

Preferences to Rewards

Fernando E. Rosas

September 8, 2026

Video: <https://www.youtube.com/watch?v=Zv0EC4-iRjQ>

What you'll learn

- Understand the difference between preferences, utility, and reward: preferences being a primary, largely uncontroversial notion, and utility and rewards being derived notions resting on specific assumptions.
- Be able to derive the relationship between various preference structures and rationality axioms.
- Critically assess alternative notions of rationality, and the consequences of dropping various classical decision theory assumptions.

Overview

This note develops a short route from preferences over complete trajectories to expected utility, reward, and discount. We begin with preference relations on deterministic trajectories and explain how completeness and transitivity yield an ordinal utility representation. We then show how lotteries, together with the von Neumann–Morgenstern axioms, produce a cardinal utility over trajectories, and we clarify the distinction between ordinal preference utility and vNM utility. Next, following Bowling et al., we add a fifth temporal axiom that is necessary and sufficient for a recursive representation in terms of local rewards and discounting. Finally, we explain why reward is not unique: different reward functions can encode the same utility or the same preference ordering, affine changes of utility induce corresponding changes of reward, and potential-based shaping provides a canonical example of reward equivalence.

1 Introduction

What does it mean for a system to have a goal? In sequential decision making, the object of evaluation is often not a single isolated choice or prize but a whole trajectory: a complete history of states, observations, actions, and consequences unfolding through

time. Before asking how goals are achieved, it is worth asking a more basic question: what does it mean to prefer some trajectories over others, and what follows from that?

Recent work in reinforcement learning has revived this fundamental question, treating preferences over histories as primitive and asking what additional assumptions are needed before one can recover scalar reward functions. The present note provides an introduction to these ideas, which proceeds in three stages.

- (1) First, one asks for a numerical representation of how whole trajectories are ranked.
- (2) Second, once one allows lotteries over trajectories, one asks when these lotteries can be ranked by the expectation of that trajectory utility.
- (3) Third, one asks what additional requirements are needed in order to decompose this expected utility into rewards assigned at each time step.

The first and second steps are closely related to the classical theory developed by von Neumann and Morgenstern ([von Neumann and Morgenstern, 1944](#)). The third asks what extra temporal structure is needed before utility over whole trajectories can be decomposed into stagewise rewards, following the line of work developed in modern reinforcement learning by [Pitis \(2019\)](#), [Shakerinava and Ravanbakhsh \(2022\)](#), and [Bowling et al. \(2023\)](#).

2 Preferences over trajectories

Let $\mathcal{O}$ be a finite set of observations and $\mathcal{A}$ a finite set of actions. A one-step interaction is then given by $t = (o, a) \in \mathcal{O} \times \mathcal{A}$. For each $n \in \mathbb{N}_{\geq 0}$, define the space of trajectories of length n by

$$\mathcal{H}_n := (\mathcal{O} \times \mathcal{A})^n.$$

We write ε for the unique trajectory of length 0. The space of all *finite* trajectories is

$$\mathcal{H}^* := \bigcup_{n=0}^{\infty} \mathcal{H}_n.$$

A typical element of $\mathcal{H}^*$ has the form $h = (o_1, a_1, o_2, a_2, \dots, o_n, a_n)$. We will keep the notation $\mathcal{H}^*$ for the set of all finite trajectories throughout.

Definition 2.1 (Preference). A preference relation on $\mathcal{H}^*$ is a binary relation $\succsim$ where

$$h \succsim h'$$

means that trajectory h is judged at least as good as trajectory h' .

From $\succsim$ we derive the usual companion relations:

$$h \sim h' \iff h \succsim h' \text{ and } h' \succsim h, \quad h \succ h' \iff h \succsim h' \text{ and not } h' \succsim h.$$

We call $\sim$ ‘indifference’, as an agent has no reason to prefer one over the other.

At this point, $\succsim$ has no properties whatsoever. One may naturally wonder what kinds of properties it is reasonable to require of $\succsim$, and what follows from them — which is what we study in the next sections.

3 When are preferences problematic?

Can a preference relation be intrinsically ‘bad’? The relevant concern here is whether it leads to some form of self-defeat, avoidable loss, or failure of coherent behaviour. The discussion in the decision-theory literature suggests at least three grades of concern.

1. *Representability failures.* A first and weakest concern is that a preference relation may fail to be representable in a convenient form. This does not, by itself, imply that the preference is irrational. Representation failures may merely be inconvenient, but they become more significant when they are symptoms of deeper issues of the kinds described next (Aumann, 1962; Fishburn, 1970).
2. *Self-defeat and avoidable loss.* A more serious concern is that preferences may guide choice poorly — as judged by the agent’s own interest. One important case is *static self-defeat*: choosing an option that is worse than another available one, or adopting a policy that is systematically improvable. A classic case of suboptimality is dominance: one option or policy dominates another when it is at least as good in every relevant respect and strictly better in some, so choosing the dominated option is a clear mistake (Kreps, 1988; Mas-Colell et al., 1995). A second case is *diachronic self-defeat*: a plan that the agent endorses now is predictably undone later in a way that leaves the agent worse off overall.
3. *Vulnerability.* The most vivid coherence arguments show that a collection of individually acceptable choices can be combined into a guaranteed loss. Dutch-book arguments play this role for credences; money-pump arguments play the analogous role for preferences. The standard example is a preference cycle

$$h_1 \succ h_2, \quad h_2 \succ h_3, \quad h_3 \succ h_1.$$

If the agent is willing to pay a small fee to move each time to a strictly preferred trajectory, then an adversary can guide it around the cycle and back to where it started, poorer than before (Gustafsson, 2010). This is why intransitivity is usually regarded as a particularly severe pathology: it is not merely hard to represent, but also vulnerable to exploitation under natural trading assumptions.

Note

Coherence is not selection.

It is useful to distinguish three kinds of formal results. A *representation result* says that if preferences satisfy certain axioms, then they can be written in a particular

mathematical form; the vNM and Savage theorems being classical examples. A *coherence result* is different: it links violations of some constraint to a penalty such as a Dutch book, money pump, dynamic inconsistency, or dominated choice. Complete-class and admissibility theorems belong more naturally in this second family than in the first, since they characterize undominated decision rules rather than utility representations. A *selection result* is different again: it adds a story about some optimization process — such as market competition, evolution, or training dynamics — and argues that systems lacking a certain property tend to be selected against. In short, representation concerns *form*, coherence concerns *vulnerability or domination*, and selection concerns *survival under pressure*.

Note that representation, coherence, and selection arguments are related, but not identical. In general, a coherence argument is a within-agent claim: it says that if a single agent violates some structural constraint, then the agent is vulnerable to a penalty such as a money pump, Dutch book, or dominated choice. A selection argument is different: it asks whether agents lacking that property would tend to disappear under some broader optimization pressure, such as market competition, training dynamics, or evolutionary selection. A coherence result may help motivate a selection story, yet it does not by itself show that realistic environments actually select against the offending preference pattern. Similarly, a representational failure with no plausible selection story may still be mathematically interesting while being less central for explaining the structure of real agents.

For the purposes of this note, the key point is that these notions of badness do not all coincide. A preference can fail standard representation without being exploitably incoherent, and not every departure from expected utility is thereby problematic. Nevertheless, the axioms studied below are useful because — as we will see — they rule out several undesirable properties.

4 Completeness, transitivity, and ordinal utility

Two structural conditions are especially important.

Definition 4.1 (Completeness). A preference relation $\succsim$ on $\mathcal{H}^*$ is *complete* if for every pair $h, h' \in \mathcal{H}^*$,

$$h \succsim h' \quad \text{or} \quad h' \succsim h \quad (\text{or both}).$$

Definition 4.2 (Transitivity). A preference relation $\succsim$ on $\mathcal{H}^*$ is *transitive* if for every $h, h', h'' \in \mathcal{H}^*$,

$$h \succsim h' \quad \text{and} \quad h' \succsim h'' \quad \implies \quad h \succsim h''.$$

Completeness says that the agent can compare any two trajectories. This is a claim about *comparability*, not about certainty: it says that the preference relation returns a verdict on every pair, not that the agent knows everything about the consequences of those trajectories. Transitivity says that these verdicts fit together consistently across chains of comparison. For example, if $h_1 \succsim h_2$ and $h_2 \succsim h_3$, then transitivity requires $h_1 \succsim h_3$ as well.

A preference relation satisfying both conditions is often called a *weak order* in economics and decision theory (Fishburn, 1970; Kreps, 1988; Mas-Colell et al., 1995). On a countable domain such as $\mathcal{H}^*$, weak orders admit an ordinal utility representation. More general representation theorems on richer domains go back to the classic work of Debreu (1954).

Proposition 4.3 (Ordinal representation on the trajectory space). *A preference relation $\succsim$ on $\mathcal{H}^*$ is complete and transitive if and only if there exists a function $u: \mathcal{H}^* \rightarrow \mathbb{R}$ such that*

$$h \succsim h' \iff u(h) \geq u(h') \quad \text{for all } h, h' \in \mathcal{H}^*.$$

Proof. If such a function u exists, then completeness and transitivity are inherited from the total order on $\mathbb{R}$. Conversely, if $\succsim$ is complete and transitive, then trajectories can be grouped into indifference classes, where each class contains all trajectories tied with one another. The quotient $\mathcal{H}^*/\sim$ of these classes is a countable total order. Any countable total order can be embedded in $\mathbb{R}$, so we may assign real numbers to the indifference classes in a way that preserves their order. Composing that assignment with the quotient map gives the desired utility function on trajectories. $\square$

This utility is *ordinal*: only the ranking matters. Any strictly increasing transformation of u represents the same preference relation. So ordinal utility lets us encode the order of deterministic trajectories, but it does not yet tell us how to compare risky mixtures of them. The numbers themselves have no independent meaning beyond the

order they induce. If u represents a preference relation, then so does $2u + 17$, or $\exp(u)$, provided the transformation remains strictly increasing. This is why ordinal utility is best thought of as a numerical *labeling of ranks*, not yet as a measure of how much one trajectory is preferred to another (Fishburn, 1970).

It is also useful to understand what happens when either condition fails.

What if completeness or transitivity fail

If completeness fails. Then some pairs of trajectories are incomparable. This may be a feature rather than a bug: the agent may genuinely refuse to rank certain alternatives because its values are plural, under-specified, or context-sensitive. But once incomparability is allowed, no single real-valued utility function can exactly represent the relation in the sense of Proposition 4.3, because any two real numbers are themselves comparable. One must then move to a different formalism, such as partial orders, sets of utility functions, or multi-criteria representations (Aumann, 1962).

If transitivity fails. Then local pairwise judgments need not assemble into a global ranking. In the simplest case one obtains a cycle

$$h_1 \succ h_2, \quad h_2 \succ h_3, \quad h_3 \succ h_1.$$

No scalar utility can represent such a cycle, since it would require

$$u(h_1) > u(h_2) > u(h_3) > u(h_1),$$

which is impossible. The logical problem is already serious: there is no single global ranking of the options. Under the additional behavioral assumption that the agent is willing to pay a small fee to move from any option to a strictly preferred one, this logical failure becomes operational through the *money pump* problem. Suppose an agent currently has h_1 and is willing to pay a small fee $\varepsilon > 0$ each time it moves to a strictly preferred trajectory. The cycle above licenses the sequence

$$h_1 \rightarrow h_2 \rightarrow h_3 \rightarrow h_1,$$

with the agent paying ε at each step. At the end it is back where it started, but poorer by 3ε . Repeating the cycle pumps away arbitrarily much money (Gustafsson, 2010).

There is also an interesting geometric view on transitivity. This paragraph is not needed for the rest of the note, so readers meeting these ideas for the first time can safely treat it as an optional aside. Fix a finite menu $V = \{h_1, \dots, h_m\} \subset \mathcal{H}^*$ and draw the complete graph whose vertices are the trajectories in V . Encode pairwise comparisons by an antisymmetric edge flow X :

$$X_{ij} = -X_{ji},$$

where $X_{ij} > 0$ means that h_i is preferred to h_j , while $X_{ij} < 0$ means the reverse. If preferences come from a utility function u , then each edge weight is just a utility difference,

$$X_{ij} = u(h_i) - u(h_j),$$

so X is a discrete gradient field. In the language of the discrete Helmholtz–Hodge decomposition, any edge flow can be split into a gradient part, which comes from a potential, and a cyclic part, which records genuine loops. On the complete comparison graph there is no extra harmonic remainder, so inconsistency is entirely captured by the cyclic component (Jiang et al., 2011). A three-cycle corresponds exactly to a nonzero discrete curl:

$$X_{ij} + X_{jk} + X_{ki} \neq 0.$$

Thus transitive preferences are precisely the *potential* part of the decomposition, with the potential given by the utility, while intransitive cycles show up as the rotational or cyclic part. If this language feels abstract, the key takeaway is simple: utility means all local comparisons come from one global ranking, whereas cycles are the leftover pattern that cannot be explained by any single scalar potential.

5 Lotteries and the von Neumann–Morgenstern axioms

Up to this point, we have just considered preferences over a countable set $\mathcal{H}^*$ without considering stochasticity. To introduce uncertainty, we expand the domain from trajectories to lotteries over trajectories.

Let $\Delta(X)$ for $X \subset \mathcal{H}^*$ denote the set of probability distributions over X . An element $\mu \in \Delta(X)$ can be expressed as

$$\mu = \sum_i p_i \delta_{x_i},$$

which should be read as a lottery that yields outcome x_i with probability p_i . The outcome x is identified with the degenerate lottery δ_x . In this way, preferences over outcomes can be viewed as a special case of preferences over lotteries. Furthermore, for $\mu, \nu \in \Delta(X)$ and $\lambda \in [0, 1]$, write

$$\lambda\mu + (1 - \lambda)\nu$$

for the compound lottery that first flips a coin with bias λ , then samples from μ or ν accordingly.

The von Neumann–Morgenstern (vNM) framework studies weak preference relations $\succsim$ over $\Delta(X)$ and asks when such preferences admit an expected-utility representation (von Neumann and Morgenstern, 1944). The framework considers four axioms on preferences over lotteries.

Axiom 1 (Completeness). For all $\mu, \nu \in \Delta(X)$, either $\mu \succsim \nu$ or $\nu \succsim \mu$ (or both).

Axiom 2 (Transitivity). For all $\mu, \nu, \xi \in \Delta(X)$, if $\mu \succsim \nu$ and $\nu \succsim \xi$, then $\mu \succsim \xi$.

Axiom 3 (Continuity). For all $\mu, \nu, \xi \in \Delta(X)$ with $\mu \succ \nu \succ \xi$, there exists $\lambda \in [0, 1]$ such that

$$\nu \sim \lambda\mu + (1 - \lambda)\xi.$$

Axiom 4 (Independence). For all $\mu, \nu, \xi \in \Delta(X)$ and $\lambda \in (0, 1)$,

$$\mu \succsim \nu \iff \lambda\mu + (1 - \lambda)\xi \succsim \lambda\nu + (1 - \lambda)\xi.$$

We already know about completeness and transitivity from the previous section; the new players are continuity and independence. Continuity says intermediate prospects admit a break-even mixture between better and worse ones. Independence says that if μ is preferred to ν , then mixing both with the same background lottery ξ should not reverse that preference.

The natural question is therefore when a preference over lotteries can be represented by the expectation of a utility function on trajectories:

$$U(\mu) = \sum_{x \in X} \mu(x)u(x).$$

This formulation says that the value of a lottery is the probability-weighted average of the utilities of its possible trajectories. In its simplest finite-outcome form, this question is answered by the celebrated vNM theorem.

Theorem 5.1 (von Neumann–Morgenstern). *Let $X \subseteq \mathcal{H}^*$ be finite. A weak preference relation $\succsim$ on $\Delta(X)$ satisfies completeness, transitivity, continuity, and independence if and only if there exists a function $u: X \rightarrow \mathbb{R}$ such that for all lotteries $\mu, \nu \in \Delta(X)$,*

$$\mu \succsim \nu \iff \sum_{x \in X} \mu(x)u(x) \geq \sum_{x \in X} \nu(x)u(x).$$

Moreover, u is unique up to positive affine transformations:

$$u'(x) = a + bu(x), \quad b > 0.$$

Proof. We first prove the easy direction: if preferences admit an expected-utility representation, then they satisfy the four axioms.

Completeness and transitivity follow immediately from the total order on $\mathbb{R}$. Continuity follows because if $\mu \succ \nu \succ \xi$ then $U(\mu) > U(\nu) > U(\xi)$, so one can choose $\lambda \in [0, 1]$ such that

$$U(\nu) = \lambda U(\mu) + (1 - \lambda)U(\xi).$$

Hence

$$\nu \sim \lambda\mu + (1 - \lambda)\xi.$$

Finally, independence follows from linearity:

$$U(\lambda\mu + (1 - \lambda)\xi) = \lambda U(\mu) + (1 - \lambda)U(\xi).$$

Therefore

$$\mu \succ \nu \iff \lambda\mu + (1 - \lambda)\xi \succ \lambda\nu + (1 - \lambda)\xi.$$

We now prove the converse direction, namely that the four axioms imply an expected-utility representation.

1. *Pick best and worst outcomes (completeness and transitivity).* Because X is finite and the restriction of $\succsim$ to degenerate lotteries is complete and transitive, either all outcomes in X are indifferent or there exist $b, w \in X$ such that

$$b \succsim x \succsim w \quad \text{for all } x \in X.$$

If all outcomes are indifferent, then the constant utility function already represents the preference relation. So we may assume $b \succ w$.

2. *Calibrate each outcome against b and w (continuity).* For $x = b$ set $u(x) = 1$, and for $x = w$ set $u(x) = 0$. For $b \succ x \succ w$, continuity gives some $u(x) \in (0, 1)$ such that

$$x \sim u(x)b + (1 - u(x))w.$$

Thus every outcome is indifferent to a lottery over the two benchmark outcomes b and w .

3. *Compare the benchmark lotteries (independence and transitivity).* We first show that for $x, x' \in X$,

$$x \succsim x' \iff u(x) \geq u(x').$$

For $\lambda \in [0, 1]$, write

$$C_\lambda := \lambda b + (1 - \lambda)w.$$

We claim that if $\alpha > \beta$, then $C_\alpha \succ C_\beta$. Indeed, if $\alpha = 1$, then $C_\alpha = b$, and independence applied to $b \succ w$ with mixing weight β and background lottery b gives

$$b \succ \beta b + (1 - \beta)w = C_\beta.$$

So it remains to consider the case $\alpha < 1$. Define

$$\rho := \frac{\alpha - \beta}{1 - \beta} \in (0, 1).$$

Since $b \succ w$, independence gives

$$\rho b + (1 - \rho)C_\beta \succ \rho w + (1 - \rho)C_\beta.$$

But the left-hand side is exactly C_α , while the right-hand side is C_β . Hence larger values of λ yield strictly better benchmark lotteries. Together with Step 2 and transitivity, this proves the displayed equivalence above.

4. *Reduce an arbitrary lottery to a benchmark lottery (independence).* Let

$$\mu = \sum_i p_i x_i.$$

By Step 2, each x_i is indifferent to $C_{u(x_i)}$. Replacing each x_i by the corresponding $C_{u(x_i)}$ inside μ , one at a time, and using independence together with transitivity, yields

$$\mu \sim \sum_i p_i C_{u(x_i)}.$$

By ordinary probability arithmetic, the compound lottery on the right reduces to

$$\sum_i p_i C_{u(x_i)} \sim U(\mu)b + (1 - U(\mu))w, \quad U(\mu) := \sum_i p_i u(x_i).$$

5. *Conclude the expected-utility representation.* Applying Step 3 to the benchmark lotteries from Step 4, we obtain

$$\mu \succsim \nu \iff U(\mu) \geq U(\nu).$$

This is exactly the desired expected-utility representation.

It remains to prove affine uniqueness, which uses continuity, independence, and the representation just obtained. Suppose u' is another expected-utility representation of the same preference relation. Let

$$a := u'(w), \quad c := u'(b) - u'(w) > 0.$$

For any $x \in X$, Step 2 gives

$$x \sim u(x)b + (1 - u(x))w.$$

Since u' also represents the same preferences, indifference implies equality of expected u' -value:

$$u'(x) = u(x)u'(b) + (1 - u(x))u'(w) = a + cu(x).$$

So u' is a positive affine transformation of u . □

Exercise 5.1 (Guided proof of the von Neumann–Morgenstern theorem).

Assume $X \subseteq \mathcal{H}^*$ is finite and that $\succsim$ on $\Delta(X)$ satisfies completeness, transitivity, continuity, and independence.

- (a) Show that there exist trajectories $b, w \in X$ such that $b \succ h \succ w$ for every $h \in X$.
- (b) Using continuity, show that for every $h \in X$ there exists $u(h) \in [0, 1]$ such that

$$h \sim u(h)b + (1 - u(h))w.$$

- (c) Use transitivity to show that if

$$h \sim u(h)b + (1 - u(h))w \quad \text{and} \quad h' \sim u(h')b + (1 - u(h'))w,$$

then

$$h \succ h' \iff u(h) \geq u(h').$$

(d) Let

$$\mu = \sum_i p_i h_i.$$

Use independence repeatedly to replace each h_i by the equivalent lottery $u(h_i)b + (1 - u(h_i))w$ and show that

$$\mu \sim \sum_i p_i (u(h_i)b + (1 - u(h_i))w).$$

(e) Reduce the compound lottery in part (d) to a simple lottery and prove that

$$\mu \sim U(\mu)b + (1 - U(\mu))w, \quad U(\mu) := \sum_i p_i u(h_i).$$

(f) Show that

$$\mu \succsim \nu \iff U(\mu) \geq U(\nu).$$

This gives the expected-utility representation.

(g) Prove affine uniqueness: if u' also represents the same preference relation in expected-utility form, then $u' = a + bu$ for some $a \in \mathbb{R}$ and $b > 0$.

(h) Prove the converse direction: if preferences admit an expected-utility representation, then they satisfy completeness, transitivity, continuity, and independence.

In contrast to our previous result, vNM utility is *cardinal* up to positive affine transformations, not merely ordinal. A general monotone transformation would destroy the expectation formula. The independence axiom forces exactly the amount of structure needed for linear averaging.

Note

Two notions of utility.

In economics it is standard to distinguish two different notions of utility (Kreps, 1988; Mas-Colell et al., 1995).

The first is *ordinal utility*, also known as *preference utility* or ‘F1’, which represents preferences over certain outcomes. This is the object obtained in Proposition 4.3: if preferences over deterministic trajectories are complete and transitive, then there exists a function u_{ord} such that

$$h \succsim h' \iff u_{\text{ord}}(h) \geq u_{\text{ord}}(h').$$

Only the ranking matters. Any strictly increasing transformation of u_{ord} represents the same preferences (Fishburn, 1970).

The second is *von Neumann–Morgenstern utility*, also called ‘F2’. This is the function u_{vNM} that appears inside the expectation operator when preferences over lotteries satisfy continuity and independence:

$$\mu \succsim \nu \iff \sum_h \mu(h) u_{\text{vNM}}(h) \geq \sum_h \nu(h) u_{\text{vNM}}(h).$$

Unlike ordinal utility, u_{vNM} is unique only up to positive affine transformations. Its numerical differences therefore carry behavioral content: they determine which mixtures an agent is willing to accept, and so encode the agent’s attitudes toward risk and gambling structure (von Neumann and Morgenstern, 1944; Mas-Colell et al., 1995).

Exercise 5.2 (Ordinal versus vNM utility). Let $x, y, z \in \mathcal{H}^*$ satisfy $x \succ y \succ z$.

- Give two different ordinal utility functions $u_1, u_2: \mathcal{H}^* \rightarrow \mathbb{R}$ that represent the same ranking of x, y, z .
- Explain why these two functions are equally good as ordinal representations.
- Suppose in addition that

$$y \sim \frac{1}{2}x + \frac{1}{2}z.$$

Show that not every strictly increasing transformation of a vNM utility can preserve this indifference.

If continuity or independence fails

It is useful to separate the roles of the four vNM axioms:

- As before, completeness and transitivity give us an ordinal utility on deterministic trajectories.
- Continuity and independence turn that ordinal picture into an *expected* utility theory over uncertain prospects.

If completeness or transitivity fail, as discussed in the previous section, we lose hope of describing the preference by a single scalar quantity. If continuity or independence fails, the preference will generally no longer admit the linear expectation form.

Failure of continuity. To build intuition, let $x \succ y \succ z$ and define

$$\mu_\lambda = \lambda x + (1 - \lambda)z.$$

Then

$$\mu_0 = z, \quad \mu_1 = x,$$

so as λ increases from 0 to 1, the lottery moves from the worse prospect z to the better prospect x . Continuity says, informally, that preferences vary smoothly along this path. In particular, it says that the intermediate prospect y can be matched by some mixture of the better and worse prospects:

$$y \sim \lambda x + (1 - \lambda)z$$

for some $\lambda \in (0, 1)$. Intuitively, the axiom requires no abrupt jumps in preference as the mixture probabilities vary.

When continuity fails, some distinctions can become *lexicographic*. For example, we could have a case where $y \succ \mu_0 = z$, but $\mu_\lambda \succ y$ for all $\lambda > 0$. This means that x has absolute priority over y : as soon as x appears with any nonzero probability, the lottery becomes strictly better than y . One way to interpret this is that some considerations are given lexical priority over others. For instance, avoiding catastrophe might outrank ordinary gains so completely that no finite improvement in ordinary reward compensates for even an arbitrarily small increase in catastrophic risk. Such preferences need not be contradictory; they are simply too sharp to be captured by a single real-valued utility whose expectations are taken in the usual way.

What breaks in this case is not all representations, but the *real-valued* vNM representation. If continuity is dropped, one is naturally led to lexicographic or non-Archimedean representations: for example, ordered pairs or vectors of utilities compared lexicographically, or utilities with infinitesimal scales (Fishburn, 1971). In those models, the first coordinate records the highest-priority consideration, and lower coordinates matter only when higher ones tie.

Failure of independence. Independence says that common lottery components should cancel. If $\mu \succ \nu$, then mixing both with the same background lottery ξ at the same rate should preserve the ranking:

$$\lambda\mu + (1 - \lambda)\xi \succ \lambda\nu + (1 - \lambda)\xi.$$

So the relative ranking of μ and ν should depend only on how they differ, not on the common part they share.

When independence fails, the value of a prospect depends on its *context*. The same local substitution can be attractive in one background and unattractive in another. This is exactly what the Allais pattern shows (Allais, 1953). Let $h_5 \succ h_1 \succ h_0$ denote trajectories yielding 5, 1, and 0 million, and consider

$$A = 1 \cdot h_1, \quad B = 0.89 h_1 + 0.10 h_5 + 0.01 h_0,$$

$$C = 0.11 h_1 + 0.89 h_0, \quad D = 0.10 h_5 + 0.90 h_0.$$

Many people prefer A to B , but prefer D to C . Under independence this is impossible, because A versus B and C versus D differ only by a common consequence. The reversal shows that certainty is treated as psychologically special: replacing a sure outcome by a tiny risk of getting nothing matters more than expected utility allows.

This is the general lesson of independence failure. Probabilities are no longer aggregated linearly against a fixed utility function on trajectories. Common branches cannot be canceled, and the whole shape of the distribution starts to matter. Decision makers may overweight certainty, distort small probabilities, care about disappointment or regret, or evaluate gains and losses relative to a reference point rather than in absolute terms. This is the route taken by prospect theory and related non-expected-utility models (Kahneman and Tversky, 1979; Machina, 1982).

From the perspective of sequential choice, independence also matters because it allows one to replace a sublottery by an equivalent one without changing the value of the larger plan. If independence fails, the value of a branch may depend on the branches surrounding it, so local and global evaluations need not line up. Hammond’s consequentialist argument shows that, together with dynamic consistency and suitable sequential assumptions, one is pushed back toward independence (Hammond, 1988). But that only shows one route to coherent planning. One may instead keep a richer, non-linear evaluation of lotteries and give up the idea that common consequences are always behaviorally irrelevant.

So the two failures have different meanings. Failure of continuity says that some priorities are infinitely sharp, leading naturally to lexicographic or infinitesimal utility scales. Failure of independence says that uncertainty is evaluated holistically rather than by linear averaging, leading to models in which background risk, certainty, or reference dependence affect choice.

Exercise 5.3 (Failures of the vNM axioms). Here we will study the consequences of different axioms.

- (a) Construct a simple incomplete preference relation on three trajectories. Why can it not be represented by a single real-valued ordinal utility?
- (b) Construct a three-cycle and explain how it gives rise to a money pump.
- (c) Give an example of a lexicographic preference over three outcomes and show that it violates continuity.
- (d) Write down the Allais pattern from Section 5 and explain which axiom it violates.

6 A fifth axiom: reward and discount

The vNM theorem gives an expected-utility representation over lotteries of whole trajectories, but it does not yet tell us that utility can be generated *locally* from stepwise rewards.

To investigate the possibility of localizing utility over time, let us denote by $T := \mathcal{O} \times \mathcal{A}$ the set of one-step transitions. For $t \in T$ and $h \in \mathcal{H}^*$, write $t \cdot h$ for the trajectory

obtained by prepending t to h , and extend this operation to lotteries by

$$t \cdot \mu := \sum_i p_i \delta_{t \cdot h_i} \quad \text{when } \mu = \sum_i p_i \delta_{h_i}.$$

Now, using a vNM utility $u(h)$ one can always define a continuation-dependent increment as

$$r(t; h) := u(t \cdot h) - u(h).$$

However, this increment may depend on the whole continuation h . What is missing is a temporal condition ensuring that the effect of prepending a transition depends only on that transition itself. Following [Bowling et al. \(2023\)](#), this can be explored with the following additional axiom.

Axiom 5 (Temporal γ -indifference). There exists a function $\gamma: T \rightarrow [0, 1]$ such that for all $\mu, \nu \in \Delta(\mathcal{H}^*)$ and all $t \in T$,

$$\frac{1}{1 + \gamma(t)}(t \cdot \mu) + \frac{\gamma(t)}{1 + \gamma(t)}\nu \sim \frac{1}{1 + \gamma(t)}(t \cdot \nu) + \frac{\gamma(t)}{1 + \gamma(t)}\mu.$$

Under expected utility, the indifference above is equivalent to

$$u(t \cdot \mu) + \gamma(t)u(\nu) = u(t \cdot \nu) + \gamma(t)u(\mu),$$

which can be rearranged as

$$u(t \cdot \mu) - u(t \cdot \nu) = \gamma(t)(u(\mu) - u(\nu)).$$

So the effect of prepending the same transition t is to rescale the utility difference between two continuations by a factor $\gamma(t)$. In that sense, $\gamma(t)$ measures how much the future still matters after the step t has occurred. When $\gamma(t) = 1$, future differences are preserved without discount at that step. When $0 < \gamma(t) < 1$, the future still matters but is damped. When $\gamma(t) = 0$, once t has occurred the continuation no longer affects the comparison. Bowling's axiom is closely related to the transition-dependent discounting discussed by [White \(2017\)](#) and [Pitis \(2019\)](#).

Bowling et al. show that adding this axiom to the four vNM axioms is necessary and sufficient for a discounted-reward representation of preferences.

Theorem 6.1 (Markov reward representation, after Bowling et al.). *A weak preference relation $\succsim$ on $\Delta(\mathcal{H}^*)$ satisfies completeness, transitivity, continuity, independence, and temporal γ -indifference if and only if there exist functions $u: \mathcal{H}^* \rightarrow \mathbb{R}$, $r: T \rightarrow \mathbb{R}$, and $\gamma: T \rightarrow [0, 1]$ such that $u(\varepsilon) = 0$,*

$$u(t \cdot h) = r(t) + \gamma(t)u(h),$$

and, for all lotteries $\mu, \nu \in \Delta(\mathcal{H}^)$,*

$$\mu \succsim \nu \iff \sum_{h \in \mathcal{H}^*} \mu(h)u(h) \geq \sum_{h \in \mathcal{H}^*} \nu(h)u(h).$$

Moreover, r is unique up to positive scale, and γ is the same function appearing in the fifth axiom (Bowling et al., 2023).

This theorem sharpens the vNM result in exactly the way reinforcement learning needs. Utility is no longer an arbitrary scalar attached to a complete trajectory; it is generated recursively from a local reward and a local weight on the future. Unrolling the recursion for a trajectory $h = (t_1, t_2, \dots, t_n)$ gives

$$u(h) = r(t_1) + \gamma(t_1)r(t_2) + \gamma(t_1)\gamma(t_2)r(t_3) + \dots + \left(\prod_{i=1}^{n-1} \gamma(t_i)\right)r(t_n).$$

Thus the fifth axiom is what allows utility over whole trajectories to be decomposed into local reward and discounting. Two familiar special cases are worth highlighting.

- If $\gamma(t) \equiv \gamma$ is constant, then

$$u(h) = \sum_{k=1}^n \gamma^{k-1} r(t_k),$$

which is the standard discounted-return objective.

- If $\gamma(t) \equiv 1$, then

$$u(h) = \sum_{k=1}^n r(t_k),$$

so utility is simply the additive cumulative reward.

In summary, the step from vNM utility over whole trajectories to RL-style reward requires one more axiom. That axiom simultaneously identifies both the local reward signal and the form of discounting.

Remark (MDPs, POMDPs, and locality)

The locality result of this section is especially natural in fully observed MDPs, where one often expects reward to depend only on the current state transition. In the present note, however, the primitive histories are sequences of observations and actions, and the corresponding local reward has the form $r: \mathcal{O} \times \mathcal{A} \rightarrow \mathbb{R}$. In partially observed settings this can be restrictive: a goal may be Markov in the hidden state, in the agent’s belief state, or in some augmented memory state, without being reducible to a function of the current observation-action pair alone. In that sense, the fifth axiom should be read as characterizing when preferences admit a reward that is local in the chosen representation of experience. If the raw observation stream is too coarse, one may need to enrich the state description before a Markov reward representation becomes available (Bowling et al., 2023).

Note

Utility is not reward.

It is helpful to keep the following four different objects apart.

- *Preference* is the primitive relation $\succsim$. It says only which trajectories or lotteries are weakly preferred to which others.
- *Preference utility* ('F1') is an ordinal representation of preferences over deterministic trajectories. Under completeness and transitivity, it is any function u_{ord} such that

$$h \succsim h' \iff u_{\text{ord}}(h) \geq u_{\text{ord}}(h').$$

It is unique only up to strictly increasing transformations ([Fishburn, 1970](#)).

- *vNM utility* ('F2') is the stronger, cardinal utility that appears when preferences over lotteries satisfy the vNM axioms. It is the function u_{vNM} for which

$$\mu \succsim \nu \iff \sum_h \mu(h) u_{\text{vNM}}(h) \geq \sum_h \nu(h) u_{\text{vNM}}(h).$$

It is unique only up to positive affine transformations ([von Neumann and Morgenstern, 1944](#)). Thus 'F2' refines 'F1': it agrees with it on the ranking of certain trajectories, but adds the extra structure needed to compare lotteries.

- *Reward* is not either of these utilities. A reward function $r(t)$ is a *local* representation introduced only after adding temporal structure. In the present framework, reward appears when the fifth axiom allows utility to be written recursively as

$$u(t \cdot h) = r(t) + \gamma(t)u(h).$$

So reward is a way of *decomposing* utility over complete trajectories into stepwise contributions. It is therefore downstream of preference, and even downstream of vNM utility. This is why different reward functions can encode the same utility or the same preference ordering, a point we return to in the next section ([Bowling et al., 2023](#)).

Exercise 6.1 (Guided proof of the Markov reward representation theorem). Assume the hypotheses of Exercise 5.1 together with temporal γ -indifference, and let u be a vNM utility representation over trajectories.

- (a) Apply expected utility to temporal γ -indifference and show that for all

lotteries $\mu, \nu \in \Delta(\mathcal{H}^*)$ and all transitions $t \in T$,

$$u(t \cdot \mu) + \gamma(t)u(\nu) = u(t \cdot \nu) + \gamma(t)u(\mu).$$

(b) Rearrange part (a) to prove that

$$u(t \cdot \mu) - \gamma(t)u(\mu)$$

is independent of the choice of μ .

(c) Define

$$r(t) := u(t \cdot \mu) - \gamma(t)u(\mu)$$

for any μ . Deduce that for every deterministic trajectory h ,

$$u(t \cdot h) = r(t) + \gamma(t)u(h).$$

(d) Show that this recursion extends to lotteries by linearity:

$$u(t \cdot \mu) = r(t) + \gamma(t)u(\mu).$$

(e) Prove the converse direction: if u , r , and γ satisfy

$$u(t \cdot h) = r(t) + \gamma(t)u(h)$$

for all t and h , and if u extends linearly to lotteries, then temporal γ -indifference holds.

(f) Unroll the recursion along a finite trajectory $h = (t_1, \dots, t_n)$ and derive

$$u(h) = r(t_1) + \gamma(t_1)r(t_2) + \gamma(t_1)\gamma(t_2)r(t_3) + \dots + \left(\prod_{i=1}^{n-1} \gamma(t_i) \right) r(t_n).$$

(g) Show that when $\gamma(t) \equiv \gamma$ is constant, the previous formula reduces to standard discounted return, and when $\gamma(t) \equiv 1$ it reduces to additive cumulative reward.

(h) Show that if $u' = a + bu$, then the corresponding reward must be

$$r'(t) = br(t) + (1 - \gamma(t))a.$$

Interpret this as a source of reward non-uniqueness.

7 Reward equivalence and shaping

The preceding theorem shows how to pass from utility over lotteries of trajectories to a local reward representation. But it also raises an important question: how unique is that reward? The answer has two parts. If one fixes a particular utility function u and a particular discount function γ , then the reward is determined. But if one changes the numerical representative of the same preference relation, or allows shaping transformations, then many different rewards can encode the same underlying ordering.

First observe that once u and γ are fixed, the reward is fixed as well. Indeed, if

$$u(t \cdot h) = r(t) + \gamma(t)u(h)$$

for all $t \in T$ and $h \in \mathcal{H}^*$, then necessarily

$$r(t) = u(t \cdot h) - \gamma(t)u(h),$$

and the right-hand side must be independent of the continuation h . So the non-uniqueness of reward does not come from ambiguity inside a fixed recursive representation. It comes from the fact that utility itself is not unique as a numerical object.

Proposition 7.1 (Multiple rewards can encode the same preference relation).

Suppose $u: \mathcal{H}^* \rightarrow \mathbb{R}$, $r: T \rightarrow \mathbb{R}$, and $\gamma: T \rightarrow [0, 1]$ satisfy

$$u(t \cdot h) = r(t) + \gamma(t)u(h)$$

for all $t \in T$ and $h \in \mathcal{H}^*$, and that preferences over lotteries are represented by the expected value of u . For any $a \in \mathbb{R}$ and $b > 0$, define

$$u'(h) = a + bu(h) \quad \text{and} \quad r'(t) = br(t) + (1 - \gamma(t))a.$$

Then $u'(t \cdot h) = r'(t) + \gamma(t)u'(h)$ for all $t \in T$ and $h \in \mathcal{H}^*$, and the expected value of u' represents the same preference relation as the expected value of u .

The proposition shows that reward non-uniqueness already appears at the level of affine reparameterizations of utility: changing the numerical representative of the same vNM preference relation generally changes the associated reward function as well. This point is especially simple under the normalization $u(\varepsilon) = 0$ used in the previous section. Then the additive degree of freedom disappears, and one is left only with positive rescalings:

$$u'(h) = bu(h), \quad r'(t) = br(t).$$

So in the normalized setting reward is unique only up to choice of units.

There is also a second, less trivial kind of non-uniqueness, in which one changes rewards while leaving the induced trajectory ordering unchanged, first studied by Ng et al. (1999).

Proposition 7.2 (Potential-based shaping changes utility only by a boundary term). *Consider a state-based setting with transitions $t = (s, a, s')$ and a potential function Φ on states.*

(a) *If $\gamma \equiv 1$ and $r_\Phi(s, a, s') := r(s, a, s') + \Phi(s') - \Phi(s)$, then*

$$\sum_{k=1}^n r_\Phi(t_k) = \sum_{k=1}^n r(t_k) + \Phi(s_n) - \Phi(s_0)$$

for any trajectory $\tau = (t_1, \dots, t_n)$ from s_0 to s_n .

(b) *If $\gamma \in (0, 1)$ is constant and $r_\Phi(s, a, s') := r(s, a, s') + \gamma\Phi(s') - \Phi(s)$, then*

$$\sum_{k=1}^n \gamma^{k-1} r_\Phi(t_k) = \sum_{k=1}^n \gamma^{k-1} r(t_k) - \Phi(s_0) + \gamma^n \Phi(s_n).$$

The previous result shows that shaping r into r_Φ changes utility only by a boundary term. In particular, if all compared trajectories share the same start state and terminal potential, then the induced preference ordering is unchanged; if the boundary term vanishes on all admissible trajectories, then even the numerical utility is unchanged.

The general lesson is that the same preferences admit many different reward functions, so no single one of them is canonical. Some transformations, such as positive scaling, merely change the numerical units. Others, such as potential-based shaping, redistribute value along the trajectory while leaving the overall preference over complete trajectories unchanged. This is why reward design is often underdetermined: what matters behaviorally is not a raw reward function in isolation, but the utility and preference structure it induces.

Exercise 7.1 (Reward shaping). Assume additive reward with $\gamma \equiv 1$ and define

$$r_\Phi(s, a, s') = r(s, a, s') + \Phi(s') - \Phi(s).$$

- (a) Show that along any finite trajectory $(t_1, \dots, t_n)$ the total shaped reward differs from the original total reward by the boundary term $\Phi(s_n) - \Phi(s_0)$.
- (b) Under what condition on the admissible trajectories does this shaping leave utility exactly unchanged?
- (c) Under what weaker condition does it leave only the induced preference ordering unchanged?

8 Conclusion

The main lesson of this note is that, in sequential settings, the primitive object is not a one-step reward but a preference over complete trajectories. From that starting point one can distinguish several layers of structure. Completeness and transitivity yield an ordinal utility over deterministic trajectories. Once lotteries are introduced, continuity and independence refine that ordinal picture into a von Neumann–Morgenstern utility, which agrees with the ranking of certain trajectories but carries strictly more information because it calibrates trade-offs between risky prospects.

This separation of layers clarifies both what expected utility achieves and what it leaves out. It explains why “utility” in economics can refer either to an ordinal representation of certain preferences or to a cardinal object suitable for evaluating lotteries. It also shows what fails when particular axioms are dropped: without completeness or transitivity, scalar representation itself can fail; without continuity or independence, one may still have meaningful preferences, but no longer the linear expectation form of vNM. In that sense, expected utility is not the whole of rationality, but a particular and mathematically powerful strengthening of it.

The final step is to see that even vNM utility is not yet reward. To obtain a local reward-and-discount representation one needs additional temporal structure, captured here by Bowling et al.’s fifth axiom. Under that axiom, utility over whole trajectories admits a recursive decomposition into local reward and discounting. But even then reward is not unique: affine changes of utility induce corresponding changes of reward, and potential-based shaping can redistribute value along a trajectory while preserving the underlying ordering. Moreover, the locality of reward depends on the chosen representation of experience, which is natural in fully observed MDPs but can be restrictive in partially observed settings unless the state is suitably enriched. The reward hypothesis, understood in this way, is therefore best read not as the claim that goals are primitively rewards, but as the claim that sufficiently structured preferences over trajectories can be represented by rewards.

References

- Maurice Allais. Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école américaine. *Econometrica*, 21(4):503–546, 1953.
- Robert J Aumann. Utility theory without the completeness axiom. *Econometrica*, 30(3): 445–462, 1962.
- Michael Bowling, John D. Martin, David Abel, and Will Dabney. Settling the reward hypothesis. *arXiv preprint arXiv:2212.10420*, 2023. URL <https://arxiv.org/abs/2212.10420>.
- Gerard Debreu. Representation of a preference ordering by a numerical function. In Robert M. Thrall, Clyde H. Coombs, and Ralph L. Davis, editors, *Decision Processes*, pages 159–165. Wiley, New York, 1954.

- Peter C. Fishburn. *Utility Theory for Decision Making*. Wiley, New York, 1970.
- Peter C. Fishburn. A study of lexicographic expected utility. *Management Science*, 17(11):672–678, 1971. doi: 10.1287/mnsc.17.11.672.
- Johan E. Gustafsson. A money-pump for acyclic intransitive preferences. *Dialectica*, 64(2):251–257, 2010. doi: 10.1111/j.1746-8361.2010.01230.x.
- Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25(1):25–78, 1988. doi: 10.1007/BF00129168.
- Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, 127(1):203–244, 2011. doi: 10.1007/s10107-010-0419-x.
- Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
- David M. Kreps. *Notes on the Theory of Choice*. Westview Press, Boulder, CO, 1988.
- Mark J. Machina. "expected utility" analysis without the independence axiom. *Econometrica*, 50(2):277–323, 1982.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- Andrew Y. Ng, Daishi Harada, and Stuart J. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning*, pages 278–287, 1999.
- Silviu Pitis. Rethinking the discount factor in reinforcement learning: A decision theoretic approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.
- Mehran Shakerinava and Siamak Ravanbakhsh. Utility theory for sequential decision making. In *Proceedings of the International Conference on Machine Learning*, 2022.
- John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- Martha White. Unifying task specification in reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

Open-ended questions

Exercise 8.1. Do the results have any prescriptive or descriptive implications? What kinds of agents, with what kinds of preferences should we design? By default, what kind of agents should we expect to obtain via typical training processes?

Exercise 8.2. Let's assume a superintelligence has preferences over trajectories. Are there weaker versions of the axioms that make it more plausible that these preferences are "safe for us" compared to preferences that lead to utilities or even rewards?

Exercise 8.3. More generally, are there interesting examples of preferences over trajectories that the students could analyse, that do or do not satisfy some of the axioms?

A Solutions

Solution to Exercise 5.1

- (a) Induction on $|X|$. A single trajectory is its own best and worst element. Given a maximal b' and minimal w' for $X \setminus \{h\}$, completeness compares h with each of them and transitivity makes the better of $\{h, b'\}$ maximal and the worse of $\{h, w'\}$ minimal for X .
- (b) If $b \sim w$ then every h is indifferent to both and any constant u works. Otherwise $b \succ w$. By (a) we have $b \succcurlyeq h \succcurlyeq w$, so continuity applied to the triple (b, h, w) gives directly some $\lambda \in [0, 1]$ with $h \sim \lambda b + (1 - \lambda)w$; set $u(h) := \lambda$. For uniqueness, independence implies the mixtures $\alpha b + (1 - \alpha)w$ are strictly increasing in α when $b \succ w$, so two different weights cannot both be indifferent to h .
- (c) Suppose $u(h) \geq u(h')$. Mixture monotonicity gives $u(h)b + (1 - u(h))w \succcurlyeq u(h')b + (1 - u(h'))w$, and chaining the two calibrating indifferences through transitivity yields $h \succcurlyeq h'$. If $u(h) < u(h')$ the same argument gives $h' \succ h$. Together: $h \succcurlyeq h' \iff u(h) \geq u(h')$.
- (d) Independence states that $\mu \sim \nu$ implies $p\mu + (1 - p)\lambda \sim p\nu + (1 - p)\lambda$. View μ as a mixture in which h_1 appears with weight p_1 against the rest; replacing h_1 by its calibrated equivalent $u(h_1)b + (1 - u(h_1))w$ therefore leaves the whole lottery indifferent. Doing this for each h_i in turn — finitely many steps — and chaining with transitivity gives $\mu \sim \sum_i p_i(u(h_i)b + (1 - u(h_i))w)$.
- (e) The right-hand side is a compound lottery whose only outcomes are b and w ; it awards b with total probability $\sum_i p_i u(h_i) = U(\mu)$. Identifying a compound lottery with the simple lottery it induces gives $\mu \sim U(\mu)b + (1 - U(\mu))w$.
- (f) By (e), μ and ν are indifferent to calibrated b – w mixtures with weights $U(\mu)$ and $U(\nu)$; by mixture monotonicity and transitivity, $\mu \succcurlyeq \nu \iff U(\mu) \geq U(\nu)$.
- (g) Apply the u' -representation to the calibration of h :

$$\begin{aligned} u'(h) &= U'(u(h)b + (1 - u(h))w) = u(h) u'(b) + (1 - u(h)) u'(w) \\ &= u'(w) + (u'(b) - u'(w))u(h). \end{aligned}$$

This is the required form $u' = a + bu$, with intercept $u'(w)$ and slope $u'(b) - u'(w)$, the latter strictly positive because $b \succ w$ and u' represents $\succcurlyeq$. (Note the statement's scalars a, b are unrelated to the trajectories b, w of part (a); the letter b does double duty here, so read $u'(b)$ as the utility of the best trajectory throughout.)

- (h) Let $\mu \succcurlyeq \nu \iff \mathbb{E}_\mu[u] \geq \mathbb{E}_\nu[u]$. Completeness and transitivity are inherited from the total order on $\mathbb{R}$. Continuity: $p \mapsto \mathbb{E}_{p\mu+(1-p)\lambda}[u] = p\mathbb{E}_\mu[u] + (1-p)\mathbb{E}_\lambda[u]$ is continuous, so upper and lower contour sets in p are closed. Independence: mixing both sides with λ adds the same $(1-p)\mathbb{E}_\lambda[u]$ and scales the difference by $p > 0$, leaving the comparison unchanged.

Solution to Exercise 5.2

- (a) For instance $u_1(x, y, z) = (2, 1, 0)$ and $u_2(x, y, z) = (100, 5, -3)$.
- (b) An ordinal representation encodes only the ranking: u represents $\succcurlyeq$ iff $u(h) \geq u(h') \iff h \succcurlyeq h'$, and both functions induce $x \succ y \succ z$. Composing with any strictly increasing map preserves exactly this information, so no ordinal criterion can distinguish u_1 from u_2 .
- (c) Under expected utility the indifference forces $u(y) = \frac{1}{2}u(x) + \frac{1}{2}u(z)$. Take the vNM utility $u(x, y, z) = (1, \frac{1}{2}, 0)$ and the strictly increasing map $f(s) = s^3$. Then $(f \circ u)(x, y, z) = (1, \frac{1}{8}, 0)$, while the lottery $\frac{1}{2}x + \frac{1}{2}z$ has expected value $\frac{1}{2}$. Since $\frac{1}{8} \neq \frac{1}{2}$, the relation represented by $f \circ u$ strictly prefers the lottery to y : the indifference is destroyed. Only positive affine transformations preserve all such midpoint identities, which is the uniqueness part of the vNM theorem.

Solution to Exercise 5.3

- (a) Let $x \succ y$, $x \succ z$, with y and z incomparable (neither $y \succcurlyeq z$ nor $z \succcurlyeq y$). Any $u: \{x, y, z\} \rightarrow \mathbb{R}$ satisfies $u(y) \geq u(z)$ or $u(z) \geq u(y)$ because the reals are totally ordered, so the relation it represents is complete; incomparability cannot be encoded by a single real-valued function.
- (b) Take $x \succ y \succ z \succ x$. Suppose you hold x and are willing to pay some small $\varepsilon > 0$ to exchange an item for one you strictly prefer. Since $z \succ x$ you pay to swap $x \rightarrow z$; since $y \succ z$ you pay to swap $z \rightarrow y$; since $x \succ y$ you pay to swap $y \rightarrow x$. You are back where you started, 3ε poorer, and the cycle can be run forever.
- (c) Give outcomes two attributes and compare lexicographically (the first attribute decides unless tied): $A = (1, 0)$, $B = (0, 1)$, $C = (0, 0)$, so $A \succ B \succ C$. Extend to lotteries by comparing expected attribute vectors lexicographically. Continuity would require some $p \in [0, 1]$ with $B \sim pA + (1-p)C$. But that mixture has attribute vector $(p, 0)$: for every $p > 0$ it beats B on the first attribute, and for $p = 0$ it is $C \succ$ -below B . No mixing weight produces indifference.

- (d) With $h_5 \succ h_1 \succ h_0$ paying 5, 1 and 0 million, the pattern of Section 5 is

$$A = 1 \cdot h_1, \quad B = 0.89 h_1 + 0.10 h_5 + 0.01 h_0,$$

$$C = 0.11 h_1 + 0.89 h_0, \quad D = 0.10 h_5 + 0.90 h_0,$$

with $A \succ B$ but $D \succ C$. Both pairs differ only by a *common consequence*: replacing 0.89 of h_1 by 0.89 of h_0 turns A into C and B into D . Independence says a ranking is unchanged when the same consequence is mixed into both sides with the same weight, so the pattern violates independence (completeness, transitivity, and continuity are all consistent with it).

Solution to Exercise 6.1

- (a) Apply u (in expected-utility form) to both sides of the axiom's indifference and multiply by $1 + \gamma(t)$:

$$u(t \cdot \mu) + \gamma(t)u(\nu) = u(t \cdot \nu) + \gamma(t)u(\mu).$$

- (b) Rearranging, $u(t \cdot \mu) - \gamma(t)u(\mu) = u(t \cdot \nu) - \gamma(t)u(\nu)$ for *all* μ, ν : the quantity does not depend on which lottery is used to compute it.

- (c) By (b) the definition of $r(t)$ is unambiguous. Taking μ to be the point mass on a deterministic continuation h gives $u(t \cdot h) = r(t) + \gamma(t)u(h)$.

- (d) $t \cdot \mu$ is the pushforward of μ under prepending t , and u is linear on lotteries, so $u(t \cdot \mu) = \mathbb{E}_{h \sim \mu}[u(t \cdot h)] = \mathbb{E}_{h \sim \mu}[r(t) + \gamma(t)u(h)] = r(t) + \gamma(t)u(\mu)$.

- (e) With the recursion and linearity, $u(t \cdot \mu) + \gamma(t)u(\nu) = r(t) + \gamma(t)(u(\mu) + u(\nu))$, which is symmetric under $\mu \leftrightarrow \nu$; hence $u(t \cdot \mu) + \gamma(t)u(\nu) = u(t \cdot \nu) + \gamma(t)u(\mu)$. Dividing by $1 + \gamma(t)$ shows both mixtures in the axiom have equal expected utility, and since u represents $\succsim$, they are indifferent.

- (f) Normalise $u(\varepsilon) = 0$ (allowed by affine freedom). Induction on n : $u(h) = r(t_1) + \gamma(t_1)u((t_2, \dots, t_n))$, and expanding the inner term yields

$$u(h) = r(t_1) + \gamma(t_1)r(t_2) + \gamma(t_1)\gamma(t_2)r(t_3) + \dots + \left(\prod_{i=1}^{n-1} \gamma(t_i)\right)r(t_n).$$

- (g) If $\gamma(t) \equiv \gamma$, the product $\prod_{i < k} \gamma(t_i)$ is γ^{k-1} and $u(h) = \sum_k \gamma^{k-1}r(t_k)$: discounted return. If $\gamma \equiv 1$ every product is 1 and $u(h) = \sum_k r(t_k)$: additive cumulative reward.

- (h) $r'(t) = u'(t \cdot \mu) - \gamma(t)u'(\mu) = (a + b u(t \cdot \mu)) - \gamma(t)(a + b u(\mu)) = b r(t) + (1 - \gamma(t))a$. The same preferences thus admit a family of reward functions: even with γ fixed, r is only determined up to a positive scaling and an additive shift modulated by $1 - \gamma(t)$.

Solution to Exercise 7.1

(a) Writing $t_i = (s_{i-1}, a_i, s_i)$, the shaped total is

$$\sum_{i=1}^n r_{\Phi}(t_i) = \sum_{i=1}^n r(t_i) + \sum_{i=1}^n (\Phi(s_i) - \Phi(s_{i-1})) = \sum_{i=1}^n r(t_i) + \Phi(s_n) - \Phi(s_0),$$

since the middle sum telescopes.

- (b) Utility is exactly unchanged iff the boundary term vanishes on every admissible trajectory, i.e. $\Phi(s_n) = \Phi(s_0)$ throughout — for example when all trajectories start in a fixed s_0 and Φ takes the value $\Phi(s_0)$ on every reachable terminal state.
- (c) Only the ordering is at stake if the boundary term is the *same constant* c for all compared trajectories: then every utility shifts by c and the induced ranking is untouched. A fixed start state together with Φ constant on the reachable terminal states suffices, whatever that constant is.

Solution to Exercise 8.1

Discussion notes, not a unique answer. Prescriptively, the money-pump and dominance arguments say that an agent which cares about a fungible resource is under pressure toward completeness and transitivity — coherence is an attractor for agents that can be exploited otherwise. Descriptively, nothing guarantees that trained systems satisfy any axiom: gradient descent on episodic objectives can produce context-dependent, intransitive, or incomparable preferences, especially off-distribution. A reasonable expectation is approximate coherence where incoherence was penalised during training, and no guarantee elsewhere. For design, the trade-off runs both ways: highly coherent agents are more predictable and analysable but are exactly the ones for which instrumental-convergence arguments bite; agents with incomplete or unstable preferences may be harder to exploit into goal-directed resource acquisition, at the cost of being harder to reason about.

Solution to Exercise 8.2

Discussion notes. The natural candidates weaken one axiom at a time. Dropping *completeness* is the most studied: an agent with incomplete preferences can remain undecided between continuing and being shut down, so it is not pushed by coherence arguments toward shutdown-resistance; money pumps also lose force because the agent may simply refuse trades between incomparable options. Weakening *continuity* permits lexicographic safety: “never cross the constraint, then optimise” cannot be represented by a single real-valued utility, which is arguably a feature. Weakening *independence* allows certainty-favouring (Allais-like) preferences, which dampen gambling-for-resources behaviour. Weakening

temporal γ -indifference removes the local reward representation: goals about the shape of a trajectory as a whole (diversity, “never do X”) stop being expressible as accumulated reward. The common pattern: each axiom kept buys a representation theorem and optimisation pressure; each axiom dropped blocks a coherence-based failure mode while making the agent’s behaviour less analysable.

Solution to Exercise 8.3

Discussion notes; instructive examples include: (i) lexicographic safety-first preferences — violate continuity (see Exercise 5.3); (ii) Pareto/multi-objective preferences, undominated but unaggregated — violate completeness; (iii) hyperbolic discounting — satisfies the vNM axioms at a single decision time yet violates temporal γ -indifference, hence admits utility but no stationary reward/discount pair, and is dynamically inconsistent; (iv) the certainty effect (Allais) — violates independence only; (v) satisficing (“anything above the threshold is equally fine”) — complete and transitive, so ordinal utility exists, but indifference plateaus interact oddly with lotteries near the threshold; (vi) trajectory-shape goals such as “visit as many distinct states as possible” — can satisfy all four vNM axioms (a utility exists) while violating the temporal axiom, illustrating exactly the gap between utility and reward.