Crash Course on Expected Utility Theory

Lasha Chochua

ISET, 2026

Table of contents

Note: A standard text on expected utility theory is Mas-Colell, Whinston, and Green (1995), Chapter 6; for a decision-theoretic treatment see Gilboa (2009), Chapters 1–3. The measurement-theoretic perspective adopted here follows Sider (2011).

1 The Problem of Representing Risky Choice

Standard consumer theory assigns utility numbers to *certain* outcomes and asks only that the numbers respect a ranking. Most economic decisions, however, involve *risk*: the agent chooses among probability distributions over outcomes rather than among outcomes directly.

Two questions immediately arise. First, can preferences over risky prospects be represented by a utility function at all? Second, if so, what do the numbers *mean* – are they ordinal labels, or do they carry cardinal information?

The answer, due to von Neumann and Morgenstern (1944), is that under four axioms on preferences there exists a utility function u on outcomes such that lottery L is preferred to lottery L' if and only if the *expected value* of u under L exceeds the expected value under L'. Moreover, u is unique up to a positive affine transformation – it lives on an *interval scale*. The present note proves this result carefully, highlighting at each step which axiom is doing the work.

2 Measurement Theory: Background

Before stating the axioms it is useful to recall what it means to represent a non-numerical structure with numbers. This perspective, developed by Krantz et al. (1971) and applied to quantity by Sider (2011), gives precise content to phrases like "utility is ordinal" or "utility differences are meaningful."

A *relational structure* is a pair $\langle A, R_1, \dots, R_n \rangle$ where A is a set and each R_i is a relation on A.

Let $\langle A, R_1, \dots, R_n \rangle$ and $\langle B, S_1, \dots, S_n \rangle$ be two relational structures. A function $f\colon A \to B$ is a *homomorphism* if it preserves every relation:

$$R_i(x_1, \dots, x_m) \iff S_i(f(x_1), \dots, f(x_m))$$

for every R_i and all $x_1, \dots, x_m \in A$.

Intuitively, f is a "structure-respecting translation." The mathematical structure $\langle B, S_1, \dots \rangle$ is useful for representing $\langle A, R_1, \dots \rangle$ precisely because a homomorphism exists – the math mirrors the non-numerical facts.

The *uniqueness* question asks: given that a homomorphism exists, how many are there? The answer determines the *scale type*:

Scale type	Preserves	Allowed transformations
Ordinal	Order only	Any monotone $\tilde{f} = g \circ f$
Interval	Ratios of intervals	Affine: $\tilde{f} = kf + \beta$, $k > 0$
Ratio	Ratios	Similarity: $\tilde{f} = kf$, $k > 0$

Standard consumer utility is an *ordinal* representation of preferences over certain outcomes: the structure is $\langle X, \succsim \rangle$ and we find a homomorphism into $\langle \mathbb{R}, \geq \rangle$. Any monotone transformation of u works equally well. The statement "bundle A gives twice the utility of bundle B" is meaningless – it is not invariant across valid representations.

The vNM theorem will establish that utility over lotteries is an *interval* representation. Utility *differences* and their ratios will be meaningful; utility *levels* and their ratios will not.

3 Setup

Outcomes. Let $X = \{x_1, \dots, x_n\}$ be a finite set of outcomes.

Lotteries. $\mathcal{L}$ denotes the set of all probability distributions over X. Write $L = (p_1, \dots, p_n)$ where $p_i \geq 0$ and $\sum_i p_i = 1$. The *degenerate lottery* δ_{x_i} gives outcome x_i with certainty.

Mixing. For any $L, L' \in \mathcal{L}$ and $\alpha \in [0, 1]$, the *mixture* $\alpha L + (1-\alpha)L'$ is the lottery that assigns probability $\alpha p_i + (1-\alpha)q_i$ to outcome x_i. Operationally: flip a coin with probability α of heads; play L if heads, L' if tails.

Primitive. A binary relation $\succsim$ on $\mathcal{L}$. Write $L \succ L'$ if $L \succsim L'$ but not $L' \succsim L$, and $L \sim L'$ if both hold.

Goal. Find $u \colon X \to \mathbb{R}$ such that

$$L \succsim L' \iff \sum_i p_i \, u(x_i) \; \geq \; \sum_i q_i \, u(x_i), \tag{1}$$

and characterise the degree of uniqueness of u.

4 The Axioms

Completeness. For all $L, L' \in \mathcal{L}$: $L \succsim L'$ or $L' \succsim L$ (or both).

Transitivity. If $L \succsim L'$ and $L' \succsim L''$, then $L \succsim L''$.

Continuity. If $L \succ L' \succ L''$, then there exist $\alpha, \beta \in (0, 1)$ such that

$$\alpha L + (1-\alpha)L'' \; \succ \; L' \; \succ \; \beta L + (1-\beta)L''.$$

Independence. For all $L, L', L'' \in \mathcal{L}$ and all $\alpha \in (0, 1)$:

$$L \succsim L' \iff \alpha L + (1-\alpha)L'' \; \succsim \; \alpha L' + (1-\alpha)L''.$$

A few remarks on what each axiom is and is not doing.

Completeness and Transitivity are the standard *weak order* conditions. They appear in any representation theorem – they are the minimum required to rank objects at all. Without completeness there are incomparable lotteries; without transitivity there may be preference cycles. Neither axiom says anything about how preferences interact with the mixing operation.

Continuity is sometimes called the *Archimedean axiom*. It rules out "infinitely good" or "infinitely bad" outcomes – no outcome is so extreme that you would not trade it for some mixture. Technically it rules out lexicographic preferences, which cannot be represented by a real-valued function at all.

Independence is the critical axiom. It says that if you mix both L and L' with the same lottery L'', your ranking of the mixture should follow your ranking of the components. The lottery L'' is a "common factor" and should be irrelevant. As we will see, Independence is the sole axiom with real bite for the scale type of the representation: it is what forces utility to be linear in probabilities, and therefore what elevates the scale from ordinal to interval.

5 Proof of the Representation Theorem

von Neumann–Morgenstern (1944). A preference relation $\succsim$ on $\mathcal{L}$ satisfies Axioms 1–4 if and only if there exists $u\colon X \to \mathbb{R}$ such that (1) holds. Moreover, u is unique up to positive affine transformation: if $\tilde{u}$ also represents $\succsim$ in this way then $\tilde{u} = ku + \beta$ for some $k > 0$, $\beta \in \mathbb{R}$.

We prove the "if" direction (axioms imply representation) in six steps, then establish uniqueness as a seventh step.

5.1 Step 1 – Best and worst lottery

Claim. There exist $\bar{L}, \underline{L} \in \mathcal{L}$ such that $\bar{L} \succsim L \succsim \underline{L}$ for all $L \in \mathcal{L}$.

Proof. Since X is finite, there are finitely many degenerate lotteries $\delta_{x_1}, \dots, \delta_{x_n}$. By Completeness we can compare any two of them; by Transitivity we can rank all of them. Set $\bar{L} = \delta_{x^*}$ (best) and $\underline{L} = \delta_{x_*}$ (worst). Repeated application of Independence extends the ranking to all lotteries in $\mathcal{L}$.

What Completeness and Transitivity are doing. These two axioms guarantee a complete, transitive ranking among the finite set of degenerate lotteries. Without them no "best outcome" can be identified and the entire construction collapses at the first step.

Assume $\bar{L} \succ \underline{L}$ throughout (the trivial case $\bar{L} \sim \underline{L}$ makes all lotteries indifferent and any constant u works).

5.2 Step 2 – The utility index – Continuity does the work

Claim. For every $L \in \mathcal{L}$ there exists a *unique* $\alpha_L \in [0, 1]$ such that

$$L \;\sim\; \alpha_L \bar{L} + (1 - \alpha_L)\underline{L}. \tag{2}$$

Proof. Define

$$A = \{\alpha \in [0, 1] : \alpha\bar{L} + (1 - \alpha)\underline{L} \succsim L\}, \quad B = \{\alpha \in [0, 1] : L \succsim \alpha\bar{L} + (1 - \alpha)\underline{L}\}.$$

From Step 1: $1 \in A$ (since $\bar{L} \succsim L$) and $0 \in B$ (since $L \succsim \underline{L}$).

Closedness. Continuity implies that preferences vary continuously in the mixing weight α: if $\alpha_n \to \alpha$ and $\alpha_n\bar{L} + (1 - \alpha_n)\underline{L} \succsim L$ for all n, then the same holds at the limit. Hence both A and B are closed.

Existence. By Completeness, $A \cup B = [0, 1]$. Since $[0, 1]$ is connected and both sets are closed, $A \cap B \neq \emptyset$. Any $\alpha \in A \cap B$ satisfies (2).

Uniqueness. Suppose $\alpha > \beta$ both satisfy (2). Then by Independence, $\alpha \bar{L} + (1-\alpha)\underline{L} \succ \beta \bar{L} + (1-\beta)\underline{L}$ (since $\bar{L} \succ \underline{L}$ and a higher weight on the better lottery is strictly preferred). But both are indifferent to L, contradicting Transitivity.

What Continuity is doing. Continuity is entirely responsible for Step 2. It ensures preferences over mixtures vary continuously in α, so the intermediate-value argument can be applied. Without it, a lottery could be better than every mixture $\alpha \bar{L} + (1-\alpha)\underline{L}$ for $\alpha < 1$ yet worse than $\bar{L}$ itself – the indifference point would not exist.

Definition. Set $U(L) \equiv \alpha_L$. This assigns a number in $[0, 1]$ to every lottery.

5.3 Step 3 – U represents $\succsim$

Claim. $L \succsim L' \iff U(L) \geq U(L')$.

Proof. By (2) and Transitivity:

$$L \succsim L' \iff U(L)\bar{L} + (1-U(L))\underline{L} \;\succsim\; U(L')\bar{L} + (1-U(L'))\underline{L}.$$

We are comparing two mixtures of $\bar{L}$ and $\underline{L}$. Since $\bar{L} \succ \underline{L}$, Independence implies a higher weight on $\bar{L}$ is strictly preferred, so the right-hand comparison reduces to $U(L) \geq U(L')$.

What the axioms are doing. Transitivity allows substitution of each lottery for its benchmark equivalent. Independence then makes the comparison of two mixtures of $\bar{L}$ and $\underline{L}$ reduce purely to comparing scalar weights.

5.4 Step 4 – Linearity in probabilities – Independence does the work

Claim. For any $L, L' \in \mathcal{L}$ and $\alpha \in [0, 1]$:

$$U(\alpha L + (1-\alpha)L') = \alpha\, U(L) + (1-\alpha)\, U(L'). \tag{3}$$

Proof. Write $p = U(L)$ and $q = U(L')$, so $L \sim p\bar{L} + (1-p)\underline{L}$ and $L' \sim q\bar{L} + (1-q)\underline{L}$.

First application of Independence. Mix both sides of the first indifference with L' at weight α:

$$\alpha L + (1-\alpha)L' \;\sim\; \alpha[p\bar{L} + (1-p)\underline{L}] + (1-\alpha)L'.$$

Second application of Independence. Substitute $L' \sim q\bar{L} + (1-q)\underline{L}$ into the right side:

$$\alpha[p\bar{L} + (1-p)\underline{L}] + (1-\alpha)L' \;\sim\; \alpha[p\bar{L} + (1-p)\underline{L}] + (1-\alpha)[q\bar{L} + (1-q)\underline{L}].$$

Chain by Transitivity and collect terms on $\bar{L}$ and $\underline{L}$:

$$\alpha L + (1-\alpha)L' \;\sim\; [\alpha p + (1-\alpha)q]\bar{L} + [1 - \alpha p - (1-\alpha)q]\underline{L}.$$

By definition of U this mixture has utility index $\alpha p + (1-\alpha)q$, establishing (3). □

What Independence is doing structurally. Independence allows us to substitute an equivalent lottery *inside* a mixture: we replaced L by $p\bar{L} + (1-p)\underline{L}$ (its benchmark equivalent) without disturbing the preference over the compound lottery. This substitution rule is precisely what forces linearity. Without Independence, knowing $L \sim p\bar{L} + (1-p)\underline{L}$ would tell us nothing about how L behaves once mixed with another lottery.

The structural parallel with mass measurement (Sider 2011) is exact. Mass has a physical *concatenation* relation $C(x, y, z)$: "z is the object obtained by combining x and y." The concatenation relation gives the mass structure enough richness to admit a ratio-scale representation. Probability mixing $\alpha L + (1-\alpha)L'$ plays the same role for lotteries. The Independence axiom says preferences must respect this operation linearly – giving the lottery structure enough richness for an interval-scale representation.

5.5 Step 5 – The expected utility formula

Claim. For any $L = (p_1, \dots, p_n)$:

$$U(L) = \sum_{i=1}^{n} p_i \, u(x_i),$$

where $u(x_i) \equiv U(\delta_{x_i})$.

Proof. Write $L = \sum_{i=1}^{n} p_i \, \delta_{x_i}$ and apply (3) by induction on n:

$$U(L) = U\Big(\sum_{i=1}^{n} p_i \, \delta_{x_i}\Big) = \sum_{i=1}^{n} p_i \, U(\delta_{x_i}) = \sum_{i=1}^{n} p_i \, u(x_i). \qquad \square$$

The representation (1) is now established. The function $u\colon X \to \mathbb{R}$ is entirely determined by how it ranks pure outcomes; the utility of any lottery is the probability-weighted average.

5.6 Step 6 – Uniqueness – the interval scale

Claim. If $\tilde{u}$ also represents $\succsim$ via (1), then $\tilde{u} = ku + \beta$ for some $k > 0$, $\beta \in \mathbb{R}$.

Proof. Let $\tilde{U}(L) = \sum_i p_i \, \tilde{u}(x_i)$. Apply $\tilde{U}$ to the indifference (2) and use linearity of $\tilde{U}$:

$$\tilde{U}(L) = U(L) \cdot \tilde{U}(\bar{L}) + (1 - U(L)) \cdot \tilde{U}(\underline{L}).$$

Rearranging:

$$\tilde{U}(L) = \tilde{U}(\underline{L}) + [\tilde{U}(\bar{L}) - \tilde{U}(\underline{L})] \cdot U(L).$$

Set $k = \tilde{U}(\bar{L}) - \tilde{U}(\underline{L}) > 0$ and $\beta = \tilde{U}(\underline{L})$. Then $\tilde{U}(L) = k\,U(L) + \beta$ for every L, so $\tilde{u} = ku + \beta$. $\square$

Why interval and not ratio. The two free parameters are $\tilde{U}(\bar{L})$ and $\tilde{U}(\underline{L})$ – the utility values assigned to the best and worst outcomes. Nothing in preferences pins these down. Two free parameters correspond to the constants k and β of an affine transformation – the signature of an *interval* scale. A ratio scale would require only one free parameter (k, with $\beta = 0$ fixed), which demands a meaningful utility zero. There is no decision-theoretic candidate for such a zero.

6 What Each Axiom Contributes

Axiom	Role in the proof
Completeness	Ensures $A \cup B = [0, 1]$ in Step 2; every lottery is comparable to the benchmark mixtures.
Transitivity	Chains indifferences in Steps 3 and 4; allows substitution of equivalent lotteries via benchmark equivalences.
Continuity	Guarantees existence of the indifference point α_L in Step 2 via the intermediate-value argument; rules out lexicographic preferences.
Independence	(i) Reduces the comparison of two benchmark mixtures to comparing scalar weights (Step 3); (ii) forces U to be linear in probabilities (Step 4), elevating the scale from ordinal to interval.

Completeness, Transitivity, and Continuity alone yield an ordinal representation. Independence is the only axiom with real bite for the *scale type*: it is what makes utility linear in probabilities and pins the representation down to affine transformations.

7 What is Meaningful Under vNM

Because u is an interval scale, only statements invariant to positive affine transformations are meaningful.

Statement	Meaningful?	Reason
$L \succsim L'$	Yes	Invariant to any affine transform
$u(x) > u(y)$	Yes	Order preserved
$u(x) - u(y) > u(y) - u(z)$	Yes	Ratios of differences invariant
$u(x) = 2\,u(y)$	No	Not invariant to change of zero
$u_A(x) > u_B(x)$ (interpersonal)	No	Scales are incomparable

Two implications are worth noting. *Risk aversion* is meaningful: the curvature of u (i.e., $u'' < 0$) is invariant to positive affine transformations, so whether an agent is risk-averse, risk-neutral, or risk-loving is a well-defined property of preferences. *Interpersonal utility comparisons* are not meaningful: adding utilities across individuals requires ratio-scale comparability, which vNM does not deliver.

8 References

- Gilboa, I. (2009). *Theory of Decision Under Uncertainty.* Cambridge University Press.
- Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement.* Academic Press.

- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- Sider, T. (2011). "Crash course on quantities and measurement theory." Lecture notes, Fall 2011.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.