that her utility follows a path from \tilde{u}_A to \tilde{u}_B . Unless Betty argues in a quite convoluted way, it makes sense to assume on one hand that Alice's path is a shortest path, and on the other hand that whenever Alice and Betty initially agree on the comparison of two lotteries, this agreement continues to hold all along the path. This is precisely what assumption **1a** below means: shortest paths preserve agreements that hold at both their endpoints.

We will now prove that for $m \geq 4$, the spherical representation is the only one that is coherent with the natural properties of the space and that respects the *a priori* symmetry between candidates.

Theorem 5 (Riemannian representation of the utility space) *We assume that $m \geq 4$. Let ξ be a Riemannian metric on $\mathcal{U}_m \setminus \{\tilde{0}\}$.*

Conditions 1 and 2 are equivalent.

1. (a) *For any non-antipodal pair of points $\tilde{u}, \tilde{v} \in \mathcal{U}_m \setminus \{\tilde{0}\}$ (i.e. $\tilde{v} \neq -\tilde{u}$), the set $\sum\{\tilde{u}, \tilde{v}\}$ of elements respecting the unanimous preferences of \tilde{u} and \tilde{v} is a segment of a geodesic of ξ ; and*
- (b) *for any permutation σ of $\{1, \dots, m\}$, the action Φ_σ induced on $\mathcal{U}_m \setminus \{\tilde{0}\}$ by*

$$(u_1, \dots, u_m) \rightarrow (u_{\sigma(1)}, \dots, u_{\sigma(m)})$$

is an isometry.

2. $\exists \lambda \in (0, +\infty)$ s.t. $\xi = \lambda \xi_0$.

Proof. Since the implication **2** \Rightarrow **1** is obvious, we now prove **1** \Rightarrow **2**. The deep result behind this is a classical theorem of Beltrami, which dates back to the middle of the nineteenth century: see [3] and [4].

The image of a 2-dimensional subspace of \mathcal{H} in $\mathcal{U}_m \setminus \{\tilde{0}\}$ by the canonical projection $\tilde{\pi}$ is called a *line* in $\mathcal{U}_m \setminus \{\tilde{0}\}$. This notion is deeply connected to the summation operator: indeed, the sum of two non antipodal points in $\mathcal{U}_m \setminus \{\tilde{0}\}$ is a segment of the line joining them. Condition **1a** precisely means that the geodesics of ξ are the lines of $\mathcal{U}_m \setminus \{\tilde{0}\}$. Beltrami's theorem then states that $\mathcal{U}_m \setminus \{\tilde{0}\}$ has constant curvature. Note that this result is in fact more subtle in dimension 2 (that is, for $m = 4$) than in higher dimensions; see [13], Theorem 1.18 and [14], Theorem 7.2 for proofs.

Since $\mathcal{U}_m \setminus \{\tilde{0}\}$ is a topological sphere, this constant curvature must be positive. Up to multiplying ξ by a constant, we can assume that this constant curvature is 1. As a consequence, there is an isometry $\Psi : \mathcal{S}_{m-2} \rightarrow \mathcal{U}_m \setminus \{\tilde{0}\}$, where \mathcal{S}_{m-2} is the unit sphere of \mathbb{R}^{m-1} , endowed with its usual round metric. The function Ψ obviously maps geodesic to geodesics, let us deduce the following.

Lemma 6 *There is a linear map $\Lambda : \mathbb{R}^{m-1} \rightarrow \mathcal{H}$ inducing Ψ , that is such that:*

$$\Psi \circ \Pi = \Pi \circ \Lambda,$$

where Π denotes both projections $\mathbb{R}^{m-1} \rightarrow \mathcal{S}_{m-2}$ and $\mathcal{H} \rightarrow \mathcal{U}_m \setminus \{\tilde{0}\}$.

Proof. First, Ψ maps any pair of antipodal points of \mathcal{S}_{m-2} to a pair of antipodal points of $\mathcal{U}_m \setminus \{\tilde{0}\}$: indeed in both cases antipodal pairs are characterized by the fact that there are more than one geodesic containing them. It follows that Ψ induces a map Ψ' from the projective space $P(\mathbb{R}^{m-1})$ (which is the set of lines through the origin in \mathbb{R}^{m-1} , identified with the set of pairs of antipodal points of \mathcal{S}_{m-2}) to the projective space $P(\mathcal{H})$ (which is the set of lines through the origin in \mathcal{H} , identified with the set of pairs of antipodal points of $\mathcal{U}_m \setminus \{\tilde{0}\}$).

The fact that Ψ sends geodesics of \mathcal{S}_{m-2} to geodesics of $\mathcal{U}_m \setminus \{\tilde{0}\}$ and condition **1a** together imply that Ψ' sends projective lines to projective lines.

As is well known, a one-to-one map $\mathbb{R}^n \rightarrow \mathbb{R}^n$ which sends lines to lines must be an affine map; a similar result holds in projective geometry, concluding that Ψ' must be a projective map. See e.g. [1] for both results.

That Ψ' is projective exactly means that there is a linear map $\Lambda : \mathbb{R}^{m-1} \rightarrow \mathcal{H}$ inducing Ψ' , which then induces $\Psi : \mathcal{S}_{m-2} \rightarrow \mathcal{U}_m \setminus \{\tilde{0}\}$.

Using Λ to push the canonical inner product of \mathbb{R}^{m-1} , we get that there exists an inner product $(\vec{u}, \vec{v}) \rightarrow \phi(\vec{u}, \vec{v})$ on \mathcal{H} that induces ξ , in the sense that ξ is the Riemannian metric obtained by identifying $\mathcal{U}_m \setminus \{\tilde{0}\}$ with the unit sphere defined in \mathcal{H} by ϕ and restricting ϕ to it.

The last thing to prove is that ϕ is the inner product coming from the canonical one on \mathbb{R}^m . Note that hypothesis **1b** is mandatory, since any inner product on \mathcal{H} does induce on $\mathcal{U}_m \setminus \{\tilde{0}\}$ a Riemannian metric satisfying **1a**.

Each canonical basis vector $\vec{e}_j = (0, \dots, 1, \dots, 0)$ defines a point in $\mathcal{U}_m \setminus \{\tilde{0}\}$ and a half-line $\ell_j = \mathbb{R}_+ e_j$ in \mathcal{H} . Condition **1b** ensures that these half-lines are permuted by some isometries of (\mathcal{H}, ϕ) . In particular, there are vectors $\vec{u}_j \in \ell_j$ that have constant pairwise distance (for ϕ).

Lemma 7 *The family $\vec{u}_1, \dots, \vec{u}_{m-1}$ is, up to multiplication by a scalar, the unique basis of \mathcal{H} such that $\vec{u}_j \in \ell_j$ and $\sum_{j < m} \vec{u}_j \in -\ell_m$.*

Proof. First, by definition of the \vec{u}_j , these vectors form a regular simplex and $\sum_j \vec{u}_j = \vec{0}$. It follows that $\vec{u}_1, \dots, \vec{u}_{m-1}$ has the required property and we have left to show uniqueness.

Assume $\vec{v}_1, \dots, \vec{v}_{m-1}$ is a basis of \mathcal{H} such that $\vec{v}_j \in \ell_j$ and $\sum_{j < m} \vec{v}_j \in -\ell_m$. Then there are positive scalars a_1, \dots, a_m such that $\vec{v}_j = a_j \vec{u}_j$ for all $j < m$, and $\sum_{j < m} \vec{v}_j = a_m \sum_{j < m} \vec{u}_j$.

Then $\sum_{j < m} a_j \vec{u}_j = \sum_{j < m} a_m \vec{u}_j$, and since the u_j form a basis, it must hold $a_j = a_m$ for all j .

Now consider the canonical inner product ϕ_0 on \mathcal{H} that comes from the canonical one on \mathbb{R}^m . Since permutations of coordinates are isometries, we get that the vectors $\vec{v}_j = \Pi(\vec{e}_j)$ (where Π is now the orthogonal projection from \mathbb{R}^m to \mathcal{H}) form a regular simplex for ϕ_0 , so that $\sum_j \vec{v}_j = \vec{0}$. It follows that $\vec{u}_j = \lambda \vec{v}_j$ for some $\lambda > 0$ and all j . We deduce that the \vec{u}_j form a regular simplex for both ϕ and ϕ_0 , which must therefore be multiple from each other. This finishes the proof of Theorem 5.

However, the implication **1** \Rightarrow **2** in the theorem is not true for $m = 3$. For each non-indifferent utility vector, let us consider its representative verifying $\min(u_i) = 0$ and $\max(u_i) = 1$. This way, $\mathcal{U}_m \setminus \{\tilde{0}\}$ is identified to edges of the unit cube in \mathbb{R}^3 , as in Figure 2(a). We use this identification to endow $\mathcal{U}_m \setminus \{\tilde{0}\}$ with the metric induced on these edges by the canonical inner product on \mathbb{R}^3 . Then conditions **1a** and **1b** of the theorem are met, but not condition **2**.

Another remark is of paramount importance about this theorem. Since we have a canonical representative \vec{u} for each equivalence class \tilde{u} as a unit vector of \mathcal{H} , we could be tempted to use it to compare utilities between several agents.

We stress on the fact that this representation cannot be used for interpersonal comparison of utility levels or utility differences.

For example, for two agents, consider the following representatives:

$$\begin{cases} \vec{u} = (0.00, 0.71, -0.71), \\ \vec{v} = (0.57, 0.22, -0.79). \end{cases}$$

The fact that $v_3 < u_3$ does not mean that an agent with preferences \tilde{v} dislikes candidate 3 more than an agent with preferences \tilde{u} . Similarly, when changing from candidate 1 to candidate 2, the gain for agent \tilde{u} (+0.71) cannot be compared to the loss for agent \tilde{v} (-0.35).

Theorem 5 conveys no message for interpersonal comparison of utilities. Indeed, utilities belonging to two agents are essentially incomparable in the absence of additional assumptions [9]. Taking canonical representatives on the $(m - 2)$ -dimensional sphere is only used to compute distances between two points in the utility space.

6 Application: Probability Measures on the Utility Space

Once the space is endowed with a metric, it is endowed with a natural probability measure: the uniform measure in the sense of this metric (this is possible because the space has a finite total measure). We will denote μ_0 this measure, which is thus the normalized Riemannian volume defined by the metric ξ_0 .

In practice, to draw vectors according to a uniform probability law over $\mathcal{U}_m \setminus \{\tilde{0}\}$, it is sufficient to use a uniform law on the unit sphere in \mathcal{H} . In other words, once one identifies $\mathcal{U}_m \setminus \{\tilde{0}\}$ with the unit sphere in \mathcal{H} , then μ_0 is exactly the usual uniform measure.

In the present case, the fact that the round sphere has many symmetries implies additional nice qualities for μ_0 , which we sum up in a proposition.

Proposition 8 *Let μ be any probability measure on $\mathcal{U}_m \setminus \{\tilde{0}\}$.*

1. *Assume that for all $\delta > 0$, μ gives the same probability to all the balls in $\mathcal{U}_m \setminus \{\tilde{0}\}$ of radius δ (in the metric ξ_0); then $\mu = \mu_0$.*
2. *If μ is invariant under all isometries for the metric ξ_0 , then $\mu = \mu_0$.*

Both characterizations are classical; the first one is (a strengthening of) the definition of the Riemannian volume, and the second one follows from the first and the fact that any two points on the round sphere can be mapped one to the other by an isometry.




Fig. 5. Two distributions of 100 agents on \mathcal{U}_4 .

In Figure 5(a), we represent a distribution with 100 agents drawn uniformly and independently on the sphere, with 4 candidates. Like for Figure 4, we represented only the unit sphere of \mathcal{H} .

The solid dark lines draw the *permutohedron*, a geometrical figure representing the ordinal aspect of these preferences. Each *facet* is constituted of all the points who share the same strict order of preference. A utility vector belongs to an *edge* if it has only three distinct utilities: for example, if the agent prefers candidate 1 to 4, 4 to 2 and 3, but is indifferent between candidates 2 and 3. Finally, a point is a *vertex* if it has only two distinct utilities: for example, if the agent prefers candidate 1 to all the others, but is indifferent between the others.

In this distribution, each agent has almost surely a strict order of preference. Each order has the same probability and agents are independent, hence this

distribution is a natural generalization of the Impartial Culture when considering expected utilities.

Since the point $\tilde{0}$ is a geometrical singularity, it is difficult to include it naturally in such a measure. If one wants to take it into account, the easiest way is to draw it with a given probability and to use the distribution on $\mathcal{U}_m \setminus \{\tilde{0}\}$ in the other cases. That being said, we have just noticed that all other non-strict orders have a measure equal to 0 ; so for a canonical theoretical model, it can be argued that a natural choice is to attribute a measure 0 to the indifference point also.

Having a distance, hence a uniform measure, allows also to define other measures by their densities with respect to the uniform measure. Here is an example of a law defined by its density. Given a vector \vec{u}_0 in the unit sphere of \mathcal{H} and κ a nonnegative real number, the Von Mises–Fisher (VMF) distribution of pole \vec{u}_0 and concentration κ is defined by the following density with respect to the uniform law on the unit sphere in \mathcal{H} :

$$p(\vec{u}) = C_\kappa e^{\kappa \langle \vec{u} | \vec{u}_0 \rangle},$$

where C_κ is a normalization constant. Given the mean resultant vector of a distribution over the sphere, VMF distribution maximizes the entropy, in the same way that, in the Euclidean space, Gaussian distribution maximizes the entropy among laws with given mean and standard deviation. Hence, without additional information, it is the “natural” distribution that should be used. Figure 5(b) represents such a distribution, with the same conventions as Figure 5(a). To draw a VMF distribution, we used Ulrich’s algorithm revised by Wood [15,18].

Qualitatively, VMF model is similar to Mallows model, which is used for ordinal preferences [11]. In the later, the probability of an order of preference σ is:

$$p(\sigma) = D_\kappa e^{-\kappa \delta(\sigma, \sigma_0)},$$

where σ_0 is the mode of the distribution, $\delta(\sigma, \sigma_0)$ a distance between σ and σ_0 (typically Kendall’s tau distance), κ a nonnegative real number (concentration) and D_κ a normalization constant. Both VMF and Mallows models describe a culture where the population is polarized, i.e. scattered around a central point, with more or less concentration.

However, there are several differences.

- VMF distribution allows to draw a specific point on the utility sphere, whereas Mallows’ chooses only a facet of the permutohedron.
- In particular, the pole of a VMF distribution can be anywhere in this continuum. For example, even if its pole is in the facet $1 > 4 > 3 > 2$, it can be closer to the facet $1 > 4 > \mathbf{2} > \mathbf{3}$ than to the facet $\mathbf{4} > \mathbf{1} > 3 > 2$. Such a nuance is not possible in Mallows model.
- In the neighborhood of the pole, VMF probability decreases like the exponential of the square of the distance (because the inner product is the cosine of the distance), whereas Mallows probability decreases like the exponential of the distance (not squared).

- VMF is the maximum entropy distribution, given a spherical mean and dispersion (similarly to a Gaussian distribution in a Euclidean space), whereas Mallows' model is not characterized by such a property of maximum entropy.

The existence of a canonical measure allows to define other probability measures easily in addition to the two we have just described. Such measures can generate artificial populations of agents for simulation purposes. They can also be used to fit the data from real-life experiments to a theoretical model, and as a neutral comparison point for such data.

To elaborate on this last point, let us stress that given a (reasonably regular) distribution μ on a space such as $\mathcal{U}_m \setminus \{\tilde{0}\}$ there is *a priori* no way to define what it means for an element \tilde{u} to be more probable than another one \tilde{v} . Indeed, both have probability 0 and what would make sense is to compare the probability of being close to \tilde{u} to the probability of being close to \tilde{v} . But one should compare neighborhoods of the same size, and one needs a metric to make this comparison. Alternatively, if one has a reference distribution such as μ_0 , then it makes sense to consider the density $f = \frac{d\mu}{d\mu_0}$, which is a (continuous, say) function on $\mathcal{U}_m \setminus \{\tilde{0}\}$. Then we can compare $f(\tilde{u})$ and $f(\tilde{v})$ to say whether one of these elements is more probable than the other according to μ . Note that in the present case, comparing the probability of δ -neighborhoods for the metric ξ_0 or the density relative to μ_0 gives the same result in the limit $\delta \rightarrow 0$, which is the very definition of the Riemannian volume.

7 Conclusion

We have studied the geometrical properties of the classical model of expected utilities, introduced by Von Neumann and Morgenstern, when candidates are considered symmetrical *a priori*. We have remarked that the utility space may be seen as a dual of the space of lotteries, that inversion and summation operators inherited from \mathbb{R}^m have a natural interpretation in terms of preferences and that the space has a spherical topology when the indifference point is removed.

We have proved that the only Riemannian representation whose geodesics coincide with the projective lines naturally defined by the summation operator and which respects the symmetry between candidates is a round sphere.

All these considerations lay on the principle to add as little information as possible in the system, especially by respecting the *a priori* symmetry between candidates. This does not imply that the spherical representation of the utility space \mathcal{U}_m is the most relevant one in order to study a specific situation. Indeed, as soon as one has additional information (for example, a model that places candidates in a political spectrum), it is natural to include it in the model. However, if one wishes, for example, to study a voting system in all generality, without focusing on its application in a specific field, it looks natural to consider a utility space with a metric as neutral as possible, like the one defined in this paper by the spherical representation.

Acknowledgements. The work presented in this paper has been partially carried out at LINCS (<http://www.lincs.fr>).

References

1. M. Audin. *Geometry*. Universitext, Springer, 2003.
2. Kenneth J. Arrow. A difficulty in the concept of social welfare. *The Journal of Political Economy*, 58(4):pp. 328–346, 1950.
3. E. Beltrami. Résolution du problème de reporter les points d’une surface sur un plan, de manière que les lignes géodésiques soient représentée par des lignes droites. *Annali di Matematica*, 1866.
4. E. Beltrami. Essai d’interprétation de la géométrie noneuclidéenne. Trad. par J. Hoüel. *Ann. Sci. École Norm. Sup.*, 6:251–288, 1869.
5. P.C. Fishburn. *Utility Theory for Decision Making*. Wiley, New York, 1970.
6. P.C. Fishburn. *Nonlinear preference and utility theory*. Johns Hopkins series in the mathematical sciences. Johns Hopkins University Press, 1988.
7. Allan Gibbard. Straightforwardness of Game Forms with Lotteries as Outcomes. *Econometrica*, 46(3):595–614, 1978.
8. F. Guénard and G. Lelièvre. *Compléments d’analyse*. Number vol. 1 in Compléments d’analyse. E.N.S., 1985.
9. Peter J. Hammond. Interpersonal comparisons of utility: Why and how they are and should be made. In *Interpersonal Comparisons of Well-Being*, pages 200–254. University Press, 1991.
10. David M. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990.
11. Colin L Mallows. Non-null ranking models. i. *Biometrika*, pages 114–130, 1957.
12. Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, jun 1995.
13. M. Spivak. *A comprehensive introduction to differential geometry. Vol. III*. Publish or Perish Inc., second edition, 1979.
14. M. Spivak. *A comprehensive introduction to differential geometry. Vol. IV*. Publish or Perish Inc., second edition, 1979.
15. G. Ulrich. Computer generation of distributions on the m-sphere. *Applied Statistics*, pages 158–163, 1984.
16. J. Von Neumann, O. Morgenstern, H.W. Kuhn, and A. Rubinstein. *Theory of Games and Economic Behavior (Commemorative Edition)*. Princeton Classic Editions. Princeton University Press, 2007.
17. John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1944.
18. Andrew T.A. Wood. Simulation of the von mises fisher distribution. *Communications in Statistics-Simulation and Computation*, 23(1):157–164, 1994.