

Nonhuman Primates Satisfy Utility Maximization in Compliance with the Continuity Axiom of Expected Utility Theory

Simone Ferrari-Toniolo,¹ Philippe M. Bujold,¹ Fabian Grabenhorst,¹ Raymundo Báez-Mendoza,^{1,2} and Wolfram Schultz¹

¹Department of Physiology, Development and Neuroscience, University of Cambridge, Cambridge, CB2 3DY, United Kingdom, and ²Department of Neurosurgery, Massachusetts General Hospital, Harvard Medical School, Boston, Massachusetts 02114

Expected Utility Theory (EUT), the first axiomatic theory of risky choice, describes choices as a utility maximization process: decision makers assign a subjective value (utility) to each choice option and choose the one with the highest utility. The continuity axiom, central to Expected Utility Theory and its modifications, is a necessary and sufficient condition for the definition of numerical utilities. The axiom requires decision makers to be indifferent between a gamble and a specific probabilistic combination of a more preferred and a less preferred gamble. While previous studies demonstrated that monkeys choose according to combinations of objective reward magnitude and probability, a concept-driven experimental approach for assessing the axiomatically defined conditions for maximizing utility by animals is missing. We experimentally tested the continuity axiom for a broad class of gamble types in 4 male rhesus macaque monkeys, showing that their choice behavior complied with the existence of a numerical utility measure as defined by the economic theory. We used the numerical quantity specified in the continuity axiom to characterize subjective preferences in a magnitude-probability space. This mapping highlighted a trade-off relation between reward magnitudes and probabilities, compatible with the existence of a utility function underlying subjective value computation. These results support the existence of a numerical utility function able to describe choices, allowing for the investigation of the neuronal substrates responsible for coding such rigorously defined quantity.

Key words: continuity axiom; decision making; expected utility theory; reward risk; utility

Significance Statement

A common assumption of several economic choice theories is that decisions result from the comparison of subjectively assigned values (utilities). This study demonstrated the compliance of monkey behavior with the continuity axiom of Expected Utility Theory, implying a subjective magnitude-probability trade-off relation, which supports the existence of numerical utility directly linked to the theoretical economic framework. We determined a numerical utility measure able to describe choices, which can serve as a correlate for the neuronal activity in the quest for brain structures and mechanisms guiding decisions.

Received Apr. 22, 2020; revised Nov. 13, 2020; accepted Nov. 19, 2020.

Author contributions: S.F.-T. and W.S. designed research; S.F.-T., P.M.B., F.G., and R.B.-M. performed research; S.F.-T. analyzed data; S.F.-T. wrote the first draft of the paper; S.F.-T. and W.S. edited the paper; S.F.-T. wrote the paper.

This work was supported by Wellcome Grants WT 095495 and WT 204811, and European Research Council Advanced Grant 293549. We thank Aled David and Christina Thompson for animal and technical support; and Federico M. Echenique and David M. Grether (Caltech) for helpful comments.

The authors declare no competing financial interests.

Correspondence should be addressed to Simone Ferrari-Toniolo at Simone.FerrariToniolo@gmail.com.

<https://doi.org/10.1523/JNEUROSCI.0955-20.2020>

Copyright © 2021 Ferrari-Toniolo et al.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

Introduction

Expected Utility Theory (EUT), the first axiomatic theory of risky choice (von Neumann and Morgenstern, 1944), demonstrates that, if a subject's behavior followed a simple set of rules, or axioms, their choices could be described by utility maximization (see Expected utility theorem). The utility concept is used by economists to mathematically define subjective values, and its maximization is a general and basic process determining the subject's survival. The core of EUT remained central to the generalized expected utility theories developed later, most notably Prospect Theory (Kahneman and Tversky, 1979; Hey and Orme, 1994; Starmer, 2000), which overcame the limits of EUT by introducing the concepts of probability weighting (PW) and reference point.

The continuity axiom of EUT, given three subjectively ranked gambles, requires the existence of an indifference point (IP) between the intermediate gamble and a probabilistic combination of the other two. This ensures that no option is considered infinitely better than any other option, making it possible to define a finite, numerical subjective value for each gamble (Jehle and Reny, 2001; Chateauneuf et al., 2008; Johnson and Ratcliff, 2014). Continuity is common to a broad spectrum of risky and riskless choice theories (Samuelson, 1948; Weber and Camerer, 1987), emerging as a fundamental construct in all economic schemes based on value computation.

Compliance with the continuity axiom can be experimentally tested by verifying the existence of indifference between the intermediate gamble and one of the probabilistic combinations. Violations according to the mathematical definition cannot be experimentally tested, as they would require infinitesimally small probability steps and thus an infinite number of tests. Nevertheless, given the biological limits of perceiving small probability differences (Tversky, 1969), tests within the finite range of noticeable probability differences could be interpreted as being biologically plausible, and finding violations in such limited-range tests could be considered relevant.

Monkeys are the evolutionary closest species to humans that allow to investigate, at the single-cell level and over multiple repetitions, the neuronal correlates of economic choice between options with finely varying reward magnitude and probability. Monkeys prefer gambles with larger expected values (EVs) to those with smaller EVs derived from both reward magnitude and probability (Musallam et al., 2004; Lau and Glimcher, 2005; Averbeck, 2015; Farashahi et al., 2018). These results suggested a trade-off between reward magnitude and probability, but their interpretation relies tacitly on the validity of the EUT continuity axiom. The axiom would formally distinguish between circumstantial and systematic trade-offs but was not tested in these studies. Further, behavioral compliance with EUT axioms would provide a conceptual foundation for the meaningful neuronal coding of subjective value and formal economic utility reported previously (Platt and Glimcher, 1999; Tremblay and Schultz, 1999; Padoa-Schioppa and Assad, 2006; Kobayashi and Schultz, 2008; Lak et al., 2014; Stauffer et al., 2014).

Here, we tested whether economic choices of rhesus monkeys comply with the EUT continuity axiom, assuming no adaptation in reference point. The animals chose between a gamble and a probabilistic combination of a more preferred and a less preferred gamble. The animals revealed their preference by selecting one of the two options. Our statistical test procedure (preferences ranked with probability, the combined option going from non-preferred to preferred) identified significant compliance with the axiom. To address the general axiom definition, we tested a broad range of magnitudes and probabilities, starting with degenerate gambles (i.e., only one outcome, probability $p = 1.0$) and advancing to gradually more complex gambles containing two or three possible outcomes.

We found that the choices of all 4 animals complied with the continuity axiom by expressing consistent preferences, which robustly identified IPs across the tested magnitude-probability (MP) space. The animals' behavior aligned closely with choice functions that modeled the MP trade-off. This result suggests the existence of an axiomatically defined utility measure for choice options. Compatible with economic theory, the pattern of subjective IPs allowed us to define a utility function that is suitable for investigating the neuronal mechanisms of risky choice.

Materials and Methods

Animals and ethical approval

Four adult male macaque monkeys (*Macaca mulatta*) were used in these experiments (weight, per animal): 12.7 kg (Monkey A), 13.8 kg (Monkey B), 10.3 kg (Monkey C), and 12.5 kg (Monkey D).

This research has been ethically reviewed, approved, regulated and supervised by the following institutions and individuals in the United Kingdom and at the University of Cambridge (UCam): the UK Home Office implementing the Animals (Scientific Procedures) Act 1986 with Amendment Regulations 2012, the local UK Home Office Inspector, the UK Animals in Science Committee, the UK National Center for Replacement, Refinement and Reduction of Animal Experiments, two UCam Animal Welfare and Ethical Review Bodies, the UCam Governance and Strategy Committee, the Certificate Holder of the UCam Biomedical Service (UBS), the UBS Director for Governance and Welfare, the UBS Named Veterinary Surgeon, and the UBS Named Animal Care and Welfare Officer.

Experimental design and statistical analysis

We presented each animal with a set of two mutually exclusive and collectively exhaustive options that appeared simultaneously and at equal joystick distance in front of the animal on a computer monitor. In each option, we set the magnitude and probability of reward independently.

We tested choices by implementing the following basic concepts:

1. Increasing the reward probability in one option, while holding constant the reward magnitude in this option and the reward probability and magnitude in its alternative, leads to monotonically increasing probability of choosing this option over its alternative, as modeled by an S-shaped psychophysical choice function (see Figs. 1e, 2).
2. Options are equally revealed preferred, and inferred to have equal utility for the animal, when the animal chooses them with equal probability ($p = 0.5$ each option).
3. Each option with its specific reward magnitude and probability is graphically represented at the intersection of the x coordinate (probability) and y coordinate (magnitude) of a two-dimensional plot (see Fig. 4).
4. An option that is as revealed preferred as another option, as shown by choice indifference, is graphically represented as a two-dimensional IP. The dots in Figure 4b–d are IPs relative to all other color-matched dots.
5. Several IPs align as an indifference curve (IC) on which each option is as revealed preferred as any other option on that same IC, despite different MP composition. The colored curves in Figure 4b, d are ICs. Options on higher ICs (farther from origin) are revealed preferred to options on lower ICs, and all options on a given IC are revealed preferred to options below that IC. Options on the red and orange ICs in Figure 4c, d are revealed preferred to options on the blue and green ICs.

The axioms of EUT are necessary and sufficient conditions for choices to maximize Expected Utility (EU). $EU = \sum_i U(m_i) \cdot p_i$, is defined as the utility $U(m)$ of the possible outcome magnitudes (m_i) weighted by their respective probabilities of occurrence (p_i). This subjectively defined quantity, EU, replaced the objective $EV = \sum_i m_i \cdot p_i$, as a key measure driving decisions (Pascal, 2004).

Compliance with the continuity axiom requires the existence of a probability at which a fixed gamble is choice indifferent against a combination of a higher and a lower gamble with that probability; the axiom implies the possibility of defining a numerical scale of subjective values. Being deterministic rules, the axioms assume perfectly constant preferences over time. In order to account for the stochasticity of choice behavior, we interpreted the axioms in a stochastic sense: option A was considered preferred to option B when the proportion of A over B choices was >0.5 (binomial test, $p < 0.05$). The standardized logistic regression coefficients were tested for statistical significance through one-sample t test. All statistical tests used (binomial test, Spearman rank correlation, one-sample t test, likelihood ratio test) were considered significant at $p < 0.05$. For multiple comparisons, we applied a false discovery rate (FDR) correction (Benjamini-Hochberg procedure) (Benjamini

and Hochberg, 1995). Data analysis was performed using custom scripts in MATLAB (version 8.3.0, R2014a; The MathWorks).

Task design

Each of the 4 animals was trained to express its preferences in > 10,000 trials between pairs of probabilistic reward options, represented as visual cues on a computer monitor. During the experiment, monkeys sat in primate chairs (Crist Instruments) in the laboratory and used arm movements to make choices between two rewarding stimuli presented on a computer monitor. Monkeys A and B moved a joystick (Biotronix Workshop) restricted to left/right movements, to control a cursor on a computer monitor vertically positioned 50 cm in front of them; Monkeys C and D made arm movements toward a touch-sensitive screen (EloTouch 1522L 15; Tyco) horizontally mounted at arm-reaching distance. The possible choice outcomes were different amounts of liquid reward (fruit juice), ranging 0.00 to 0.50 ml (Monkeys A and B) or 0.05 to 0.90 ml (Monkeys C and D). A computer-controlled solenoid valve delivered juice reward from a spout in front of the animal's mouth. Task event-times were sampled and stored at 1 kHz on a Windows 7 computer running custom MATLAB (The MathWorks) code, using Psychtoolbox 3.

During the initial training period (1–2 months), monkeys learned to select the side of the screen where a single visual stimulus was presented. Initially, the stimulus represented a sure reward, with magnitude varying among up to three levels (between 0.05 and 0.5 ml) in pseudo-randomly intermingled trials. Then, to learn the probabilistic information, gambles were introduced in a similar scheme, alternating up to three probability levels (between 0.1 and 0.9). A further training period (2–4 months) included choices between two learned stimuli. Gradually, further magnitude and probability levels and more complex gambles were added to the choice set. Behavioral data acquisition began when preferences were considered stable and clearly differed among the newly introduced gambles.

Task design for Monkeys A and B. The reward amount (magnitude) was represented though the vertical position of a horizontal white bar within a frame, composed by two thin vertical gray lines. A single option could contain up to three possible outcomes, each with a specific probability. The probability associated with each outcome was cued through the width of the horizontal bar. Each choice option could be either a safe option (i.e., a sure reward or “degenerate gamble,” with probability $p = 1$), presented as a single horizontal bar filling the full width of the frame, or a probabilistic distribution of rewards (i.e., a risky gamble) presented as multiple horizontal bars. The horizontal position of the bars representing nonsafe outcomes were randomly shifted horizontally within the frame to avoid that animals only considered a particular portion of the stimulus.

To initiate a trial, the monkey held a joystick in the central position for a variable time interval (1–1.5 s). Two visual cues representing the choice options appeared to the left and right sides of the computer

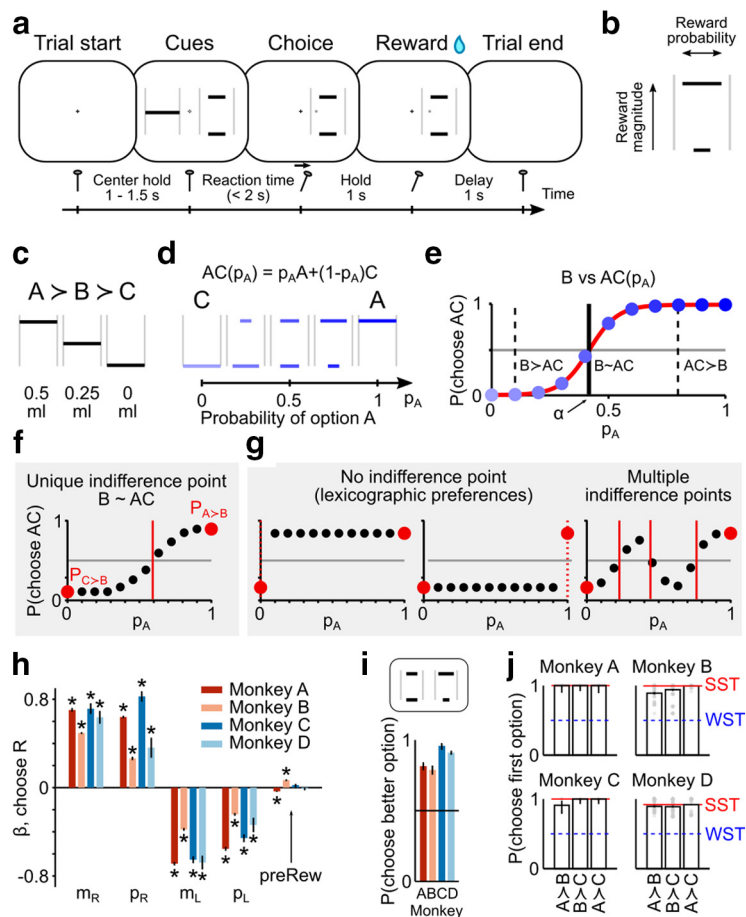


Figure 1. Experimental design and consistency of choice behavior. **a**, Trial sequence. Monkeys chose between two options by moving a cursor (gray dot) with a joystick to one side of the screen. After a delay, the reward corresponding to the selected cue was delivered. **b**, Visual cues indicated magnitude and probability of possible outcomes through horizontal bars' vertical position and width, respectively. **c–e**, Continuity axiom test. The continuity axiom was tested through choices between a fixed gamble B and a probabilistic combination of A and C (AC). A, B, and C were ordered reward magnitudes (**c**); AC was a gamble between A and C, with probabilities p_A and $1 - p_A$, respectively (**d**); different shades of blue represent different p_A values (darker for higher p_A). The continuity axiom implies the existence of a unique AC combination ($p_A = \alpha$) corresponding to a choice indifference between the two options ($B \sim AC$, vertical line in **e**), with the existence of a p_A for which $B > AC$ and of a different p_A for which $AC > B$ (vertical dashed lines). The value of α was identified by fitting a softmax function (Eq. 2, red line) to the proportion of AC choices (blue dots). **f, g**, Compliance and violation. Choice pattern compatible with the continuity axiom (**f**) and possible axiom violations (**g**). Red dots represent the proportion of AC choices when $p_A = 0$ or 1, corresponding to the axiom's initial requirement ($A > B$ and $B > C$, implying $P(A > B) > 0.5$ and $P(C > B) < 0.5$). **h–j**, Consistency of choice behavior. The standardized β coefficients from logistic regressions of single trials' behavior (**h**) showed that the main choice-driving variables were reward magnitude (m_R , m_L) and probability (p_R , p_L) for all animals, both for left (L) and right (R) choices; previous trial's chosen side (preCh_R) and reward (preRew_R) did not consistently explain animals' choices (error bars indicate 95% CI across sessions). $*p < 0.05$ (one-sample t test, FDR-corrected). No. of sessions per animal: 100 (A), 81 (B), 24 (C), 15 (D). In choices between options with different probability of delivering the same reward magnitude, the better option was preferred on average by all animals, demonstrating compliance with FSD (**i**) (error bars indicate binomial 95% CI; no. of tests per animal: 28 (A), 24 (B), 15 (C), 23 (D); average no. of trials per test: 12 (A), 13 (B), 11 (C), 34 (D)). In choices between sure rewards (bars represent average across all sessions; gray dots represent single sessions; error bars indicate binomial 95% CI) animals preferred A to B, B to C, and A to C (**j**), complying with both WST and SST (WST: proportion of choices of the better option > 0.5 (blue dashed line); SST: proportion of A over C choices (red line) \geq other choice proportions).

monitor. The monkey indicated the preferred option within 2 s by moving the joystick to the side of one option, at which time the unselected option disappeared. After holding the joystick for at least 1 s, the reward corresponding to the selected option was delivered (see Fig. 1a). Visual cues were presented on a blank screen, indicating the amount (magnitude) and probability of receiving a reward (fruit juice) through white horizontal lines: each line's vertical position indicated a reward amount, while the line width was proportional to the probability of obtaining that reward (see Fig. 1b).

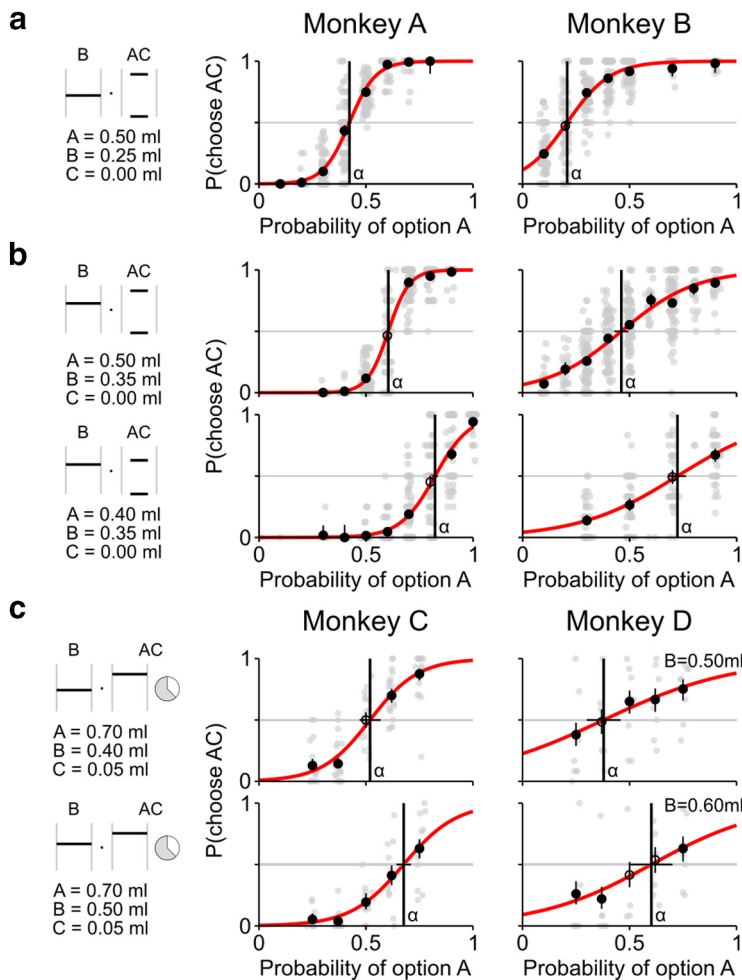


Figure 2. Experimental test of the continuity axiom. *a–c*, Compliance with the continuity axiom. The axiom was tested through choices between a gamble B and a varying AC-combined gamble (left, visual stimuli for an example choice pair with $p_A = 0.5$ (*a,b*) or $p_A = 0.375$ (*c*)); increasing p_A values resulted in gradually increasing preferences for the AC option. In each plot: Gray dots represent the proportion of AC choices in single sessions. Black circles represent the proportions across all tested sessions. Vertical bars represent the binomial 95% CIs. Filled circles represent significant difference from 0.5 (binomial test, $p < 0.05$). The tests were repeated using different A and B values (*b*) as well as non-zero C values in a modified task (*c*). All 4 animals complied with the continuity axiom by showing increasing preferences for increasing probability of gamble A (rank correlation, $p < 0.05$), with the AC option switching from nonpreferred ($P(\text{choose AC}) < 0.5$) to preferred ($P(\text{choose AC}) > 0.5$) (binomial test, $p < 0.05$). Each IP (α , vertical line) was computed as the p_A for which a data-fitted softmax function (Eq. 2) had a value of 0.5 (horizontal bars represent 95% CI); α values shifted coherently with changes in A and B values in all 4 animals, indicating a continuous MP trade-off relation. For single sessions' IP values, see Figure 3.

Task design for Monkeys C and D. The reward magnitude was represented by the vertical position of a horizontal black bar within a vertically oriented white rectangle. The probability of a reward was conveyed through a circular stimulus, presented adjacent to the bar stimulus, composed of two sectors distinguished by black-white shading at horizontal and oblique orientation; the amount of horizontal shading indicated the probability of obtaining the cued reward magnitude. On each trial, the animal made a choice between two gambles, one of which was a degenerate gamble ($p = 1$), presented randomly in left-right arrangement on the monitor. For risky gambles, the cued reward magnitude could be obtained with $P = \text{cued probability}$ and a fixed small reward (0.05 ml) could be obtained with $p = 1 - \text{cued probability}$.

Each trial started when the background color on the touch screen changed from black to gray. To initiate the trial, the animal was required to place its hand on an immobile, touch-sensitive key. Presentation of the gray background was followed by presentation of an ocular fixation spot (1.3° visual angle). After 500 ms, both choice options appeared in left-right arrangement on the monitor, followed after 750 ms by appearance of two

blue rectangles below the choice options at the margin of the monitor, close to the position of the touch-sensitive key. The animal was then required to touch one of the targets within 1500 ms to indicate its choice. Once the animal's choice was registered, the unchosen option disappeared, and after a delay of 500 ms, the chosen object also disappeared and a liquid reward was given to the acting animal. Reward delivery was followed by a trial-end period of 1000 ms, which ended with extinction of the gray background.

Logistic regression

To identify the key variables driving choice, we analyzed single trials' data from each session using the following logistic regression:

$$\text{Logit}(P_R) = \beta_0 + \beta_1 m_L + \beta_2 p_L + \beta_3 m_R + \beta_4 p_R + \beta_5 \text{preCh} \cdot \text{preRew} + \varepsilon \quad (1)$$

where P_R is the probability of choosing the right-side option; m and p represent the reward magnitude and probability of options, respectively, presented on the left (L) or right (R) side of the screen; preCh represents the previous trial's choice (−1 for left-side and 1 for right-side choices) while preRew corresponds to the reward magnitude obtained in the previous trial; the product $\text{preCh} \cdot \text{preRew}$ thus increases for larger rewards obtained when choosing the right-side option in the previous trial; ε is the error term. We obtained estimates of the regression coefficients (β_i) for each session. Coefficients were standardized by multiplying each β with the ratio of the corresponding independent variable's SD over the SD of the predicted variable. Standardized β s were tested for statistical significance through one-sample t test.

To explicitly investigate variability across sessions, we repeated this analysis on the complete trials dataset (not separated by session), and compared the regression accuracy between the given model (Eq. 1) and a second model, which included session number as a random variable. We compared results from the two models (fixed effect and mixed effect) using a likelihood ratio test ($p < 0.05$) and computed the pseudo- R^2 as a goodness-of-fit measure to compare the accuracy of the two models.

Axioms of EUT

The axioms of EUT are necessary and sufficient conditions for choices to be described by the maximization of EU: if the axioms are fulfilled, a subjective value corresponding to the EU can be assigned to each choice option, and the option with the highest EU is chosen (von Neumann and Morgenstern, 1944).

- Formally,
- I. **Completeness:** $\forall A, B$ either $A \succ B$, $B \succ A$, or $A \sim B$
 - II. **Transitivity:** $A \succ B, B \succ C \Rightarrow A \succ C$
 - III. **Continuity:** $\forall A \succ B \succ C, \exists! p \in (0, 1)$ such that $pA + (1 - p)C \sim B$
 - IV. **Independence:** $\forall A \succ B \Rightarrow pA + (1 - p)C \succ pB + (1 - p)C; \forall C, \forall p \in (0, 1)$

Where A, B, and C are gambles corresponding to known probability distributions over outcomes, " \succ " is the preference relation, and " \sim " represents indifference. The operation $pA + (1 - p)C$ corresponds to combining the two gambles A and C with probabilities p and $(1 - p)$ respectively, thus representing itself a gamble different from A or C alone.

The continuity of preferences (axiom III) can also be expressed as follows:

III-a. *Monotonicity*: $\forall A \succ B \Rightarrow A \succ \alpha A + (1 - \alpha)B \succ B; \forall \alpha \in (0, 1)$
 III-b. *Archimedean property*: $\forall A \succ B \succ C, \exists! p_1, p_2 \in (0, 1): p_1A + (1 - p_1)C \succ B \succ p_2A + (1 - p_2)C$

Such an alternative expression (III-a, III-b) does not include any equality (i.e., IP) and is thus better suited for experimental hypothesis testing compared with III.

Complete (I) and transitive (II) preferences are necessary for univocally and consistently ranking all choice options, representing a “weak ordering” condition. In this case, each possible choice option can be given a specific rank level, so that an option with higher rank will be preferred to one with lower rank. Although these rank levels can be defined as numbers, they have no cardinal meaning: any monotonic transformation of these values would still represent preferences. Such rank levels would give no information about the strength of preferences and could not predict choices between options defined as combinations of gambles.

Conversely, if preferences are also continuous (axiom III), they can have a meaningful numerical utility representation. Thus, if A is preferred to B, the utility of option A (U_A , a real number) is larger than the utility of B (U_B) and vice versa if $U_A > U_B$, option A is preferred over option B:

$$A \succ B \iff U_A > U_B$$

The independence axiom (IV) allows to go one step further, defining how to compute the utility of any gamble G from its attributes (magnitudes m_i and associated probabilities p_i) as follows:

$$U_G = EU(\{m_i, p_i\}) = \sum_i U(m_i) \cdot p_i$$

making it possible to predict choices between any possible choice options.

Expected utility theorem

Following the four axioms, the EU theorem states that given any two options A and B, A will be preferred to B if and only if the EU of A is larger than the EU of B as follows:

$$A \succ B \iff EU(A) > EU(B)$$

where

$$EU(X) = \sum_i U(m_i) \cdot p_i$$

with X representing a gamble with outcomes m_i and associated probabilities p_i and $U(m)$ representing the utility associated with the magnitude m . The EU of a gamble thus corresponds to the average utility of a gamble, weighted by the reward probabilities, representing the subjective equivalent to the objective (mathematical) $EV = \sum_i m_i \cdot p_i$.

The EU theorem links preferences to subjective evaluations: if option A is preferred to option B, the EU of option A will be greater than the EU of option B; vice versa, if the EU of A is greater than the EU of B then A will be preferred to B.

Lexicographic preferences

Lexicographic preferences represent a possible violation of the continuity axiom. Lexicography refers to the way words are ordered based on their component letters: the first letter defines which word comes first in the dictionary, unless words have the same first letter in which case the second letter will define the order, and so on. In choice theory, lexicographic preferences correspond to a decision strategy where the preference for one option is only based on one attribute, while a second attribute is considered only when the first attribute has the same value in both options. In risky choices, the attributes of an option correspond to reward magnitude and probability; in this context, lexicographic preferences imply that the option with the highest magnitude would always be

chosen, independent of its probability, unless the two options had the same magnitude, in which case the option with the highest probability would be chosen. Inverting the roles of magnitude and probability would also result in lexicographic choices (see Fig. 1g).

Lexicographic preferences, while complying with the completeness and transitivity axioms, represent a violation of the continuity axiom. They imply that reward magnitude and probability are not combined into a subjective value, indicating an underlying choice mechanism (and its neural implementation) incompatible with EUT and with the concept of utility.

Testing deterministic axioms

To experimentally test an axiom, which is an absolute rule that must hold for any possible gamble, it is necessary to test the largest possible number of different cases. We thus generalized our results by repeating the continuity test using different initial gambles: in one set of tests A, B, and C were defined as sure rewards (degenerate gambles) varying over the range 0-0.9 ml; in a different set of tests, B was defined as a probabilistic two-outcome gamble; in a final set of tests A, B, and C were all defined as two-outcome gambles, resulting in the AC option being a three-outcome gamble.

The EUT axioms were originally defined as deterministic rules, which assume that preferences do not change over time. In order to account for the variability in choice behavior (repeated choices between the same pair of options can yield different results), we interpreted the axioms in a stochastic sense: option A was considered preferred to option B when the proportion of A over B choices was larger than 0.5 ($P(A \succ B) > 0.5$; binomial test, $p < 0.05$).

Stochastic transitivity

We tested two stochastic forms of the transitivity axiom: weak stochastic transitivity (WST) and strong stochastic transitivity (SST).

In repeated presentations of the same choices, compliance with WST corresponds to choosing A over B, B over C, and A over C in >50% of the trials as follows:

$$\text{WST: } P(A \succ B) > 0.5, P(B \succ C) > 0.5 \Rightarrow P(A \succ C) > 0.5$$

WST thus represents a simple extension of the transitivity rule to stochastic choices.

SST further requires that preference of A over C must be stronger than the other two preferences. Interpreting the proportion of choices for one option as indicating the strength of preference for that option relative to another option, SST can be defined as follows:

$$\text{SST: } P(A \succ B) > 0.5, P(B \succ C) > 0.5 \Rightarrow P(A \succ C) > \max\{P(A \succ B), P(B \succ C)\}$$

We tested the compliance with WST and SST in choices between different reward amounts. Each test was identified by a triplet of reward magnitudes (degenerate gambles A, B, and C), which were then presented in a randomly intermingled sequence of repeated choice pairs (A vs B, B vs C, or A vs C). We used a range of possible values for A, B, and C rewards: in Monkeys A and B, reward A included values between 0.15 and 0.50 ml (in 0.05 ml increments), reward B was 0.05, 0.15, 0.20, 0.25, 0.35, or 0.45 ml, and reward C was 0, 0.05, 0.15, 0.25, or 0.35 ml. In Monkeys C and D, reward A had values between 0.30 and 0.90 ml, reward B was between 0.30 and 0.90 ml, and reward C was between 0.10 and 0.60 ml (all in 0.10 ml increments). The total number of tested triplets per monkey was as follows: 14 (A), 48 (B), 4 (C), 84 (D).

Testing the continuity axiom

We implemented a behavioral and statistical test of the continuity axiom as follows:

We defined three starting gambles and verified that monkeys complied with the transitivity axiom; this allowed us to define A, B, and C as the most preferred, middle, and least preferred gamble, respectively.

In each trial, monkeys chose between the middle gamble B and a probabilistic combination (AC) of the most and least preferred gambles: $AC = p_A A + (1 - p_A)C$, where p_A is the probability of obtaining the most preferred gamble.

We statistically tested compliance with the axiom according to definitions III-a and III-b: we defined a series of AC combinations with specific probabilities (p_A between 0 and 1, in 0.1 increments) and measured the proportion of choices for the AC option (P_{AC}). We then tested whether monkeys preferred the middle gamble to the AC combination for at least one p_A ($P_{AC} < 0.5$; binomial test, $p < 0.05$, FDR-corrected) while also preferring the AC combination to the middle gamble in at least one case ($P_{AC} > 0.5$), as required by the Archimedean property (III-b); compliance with the Monotonicity rule (III-a) was ensured by the tested compliance with FSD, and tested in each continuity test by showing increasing preferences for increasing probability of gamble A (rank correlation, $p < 0.05$).

After testing for the existence of an IP (α), its numerical value was determined by fitting a softmax function to the choice data through nonlinear least squares fit. The softmax function was defined as follows:

$$P_{AC}(p_A) = 1/(1 + \exp(-(p_A - \alpha)/\tau)) \quad (2)$$

where τ (softmax “temperature” parameter) represents the steepness of the preference function (steeper for lower τ values).

Data fitting of IPs

To obtain a set of curves approximating the ICs, for each middle gamble B, we fitted the IPs corresponding to varying gamble A magnitudes, using three different functions as follows:

$$\text{Linear} : f(x) = ax + b$$

$$\text{Power} : f(x) = a(x - c)^b$$

$$\text{Hyperbolic} : f(x) = a + b/(x - c)$$

where x represents the reward magnitude, $f(x)$ the IC, that is, the reward probability as a function of reward magnitude. A nonlinear least squares method was used to minimize the error in the probability domain (x axis in Figs. 4, 6b)

Economic choice models

We modeled the probability of choosing one option using a standard discrete choice model. The probability of choosing gamble A in choices between any two gambles (choice set: {A,B}) was defined through a binary logistic model as follows:

$$P(A|\{A, B\}) = 1/(1 + \exp(-(V_A - V_B)/\tau)) \quad (3)$$

where V_A and V_B represent the subjective values of gamble A and B, respectively, τ the temperature parameter. The gamble value was defined following EUT as $V = EU = \sum_i U(m_i) \cdot p_i$, with utility $U(m)$ being a parametric function of reward magnitude m .

A maximum likelihood estimation (MLE) procedure was used to estimate the free parameters (τ and utility-function parameters) to best approximate the measured proportions of choices.

The MLE procedure involved computing and maximizing the log-likelihood (LL) in the parameters space.

$$\begin{aligned} LL(\text{param}|\text{choices}) &= \sum_{\text{trials}} A_{\text{chosen}} \cdot \log(P(A, \text{param})) \\ &+ (1 - A_{\text{chosen}}) \cdot \log(1 - P(A, \text{param})) \end{aligned}$$

where the sum is defined across all trials in one session, A_{chosen} takes the value of 1 if gamble A was chosen in one trial, zero otherwise, and $P(A, \text{param})$ is the discrete choice model defined above with parameters param . We minimized the negative LL using the *fminsearch* Matlab function.

We defined three possible utility functions as follows:

$$\text{Linear utility} : U(m) = \frac{m}{m_0}$$

$$\text{Power utility} : U(m|a) = \left(\frac{m}{m_0}\right)^a$$

$$\text{S-shaped utility} : U(m|a, b) = \exp\left(-b \left(-\log \frac{m}{m_0}\right)^a\right)$$

with $m_0 = 0.5ml$, representing the maximum reward magnitude, thus normalizing all utility functions between 0 and 1.

A linear utility function could only explain choices based on the EV of the options: it would perform as the best model only if monkeys were choosing by comparing the objective, mathematical EV of the options. A power utility function would instead be able to describe choices with a specific risk preference: either risk seeking or risk aversion. Finally, an S-shaped utility function could accommodate a more complex pattern of risk attitudes, with the possibility of both risk seeking and risk aversion for different reward magnitudes. As the S-shaped function, we used the two-parameter *Prelec* function, which is typically used as a PW function, but can also represent a plausible shape for the utility function.

Using the same binary logistic model with a different definition of the gamble value allowed us to test models from different economic choice theories. In prospect theory (PW model), the gamble value is defined as $V = U(m) \cdot w(p)$; we used the two-parameter *Prelec* function as the PW function $w(p)$. In the additive model, we defined a gamble's value as $V = \omega_m \cdot U(m) + \omega_p \cdot p$, with ω_m and ω_p as additional free parameters acting as “weights” for the magnitude and probability terms. According to the mean-variance approach, the value definition does not rely on the utility concept: $V = EV + \beta \cdot \text{Risk}$, where EV and Risk are the first two moments of the gamble's probability distribution, and β is a free parameter. We computed Risk as the EV of the squared deviation from the mean: $\text{Risk} = \sum_i p_i \cdot (m_i - EV)^2$.

In order to construct the full indifference map predicted by a model, for each IC, we numerically computed the IPs corresponding to finely spaced magnitude levels: for a selected model (using the average recovered parameters across all sessions), the subjective value of the B gamble was computed (V_B); after increasing the magnitude by 0.001 ml (starting from the B gamble magnitude), the subjective value was then computed for a series of probabilities (step 0.001), and the probability corresponding to the value closest to V_B was identified as the IP. This procedure, repeated for all B gambles, allowed us to obtain a distance measures between all the modeled and measured IPs, in the probability domain, which was used as one of the quantities for model comparison.

To compare the six tested models (EUT with three possible utility functions, PW, additive model, and mean-variance), we defined four accuracy metrics: (1) the square root of the mean squared error, representing the average distance between modeled and measured IPs, in probability units; (2) the Bayesian information criterion (BIC); and (3) the Akaike information criterion (AIC), both introducing a penalty term when increasing the number of model parameters; the variance in the differences of modeled versus measured preferences (i.e., the proportion of AC vs B choices across all continuity tests). For each of these four measures, a lower value represented a better model compared with a higher value.

Results

Rationale

In EUT, decisions are modeled “as if” subjects had an internal utility representation, making no assumptions about the brain processes underlying choice (Camerer, 2008). To investigate the possibility of EUT describing the actual neuronal mechanisms of choice, our approach is to (1) verify that subjects follow the model's assumptions, (2) infer the utility measure defined in EUT, which is not directly measurable, and (3) identify the neuronal

substrates coding such quantity, if they exist. To fulfill the first point, we need to verify compliance with the assumptions of EUT (i.e., the axioms). If the assumptions are satisfied, the utility measure can be elicited following econometric methods. These crucial steps identify the subjective quantities, as opposed to the objective, physical ones, that can be used to describe preferences. The third point, which represents the ultimate goal of our research, involves the identification of utility-coding neuronal substrates by correlating the neuronal activity with the utility measure rigorously defined in the previous points.

Here, we focused on the initial step: testing the basic assumptions of EUT to infer the existence of a utility measure. We based our experimental tests on the third EUT axiom, the continuity axiom, which is crucial for defining a numerical utility measure. The fourth axiom (independence) is needed to mathematically define the subjective value as an option's EU. Violations of the independence axiom have been reported in human and animal experiments, highlighting the limits of EUT (Allais, 1953; Kahneman and Tversky, 1979; Battalio et al., 1985; Conlisk, 1989). Nevertheless, the continuity axiom remained a necessary condition in all major generalized EU theories developed since the 1940s, which share the axiom's main implication, that is, the definition of a scale of numerical subjective values (Harless and Camerer, 1994; Starmer, 2000; Jehle and Reny, 2001).

We conducted our study on macaque monkeys, a close evolutionary neighbor to humans, who are able to perform skilled problem solving after relatively short training periods. For each animal, it was possible to elicit hundreds of choices per day, for several months. This allowed us to systematically investigate a broad range of choice problems, with a fine resolution in the reward magnitude and probability domains, which is difficult to achieve in typical economic studies on humans. In particular, the idea of testing an axiom only makes sense if we are able to test compliance with it in an extensive and comprehensive set of choice situations. Indeed, this is what drove our definition of the choice sets in the current experiment: extending the axiom test to gradually more complex choice cases, with the highest attainable resolutions. As a downside, because of the high experimental costs and long durations of primate research, we achieved a lower number of participants compared with typical human studies. We believe that this limitation was counterbalanced by the robustness of the results that were obtained in each animal, thanks to the elevated number of repetitions and choice problems tested. The currently developed systematic choice task should stimulate future neuronal investigations of decision variables in primates using well-established formal economic theory.

Design

To test the continuity axiom of EUT in nonhuman primates, we trained 4 monkeys to perform a binary choice task. In each trial, the animal chose between two options, presented simultaneously on a computer monitor (Fig. 1*a*), offering liquid rewards varying in amount and probability (Fig. 1*b*). The continuity axiom states that, given any three ranked gambles (A, B, and C, ranging widely), a decision maker should be indifferent between the middle gamble (B) and a probabilistic combination of the two other gambles (AC). Formally,

$$\forall A \succ B \succ C, \exists! \alpha \in (0, 1) \text{ such that } \alpha A + (1 - \alpha) C \sim B \quad (4)$$

where \succ defines a preference relation and \sim defines indifference; α is the specific probability associated to gamble A for

which indifference occurs. The axiom should be satisfied for any arbitrary set of gambles A, B, and C.

The axiom was originally defined as a “plausible continuity assumption” (von Neumann and Morgenstern, 1944). Thought experiments intuitively clarified how the continuity axiom could be violated (Georgescu-Roegen, 1954; Levin, 2006; Chateaufort et al., 2008), for example, when options had infinitely different values. Nevertheless, the axiom was considered a reasonable condition and not experimentally tested.

To experimentally test the axiom, we first defined three gambles (Fig. 1*c*) for which the monkey had well-defined preferences ($A \succ B$ and $B \succ C$ in the majority of trials; binomial test, $p < 0.05$). We then combined the most and least preferred gambles (A and C, respectively) with probability p_A (varying between 0.1 and 0.9 in 0.1 steps), obtaining the family of gambles $AC(p_A)$ (Fig. 1*d*). Finally, we presented choices between B and one of the AC combinations and tested for the existence of indifference between B and a probabilistic combination of gambles A and C, with probability $p_A = \alpha$ such that $B \sim \alpha A + (1 - \alpha) C$ (Fig. 1*e*). Compliance with the continuity axiom would thus be demonstrated by the existence of a unique α between 0 and 1 (Fig. 1*f*), while violations would occur if α were not identifiable or when multiple α existed (Fig. 1*g*).

Therefore, we varied the behavioral test with the A, B, and C gambles in several ways: (1) we varied the safe reward amounts of the degenerate gambles A, B, and C between tests (see paragraph Compliance with the continuity axiom); (2) we used a risky B gamble but kept varying the safe reward amounts of the degenerate gambles A and C between tests (see paragraph ICs in the MP space); and (3) we used only risky A, B, and C gambles and varied A and C between tests (see paragraph Continuity axiom test in the Marschak-Machina triangle). The first, more basic manipulation (1) was tested in 4 monkeys, while the further two, more specific variations (2 and 3) were tested in only 2 of the animals (Monkeys A and B).

We used pseudo-random repetitions of all presented choice pairs to account for the stochasticity of choice behavior and as a basis for future recordings of neuronal activity. Although the EUT axioms were originally defined as deterministic rules (von Neumann and Morgenstern, 1944), we extended their definition to the stochastic domain (Hey and Orme, 1994; Loomes, 2005) because the statistical analysis of neuronal responses requires the use of multiple trials. Therefore, we made our design compatible with basic assumptions of stochastic choice theories (Luce, 1959; McFadden and Richter, 1990; McFadden, 2005). We quantified the animal's choice from the probability of choosing one option over its alternative, rather than using traditional single-shot economic tests.

Basic choice behavior

We investigated the consistency of choice behavior to make sure that the 4 tested monkeys understood the reward-cue associations and were able to express their preferences.

To assess the contribution of magnitude and probability information to decisions, we performed a logistic regression on single trials' choice data, using the chosen side as the dependent variable and the options' probabilities and magnitudes as independent variables. An additional regressor controlled for the effect of past trials: the product of the previous trial's chosen side and obtained reward (see Logistic regression, Eq. 1). Standardized regression coefficients (β) corresponding to reward magnitude and reward probability were significantly different from zero (one-sample *t* test,

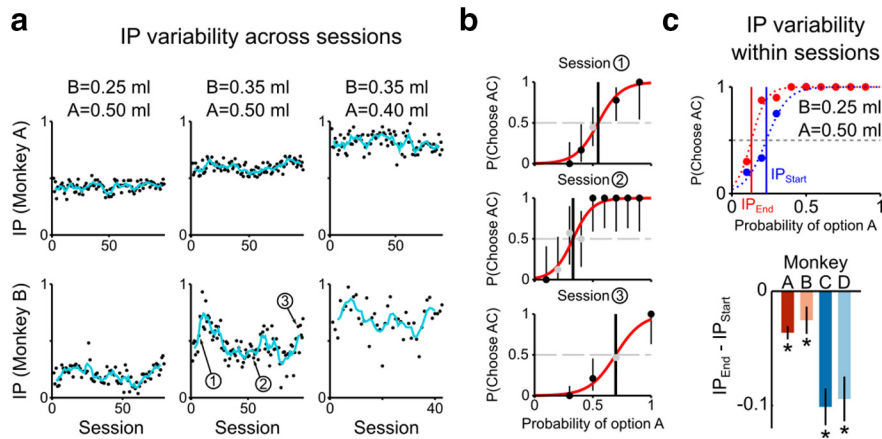


Figure 3. Variability of IP across and within sessions. **a**, Evolution of the IP as a function of session number (blue curve indicates 5-session moving average), for the three continuity axiom tests presented in Figure 2a, b (Monkeys A, B). **b**, Choice proportion for the AC gamble in three example sessions (as indicated in **a**). Vertical bars represent binomial 95% CI, indicating a significant preference (black dot) when not crossing the horizontal dashed line. **c**, Top, Preferences in the first half (blue) and in the second half (red) of a session (one example session, Monkey B). Bottom, Average IP change across all sessions. Error bars indicate SE. * $p < 0.05$, IP difference significantly different from zero (t test).

$p < 0.05$, FDR-corrected) in all 4 animals (Fig. 1h), indicating that both variables were choice-driving factors. Compared with such coefficients (average absolute value across animals: 0.54 ± 0.18 SD), the past trials' β was much smaller (average absolute value: 0.032 ± 0.025 SD) and not consistently significant across animals, confirming that choices were mainly driven by the two cued attributes. A significant intercept implied a side bias for Monkeys A, B, and C (one-sample t test, FDR-corrected; per animal: $p = 1.7 \times 10^{-10}$ (A), $p = 5.1 \times 10^{-3}$ (B), $p = 7.3 \times 10^{-5}$ (C), $p = 4.1 \times 10^{-1}$ (D)). The side bias was accounted for by presenting each option on both sides of the screen, the same number of times.

As a direct test of consistent choice behavior, we verified compliance with first order stochastic dominance (FSD). FSD is the probabilistic analog of “more is better,” and represents a basic requirement of EUT and of the continuity axiom in particular (see Axioms of EUT): a gamble should be preferred if it contains outcomes at least as good as another gamble, with at least one strictly better outcome. FSD implies that an option with a more probable outcome should be preferred to one with a less probable outcome of the same reward amount. The higher probability gamble stochastically dominates the lower probability gamble and should thus be preferred. Because of choice stochasticity, a number of dominated choices are naturally expected, but to comply with FSD their proportion must be significantly < 0.5 . We tested FSD in choices between a gamble and a safe option as well as between two gambles, using reward magnitudes (fixed for each presented pair of options) between 0.1 and 0.9 ml and reward probabilities ranging from 0.05 to 0.97 (step 0.02, Monkeys A and B) and from 0.25 to 0.75 (step 0.125, Monkeys C and D). We found that all 4 animals complied with FSD by preferring the dominant option in $> 50\%$ of trials across all FSD tests (Fig. 1i; binomial 95% CI > 0.5). We inferred from this behavioral compliance with FSD that the animals attributed higher reward value to higher reward probability, as prerequisite for testing the integration of reward probability and magnitude with the continuity axiom.

A further prerequisite for testing the continuity axiom is compliance with the completeness and transitivity axioms. Completeness ensures that subjects have well-defined preferences for any presented pair of options. In line with general notions of discrete choice models (McFadden, 2001), in every trial, our choice set had a finite number of offered alternatives

(two) with mutually exclusive (only one option could be selected) and collectively exhaustive options (the set included all possible options). Animals were thus induced to express complete preferences. Still, they could choose not to select any option, avoiding expressing a preference, which would violate the completeness axiom. This was not consistently observed, except rarely for low-valued options pairs (which were excluded from subsequent testing). Thus, we tested the animals' choices while they complied with the completeness axiom.

The transitivity axiom ensures that all gambles can be univocally ranked. In line with stochastic choice theory, we tested two stochastic forms of transitivity, weak (WST) and strong (SST) stochastic transitivity (see Stochastic transitivity), using combinations of the A, B, and C magnitudes ranging from 0 to 0.5 ml (step: 0.05 ml) for Monkeys A and B, and from 0 to 0.9 ml for Monkeys C and D (step: 0.1 ml). In choices from all tested triplets, the 4 animals complied with both WST and SST (Fig. 1j). Individual transitivity tests revealed compliance with WST in all 141 tested magnitude combinations and compliance with SST in 125 (89%) tested triplets (average number of trials per test, per animal: 21 (A), 96 (B), 36 (C), 105 (D)). This compliance with the transitivity axiom indicated that the animals made consistent choices and thus ranked the tested gambles unequivocally.

Compliance with the continuity axiom

Following the formal definition of the continuity axiom (Eq. 4), we assessed the existence of a unique IP in choices between a fixed gamble and a probabilistic combination of the other two gambles. We defined three degenerate gambles A, B, and C with three different reward amounts; in each trial, the animal chose between the middle gamble (B) and the probabilistic combination of the most and least preferred gambles (A and C, respectively). Thus, we tried to obtain a p_A at which choice indifference occurs: $\alpha = p_A$ such that $B \sim p_A(A) + (1 - p_A)C$.

All 4 animals preferred the middle gamble to at least one of the AC combinations, while also preferring at least one of the AC combinations to the middle gamble (Fig. 2): for different p_A values, the proportion of choices for the AC combination was significantly $<$ or $>$ 0.5 (binomial test, $p < 0.05$) (Fig. 2a,c), following an increase in preference with increasing p_A (rank correlation, $p < 0.05$). Such a switch of revealed preference depending on probability p_A indicated the existence of a unique IP and thus

compliance with the continuity axiom (see Testing the continuity axiom).

We defined the A, B, and C gambles as degenerate gambles of varying reward magnitudes. All tested triplets showed a pattern of AC preferences compatible with the continuity axiom: the existence of both preferred and nonpreferred AC combinations, together with gradually increasing preferences of the AC option, implied the presence of an IP. Monkeys complied with the continuity axiom when defining the C gamble as 0 ml (Monkeys A and B; Fig. 2*a,b*) as well as when using a non-zero C gamble (Monkeys C and D; Fig. 2*c*) in a different task (see Task design), thus confirming the robustness of our results.

Importantly, different triplets of gambles produced IPs varying in a meaningful and consistent manner: increasing only the reward magnitude of the middle gamble (B) (e.g., from 0.25 to 0.35 ml; Fig. 2*a,b*) produced larger α values; decreasing the magnitude of the most preferred gamble (A) (e.g., from 0.50 to 0.40 ml; Fig. 2*b*) resulted in higher α values (nonoverlapping 95% CI, $p < 0.05$). Such a pattern reflected the notion, central to the continuity axiom, of α being a measure of the subjective value of the middle gamble: the more B was considered close to A in value, the higher its α ; the closer B was considered to C, the lower its corresponding α . While shifting consistently for different initial gambles, the α values were also different across animals, denoting the subjective quality of the measured IPs. In conclusion, testing the continuity axiom showed a coherent pattern of IPs, highlighting the joined contribution of reward magnitude and probability to the definition of subjective values.

Lexicographic preferences represent a possible continuity axiom violation (Fig. 1*g*). Lexicography refers to the way words are ordered based on their component letters. In analogy, lexicographic preferences in risky decision-making correspond to choices based on one component at a time (either reward magnitude or probability). They represent a specific choice heuristic in which the gamble components are considered separately and are not combined into a single quantity. This corresponds to a choice mechanism incompatible with the definition of numerical subjective values: lexicographic choices cannot be described by assigning a numerical value to each gamble, as in EUT (see Lexicographic preferences). By showing the existence of a coherent set of IPs, our data demonstrated that preferences were not lexicographic, implying that animals considered and combined magnitude and probability information.

Although compliance with the axiom was robust when considering the average IPs, there appeared to be a clear variability of IPs across sessions. When looking at the variation of the IPs over time (Fig. 3*a*), we found that they showed a combination of a random variation and linear trend periods. This was evident in Monkey B's choice behavior, where different sessions produced clearly different IP values for the same continuity test (Fig. 3*b*). This resulted from values gradually varying by session, reflecting gradual changes in the risk attitude over long periods of time

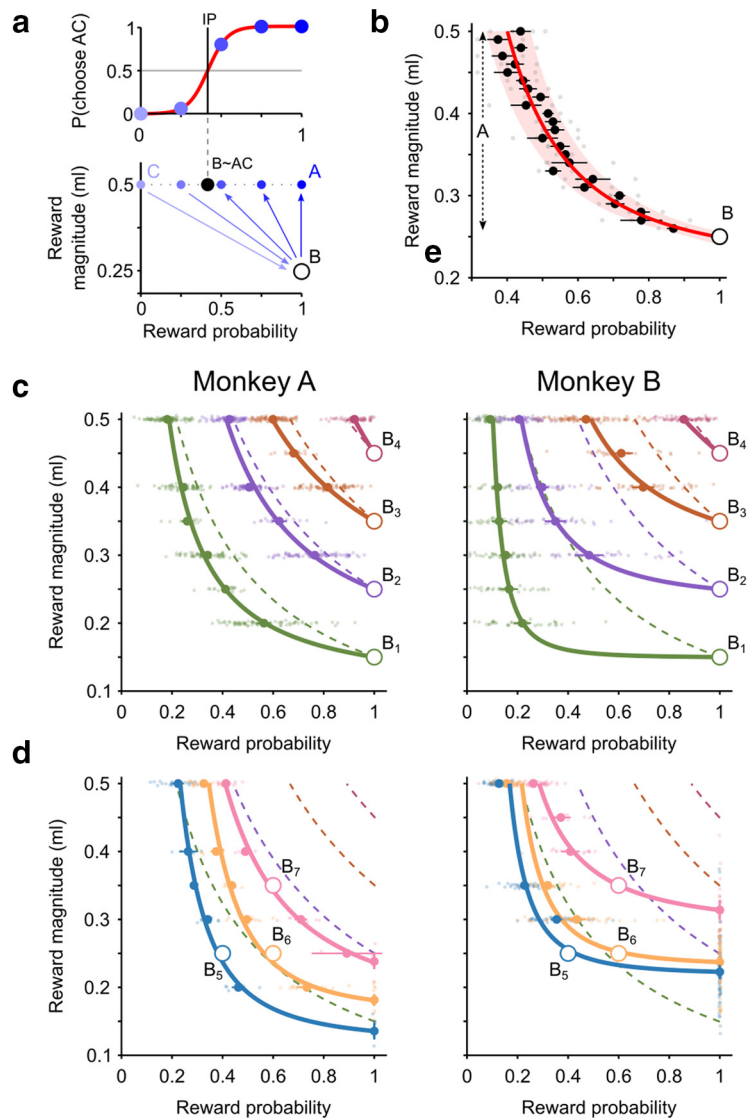


Figure 4. ICs in the MP space. *a*, Representation of the continuity axiom test in the MP space. The gambles used for testing the axiom can be mapped into the MP diagram. Preference in choices between B (circle) and combinations of A and C (graded blue dots) is represented by an arrow pointing in the direction of the preferred option (bottom), consistently with the proportion of choices for the AC option (top). Each continuity axiom test resulted in an IP (vertical line, top), represented as a black dot in the MP space (bottom). *b*, ICs. ICs (gray dot indicates single session; black dot indicates average across sessions; error bar indicates SE) obtained using different A values (step 0.01 ml) shifted continuously, producing an IC in the MP space (area: 95% CI). Curve indicates best fitting power function. Data from Monkey A (5 sessions, 1781 trials). *c, d*, Indifference map. ICs for different B values (colored thick curves) described the gradual variation of the average IPs (colored dots, with SE bars) for each B. ICs were modeled as power functions (for comparison with hyperbolic and linear fits, see Extended Data Figure 4-1). Small dots indicate IPs measured in single sessions. Both sure rewards (*c*) and probabilistic gambles (*d*) as B options, produced coherent indifference maps, with smooth and non-overlapping ICs. Dashed curves indicate points with the same objective EV, highlighting the subjective quality of ICs.

(days/months), possibly because of factors external to the experiment-defined variables (including learning of the reward-stimuli association and adaptation to uncontrolled variables). Importantly, the observed variability did not affect the ability to robustly identify the IP in a single session (Fig. 3*b*), a necessary requirement for the possibility of correlating such behavioral measures with the activity of single neurons or neuronal populations.

By including the session number as a random variable to the previously defined logistic regression, we compared the performance of a mixed-effect model (including both fixed and random effects) with that of the basic fixed effect one (Eq. 1). Adding the

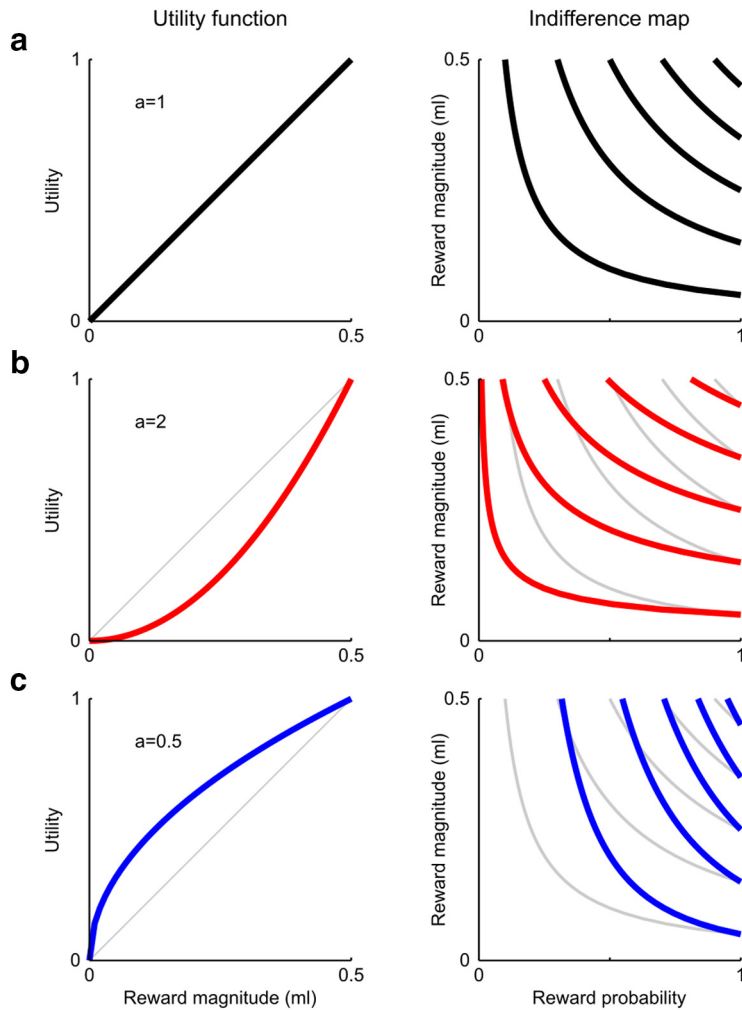


Figure 5. Theoretical relation between utility function and indifference map. Sample indifference maps obtained from different utility functions, with a representing the single parameter of a power function ($U(m) = (m/m_0)^a$, with $m_0 = 0.5\text{ml}$). The IC for a degenerate gamble B was analytically obtained as a function of magnitude values, from the equation $EU_B = U(m) \cdot p$. By solving the equation for p , the IC equation can be obtained as follows: $p(m) = (m_B/m)^a$, where m_B is the magnitude of a degenerate gamble B. The three utility shapes are directly related to different risk attitudes: risk neutrality for linear utility (a), risk seeking for convex utility (b), and risk aversion for concave utility (c). In all plots, gray curves indicate risk neutrality (i.e., choices based on the objective EV of gambles). The indifference map globally warps according to the risk attitude.

random effect resulted in a better description of choices (likelihood ratio test, $p < 0.05$), confirming that random variations of choice-driving parameters were observed across sessions. Nevertheless, the increase in the goodness of fit compared with the model not containing the random effects was only marginal (difference in pseudo- R^2 between the mixed-effect model and the fixed effect model, per monkey: 0.01 (A), 0.02 (B), 0.02 (C), 0.01 (D)) with no evident differences in the regression coefficients.

When looking at the variability within single sessions, by measuring the IP difference between the last half and the first half of each session, we found that the IPs significantly decreased over one session's duration (paired t test, $p < 0.05$) in all 4 monkeys (Fig. 3c). This result suggested the contribution of the total fluid intake as one of the factors affecting IPs, although its effect size (mean IP difference, per monkey: -0.03 (A), -0.02 (B), -0.10 (C), -0.10 (D)) was smaller than the smallest intervals of probability values used (0.10 for Monkeys A, B; 0.125 for Monkeys C, D).

Overall, these results support the core ideas arising from the continuity axiom: subjective values, which define preferences, are quantities (numbers) that depend on reward magnitudes and are modulated by reward probabilities. In other words, probabilities modify the subjective reward values in a graded and continuous way; a variation in reward magnitude can be compensated by a change in reward probability and vice versa, establishing a continuous trade-off relation between magnitudes and probabilities.

ICs in the MP space

To confirm the existence of a continuous trade-off relation between reward magnitudes and probabilities, as implied by the continuity axiom, we represented the animals' IPs (measured through a softmax fit, Eq. 2) in a two-dimensional diagram with reward magnitude and probability as variables. Such MP space was used to represent the continuity axiom tests, conducted in Monkeys A and B, in which the B gamble is either a degenerate gamble or a true (nondegenerate) gamble, with degenerate A and C gambles ($C = 0$ ml). Each gamble used in a continuity test (B and AC combinations) corresponded to a single point in the MP space (Fig. 4a). Compliance with the continuity axiom was manifested as choice indifference between the B gamble and an AC combination, identifying a single point in the MP space ($B \sim AC$ in Fig. 4a).

To test compliance with the continuity axiom for an extended set of degenerate gambles, we held the B gamble fixed but varied the magnitude of the A gamble. This test yielded a set of IPs that lined up as an IC. Importantly, there were no discontinuities ("jumps") in the IC while varying the A magnitude in 0.01 ml steps. This was confirmed through a leave-one-out regression, by measuring the percentage of left-out data points falling within the CI identified by all other points: 96% of average IPs (24 of the 25) and 94% of single session's IPs (72 of 77) fell within the 95% CI of the nonlinear IC regression (power function, see below). Such very gradual change in the IP fulfilled a fundamental requirement of the continuity axiom: as the magnitude of the A gamble increased, IPs gradually decreased without any apparent discontinuity (Fig. 4b).

Repeating the IC elicitation procedure for different degenerate B gambles yielded a set of ICs (i.e., an indifference map), which captured the full pattern of relations between reward magnitudes and probabilities. To measure each animal's indifference map, we performed 14 continuity axiom tests, by systematically varying the magnitudes of gambles A and B between 0.15 and 0.50 ml in 0.05 ml steps. For each middle gamble (B_1 to B_4), we varied the value of gamble A, obtaining a set of IPs in each session (average sessions per continuity test, per animal: 64 (A), 48 (B)), thus confirming the compliance with the continuity axiom for a large set of A and B magnitudes. We modeled the resulting IC through a power function, which we identified as the best fitting function compared with linear and hyperbolic ones

(Extended Data Fig. 4-1). The fitted IC followed the gradual shift in IPs observed when varying the reward magnitude of gamble A. The indifference map, obtained by including ICs corresponding to all tested B gambles, captured the full pattern of relations between IPs, highlighting their smooth and continuous transitions (Fig. 4c).

As the EUT axioms should apply to any arbitrary set of gambles, we further tested compliance by using a set of truly risky B gambles (B_5 to B_7). These 14 experimental tests involved choices between pairs of probabilistic gambles with no option of getting a sure reward, making it a more general and more complex choice situation (average sessions per continuity test, per animal: 12 (A), 40 (B)). Nevertheless, IPs were still consistently observed, and the resulting ICs had qualitatively similar shapes to the ones involving a degenerate gamble (Fig. 4d).

Together, these results confirm compliance with the continuity axiom in a broad class of choice situations and highlight the existence of an orderly trade-off relation between reward magnitudes and probabilities: even a small decrease in reward magnitude was compensated in revealed preference by an increase in reward probability and vice versa.

Economic modeling of ICs

We investigated whether our results were compatible with theoretical economic models of choice in the framework of EUT, particularly in relation to the existence of a utility function able to represent choices in agreement with the EU theorem. According to EUT, a gamble's value stems from the product of the reward's utility and its associated probability; this assumes the existence of a utility function over magnitude values, which fully defines the shape of the whole indifference map, uniquely identifying the subjective MP trade-off relation (Fig. 5). Assuming a linear utility function results in choices depending only on the objective quantities: the EU model incorporates the objective EV model, which represents the objectively optimal preference pattern (Fig. 5a), as a special case.

We estimated the utility function using single-trial choices from each session, through an MLE method. We defined a discrete choice model in standard fashion (McFadden, 2001), with the probability of choosing one option described by a logistic function (Eq. 3), dependent on the difference in EU between the two options. Each gamble's EU was computed as the utility of the reward multiplied by its probability (see Economic choice models).

We compared MLE results from three utility functions: linear, power, and S-shaped. The power utility function captured the monkeys' choice behavior better than the linear one (difference in BIC: 51.7 ± 39.1 SD, $p = 2.9 \times 10^{-18}$, Monkey A; 54.9 ± 29.5 SD, $p = 1.7 \times 10^{-27}$, Monkey B; one-sample t

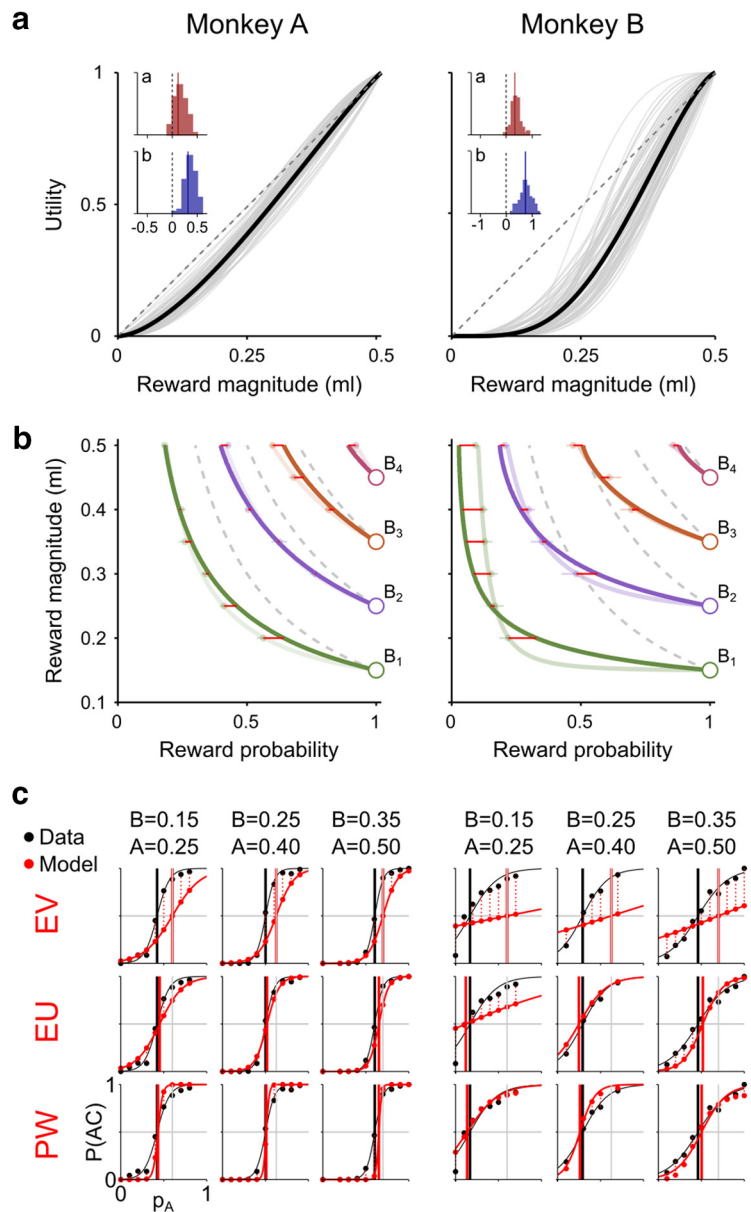


Figure 6. ICs are compatible with economic models. **a**, Utility functions. Single sessions' utility functions (gray) and averages (black) estimated through MLE using single-trial choice data. The two estimated parameters (a and b in inset, histograms of log values) were both significantly positive, indicating S-shaped utility functions. **b**, EUT-predicted ICs. Indifference map reconstructed using the estimated utility functions. Light-colored curves indicate the measured indifference map (see Fig. 4c). Red horizontal lines indicate the distance between measured and modeled IPs. Dashed gray curves indicate points with equal EV, corresponding to a linear utility model. For the theoretical relation between the utility function and the ICs, see Figure 5. **c**, Comparison of modeled and revealed preferences. Percentage of choices for the AC gamble ($P(AC)$) measured (black) and modeled using three models (red), for three example A-B-C triplets (top, A and B in ml, $C = 0$ ml). The EV model could only predict IPs equal to the EVs (gray vertical lines), with a larger error in the prediction of the $P(AC)$ (vertical dotted red lines) compared with the EU model. The PW model, which included a subjective PW, was better at capturing the revealed preferences only in specific cases (e.g., $B = 0.15$, $A = 0.25$ in Monkey B).

test), whereas the S-shaped utility function outperformed the power-shaped one (BIC difference: 12.2 ± 12.3 SD, $p = 2.1 \times 10^{-13}$, Monkey A; 8.7 ± 12.2 SD, $p = 1.4 \times 10^{-8}$, Monkey B). The two recovered parameters for the S-shaped utility functions (Fig. 6a, histograms) were both significantly different from one ($p < 10^{-15}$ in both monkeys; one-sample t test), confirming that utility functions were nonlinear and had a significant inflection point (i.e., a change in curvature), resulting in an S-shaped curve.

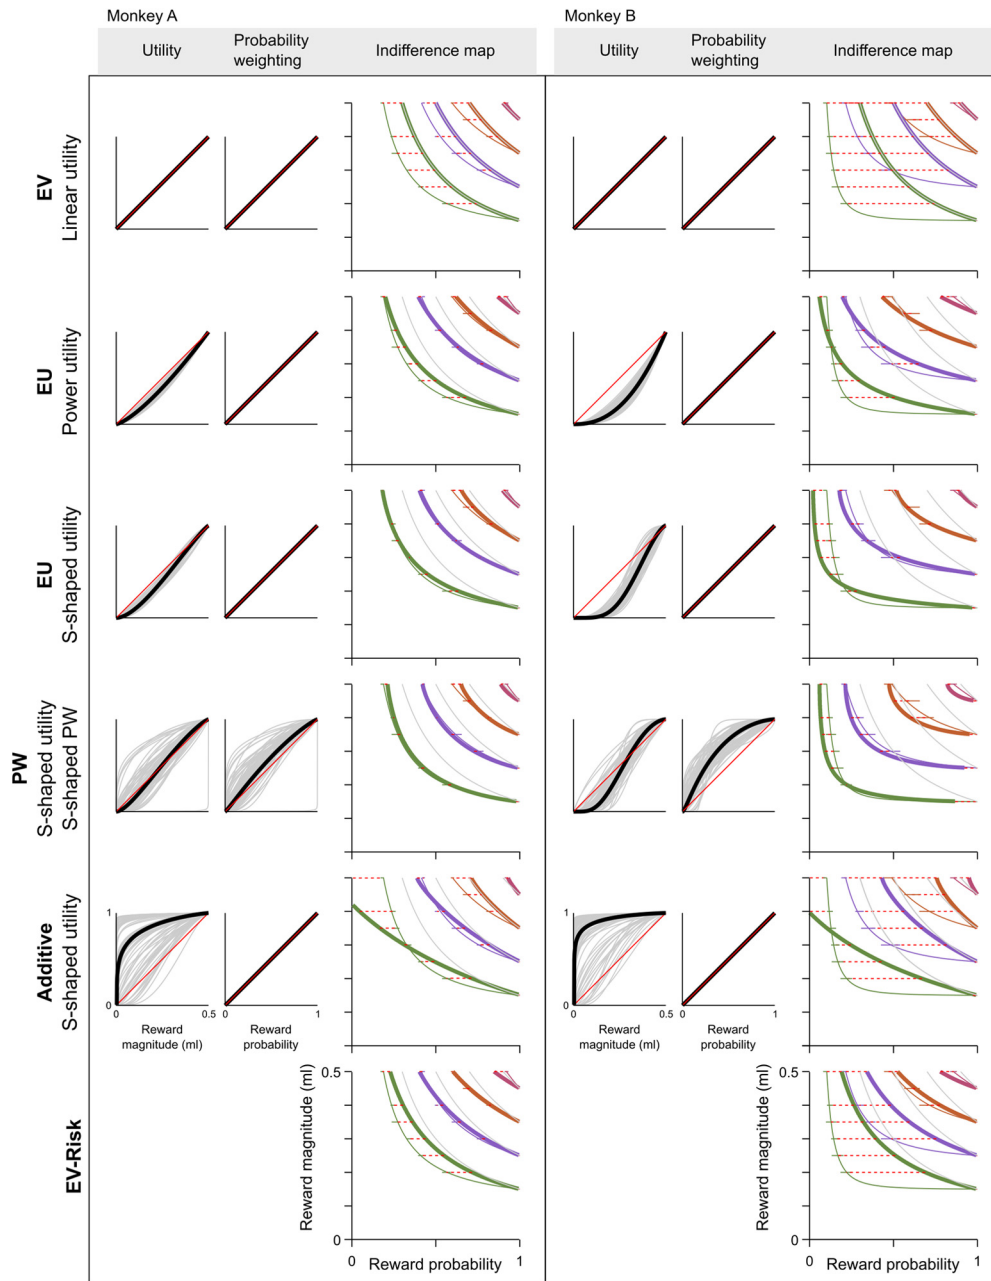


Figure 7. Comparison of economic models. Recovered utility function, PW function, and corresponding indifference map for each economic model (rows). Gray curves indicate single-session estimates. Black curves indicate the corresponding means, plotted by averaging the recovered parameters across all sessions. Red lines indicate linear utility and PW functions, for comparison. The EV, EU, and additive models assume linear PW. The mean-variance model (EV-Risk) does not have a utility representation. Other conventions and symbols are as in Figure 6b.

We used the recovered S-shaped utility functions (Fig. 6a) to construct the corresponding indifference map: for each gamble B, we computed its EU and obtained an IC as the set of points with equal EU in the MP space. It was thus possible to define a whole indifference map using a single utility function. Such a map, modeled from the MLE-estimated utility function, closely matched the behavioral IPs and the previously fitted ICs (Fig. 6b), which had been measured for each B gamble independently and had no link to the economic theory. The average distance between the modeled IPs and the behavioral IPs (Fig. 6b, red lines) was smaller for the EU model than for the objective-EV model (Fig. 6b, dashed curves) (square root of the mean squared error: 0.028 (EU model), 0.108 (EV model) for Monkey A; 0.052 (EU model),

0.274 (EV model) for Monkey B). Thus, the EU model was better at capturing the shape of the indifference map compared with the objective EV model by 3.9 times in Monkey A and 5.3 times in Monkey B. We quantified the ability of the model to describe the actual preferences (proportion of choices for the AC option), from which the IPs were calculated, using the variance in the deviation between predicted and measured proportion of choices (Fig. 6c, vertical dotted lines). A lower average variance for the EU model indicated that it was better at describing preferences compared with the EV model (*Var*; Extended Data Fig. 6-1). This was also confirmed using a standard model comparison analysis (*BIC* and *AIC* scores; Extended Data Fig. 6-1). These results indicate that the nonlinearity in the utility function was able to

capture the subjective quality of the IPs (Fig. 6*b*) and of the revealed preferences (Fig. 6*c*).

Nonlinear PW is a further subjective factor explaining economic choices, as proposed in prospect theory and other generalized EU theories (Harless and Camerer, 1994; Hey and Orme, 1994). We found that a model incorporating utility and weighted probabilities (PW model) improved the description of the measured IPs by 4.5 times in Monkey A and 4.4 times in Monkey B, compared with the EV model. Therefore, the PW model had similar descriptive power compared with the EU model. Overall, the PW model outperformed the EU model in 3 of 4 model-accuracy metrics, in both monkeys (Fig. 7; Extended Data Fig. 6-1), suggesting that, in the tested choice situation, adding the subjective weighting of probabilities marginally improved the description of preferences compared with EUT, representing a possible refinement to our EU model for describing preferences (Fig. 6*c*).

We explored the possibility of an additive model being able to represent the animals' choice behavior. Additive models, in which magnitude and probability information are normalized and added to each other (see Economic choice models), have been shown to account for choices under uncertainty (i.e., when risk is not explicitly known) in both humans and monkeys (Farashahi et al., 2019). Fitted to our data, the additive model was outperformed by the other utility-based models (Fig. 7; Extended Data Fig. 6-1), confirming the finding that risky choices are better explained by a multiplicative model rather than by an additive one.

The mean-variance approach, an alternative economic model that approximates EUT without relying on the concept of utility (Levy and Markowitz, 1979), defines a gamble's value as the sum of the corresponding EV and risk components (see Economic choice models). When fitting a mean-variance model to our data, the ICs could not be predicted as well as with any of the utility-based models (Fig. 7; Extended Data Fig. 6-1), with an improvement in the IC description over the EV model of 1.6 (Monkey A) and 1.4 (Monkey B) times, well below the performance of utility-based models.

In support for the existence of a utility-compatible mechanism producing the indifference map, we investigated the variation of IPs across sessions. We computed the Pearson's correlation coefficient (ρ) for all possible pairs of IPs. A significant ρ , both within each IC (one-sample *t* test, per animal: $p = 1.8 \times 10^{-5}$ (Monkey A); $p = 5.4 \times 10^{-4}$ (Monkey B)) and across different ICs ($p = 8.2 \times 10^{-5}$ (Monkey A); $p = 1.5 \times 10^{-3}$ (Monkey B)), confirmed that the variation of each IP was associated with a variation of other IPs (average ρ , per animal: 0.19 ± 0.28 SD (A); 0.15 ± 0.26 SD (B)). Across sessions, the indifference map changed shape as a whole: IPs were not varying independently from each other, but were linked by a common underlying root, identifiable as the utility function.

In conclusion, the economic modeling of ICs and the correlation among IPs support the idea of choices resulting from a utility maximization process: the product of a subjectively defined utility

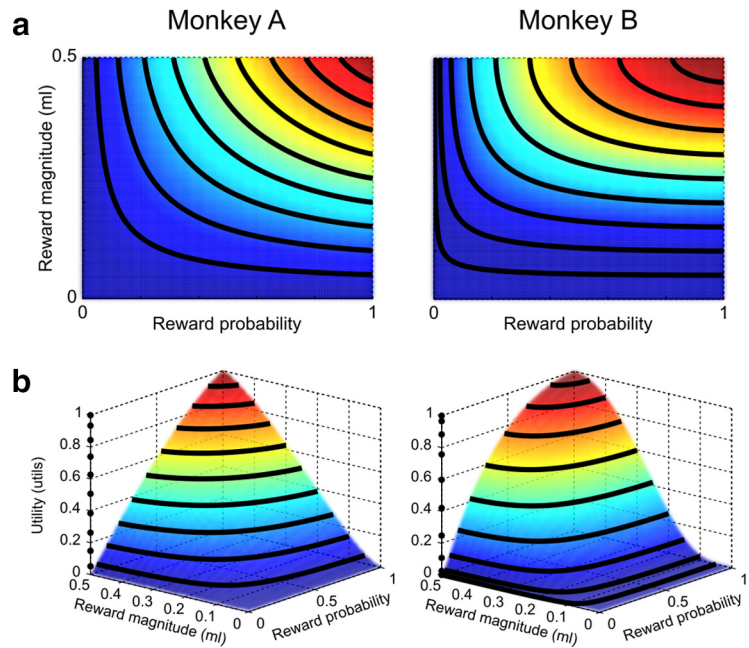


Figure 8. Three-dimensional utility representation. *a*, Indifference map resulting from the best-fitting PW model. Black curves indicate the iso-utility points (i.e., the ICs) for 9 equally spaced EV levels (0.05–0.45 ml, step: 0.05 ml). Colors (blue to red) represent utility values (from 0 to 1). *b*, Three-dimensional representation of utility values, highlighting the continuous relation between objective quantities (m , p) and subjective values (utility): for any two-outcome gamble, corresponding to a $\{m, p\}$ pair, a utility level can be mathematically computed using the two-dimensional function defined by the best-fitting PW model ($V(m, p) = U(m) \cdot w(p)$). This relation can be used to identify utility-coding neurons: the activity of a neuron coding utility should follow the three-dimensional surface across the whole tested region of the MP space.

function with reward probabilities (possibly subjectively weighted) is able to describe the choice behavior between any pair of gambles and in particular the smooth trade-off relation between reward magnitudes and probabilities (Fig. 8).

Continuity axiom test in the Marschak-Machina triangle

The Marschak-Machina triangle (Marschak, 1950; Machina, 1982) has been extensively used in economic studies of human behavior for evaluating and comparing different generalized EU theories (Camerer, 1989; Sophor and Gigliotti, 1993). This approach graphically displays continuity tests by showing IPs in choices between test gamble B and probabilistic AC combinations containing multiple possible outcomes.

We further tested the continuity axiom using A, B, and C gambles defined as two-outcome gambles, which resulted in AC combinations being three-outcome gambles. To present such gambles to the animal, we used visual cues with three horizontal lines, which simultaneously represented all possible reward outcomes and their probabilities (Fig. 9*a*, inset). The Marschak-Machina triangle represents gambles with three fixed outcome magnitudes (defined in our experiment as 0, 0.25, and 0.5 ml) and any combination of associated probabilities (p_1 , p_2 , and p_3 , defined as the probabilities associated with the low, middle, and high outcome magnitudes, respectively). The x and y coordinates correspond to the probability of obtaining the lowest (p_1) and highest (p_3) outcome, respectively (Fig. 9*a*).

We defined A as a gamble with 0.25 ml and 0.5 ml as possible outcomes, whereas both B and C had 0 ml and 0.25 ml as possible outcomes; AC combinations then corresponded to gambles with three possible reward magnitudes: 0, 0.25, and 0.5 ml (Fig. 9*b*). In the Marschak-Machina triangle, gamble A lay on the y

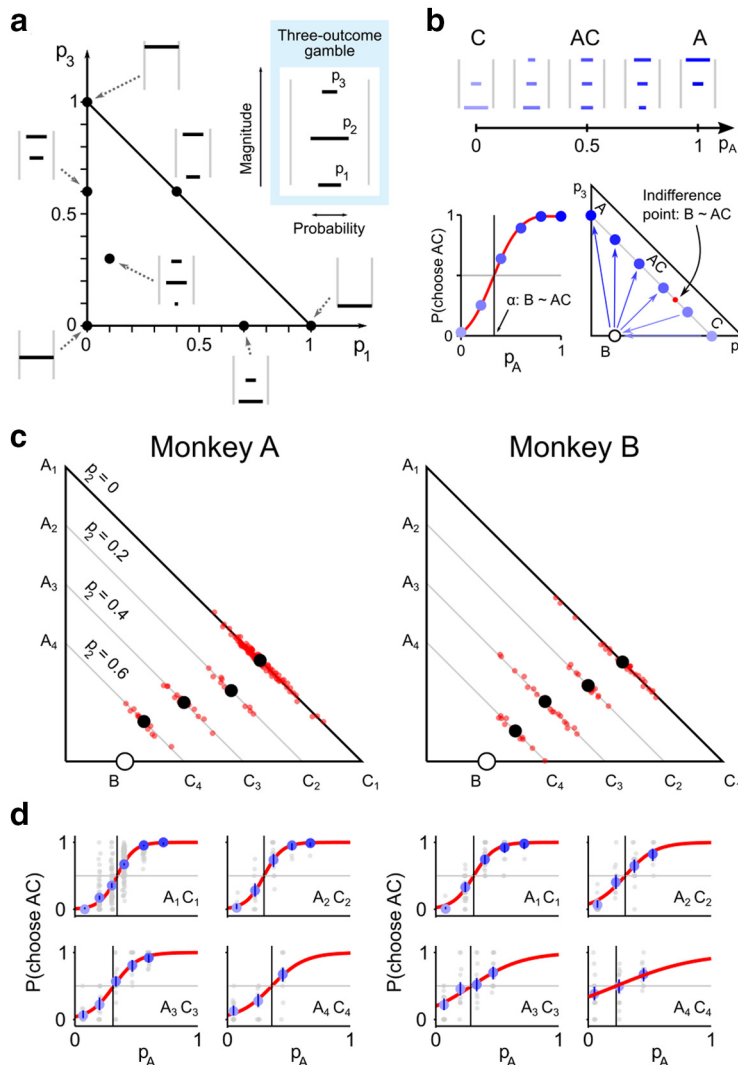


Figure 9. Continuity axiom test in the Marschak-Machina triangle. **a**, Three-outcome gambles. Gambles with three fixed outcome magnitudes and any combination of outcome probabilities can be represented in the Marschak-Machina triangle. The visual cue (inset) for three-outcome gambles included three horizontal lines indicating the three possible outcome magnitudes (vertical position) and the respective probabilities p_1 , p_2 , and p_3 (line width). **b**, Scheme of the continuity axiom test. Three-outcome gambles used to test the axiom (top) can be represented in the Marschak-Machina triangle (bottom right) together with the B gamble (circle) and the resulting IP (red dot). Arrows point toward the preferred option, consistently with the proportion of AC choices (bottom left). **c**, Continuity axiom test in the Marschak-Machina triangle. Average IPs (black dots) and from single sessions (red dots) were consistently elicited, indicating compliance with the continuity axiom in choices between two- and three-outcome gambles. **d**, Revealed preferences in choices between two- and three-outcome gambles. Average measured percentage of AC choices as a function of the probability of obtaining the A option (graded blue dots). Each average IP (black vertical line) corresponds to a black dot in **c**. Other symbols are as in Figure 3.

axis while gambles B and C lay on the x axis. Consequently, the $AC(p_A)$ combinations lay inside the triangle, on a straight line between A and C, the position between bottom right and top left being proportional to the probability p_A (Fig. 9b, bottom). Satisfaction of the continuity axiom would be manifested as a point on the line between A and C where the animal is indifferent between the B gamble and the AC combination (Fig. 9b, bottom, labeled $B \sim AC$).

We defined four pairs of A and C gambles (A_1C_1 to A_4C_4 , associated with increasing probability of the middle outcome (p_2) between $p_2 = 0$ and $p_2 = 0.6$, in 0.2 increments); for each A-C pair, we tested the continuity axiom using a fixed middle gamble B, for a total of four tests (Fig. 9c). Results showed the

existence of IPs in all tested cases (Fig. 9d), confirming compliance with the continuity axiom in choices between two- and three-outcome gambles. Because reward magnitudes are fixed in the Marschak-Machina triangle, while probabilities vary across the full range, the pattern of IPs confirmed the role of reward probabilities as modifiers for the EU: a gradual change in A-C (in terms of p_2) led to a continuous increase in IPs (in terms of the probability of the highest outcome, p_3), also demonstrating the possibility of constructing ICs within the Marschak-Machina triangle (Fig. 9c).

Through the unique graphical representation of the Marschak-Machina triangle, we showed that monkeys complied with the continuity axiom in choices involving three-outcome gambles, thus supporting the idea of a choice mechanism based on numerical subjective values also in more complex choice scenarios.

Discussion

This study demonstrates compliance of monkey behavior with the continuity axiom of EUT, implying a MP trade-off relation and determining a numerical utility measure able to describe choices. Our results build on studies reporting monkeys' preferences for gambles with larger EV to those with smaller EVs derived from both reward magnitude and probability (Musallam et al., 2004; Lau and Glimcher, 2005; Averbach, 2015; Farashahi et al., 2018). The interpretation of an apparent trade-off between reward magnitude and probability in these studies relies on the validity of EUT axioms in general and on the continuity axiom in particular. By showing compliance with the continuity axiom, we demonstrate that the trade-off is systematic and quantitatively complies with theoretically ideal choices. In addition, the compliance with the continuity axiom fulfills a necessary condition for the validity of neuronal coding of subjective value and formal economic utility (Platt and Glimcher, 1999; Tremblay and Schultz, 1999; Padoa-Schioppa and Assad, 2006; Kobayashi and Schultz, 2008; Lak et al., 2014; Stauffer et al., 2014). These results demonstrate compliance with the continuity axiom for biologically plausible reward magnitudes and probabilities. The animals' trade-off choices followed closely the IPs modeled by theoretical choice functions, suggesting the existence of an axiomatically defined utility measure for choice options' values. The pattern of subjective IPs allowed us to define a utility function that is suitable for investigating neuronal mechanisms of economic choice according to the rigorous definitions of economic theory.

The continuity axiom, a necessary condition for the existence of a numerical utility, states that given three subjectively ordered gambles, a decision maker will be indifferent between the middle gamble and a probabilistic combination of the two other

gambles. We experimentally tested the continuity axiom in choices between a two-outcome gamble and a safe option. Four monkeys exhibited a choice behavior consistent with the continuity axiom, making choices compatible with the existence of a unique IP. We generalized our results to more complex choice situations in 2 monkeys, confirming compliance with the axiom in choices between two- and three-outcome gambles, representable in the Marschak-Machina triangle. We showed how the IPs identified through the axiom test procedure could be interpreted as subjective evaluations of the choice options and used to construct an indifference map. Such a map revealed a congruent, subjective trade-off relation between reward magnitudes and probabilities, which supported the idea of choices being the result of a utility maximization process compatible with EUT.

The four axioms of EUT represent the necessary conditions for the existence of a precisely defined utility quantity. In particular, the continuity axiom permits the definition of a numerical utility, whereas the independence axiom defines how to compute the utility measure. In our quest for investigating a utility-based brain mechanism driving human decisions, we need to clarify if and to what degree the economic theories are generalizable across primates. Although the continuity axiom has not been tested in human subjects, it is accepted as a reasonable condition. On the other hand, humans have been shown to violate the independence axiom of EUT, which led to the creation of alternative economic choice theories. Although it is still unknown whether nonhuman primates violate the independence axiom similarly to humans, as a first step we showed that they comply with a more basic assumption, the continuity axiom. By sequentially testing the EUT axioms, we can verify up to which point their behavior can be described by the economic theory, and if monkeys' preferences reflect the characteristics of human decision-making. This approach can shed light on the existence of a common choice mechanism across primates.

Past studies have shown that decisions in monkeys reflect both the magnitude and probability information of the choice options (Platt and Glimcher, 1999; Lak et al., 2014), leaving two open questions: how are magnitude and probability, two physical quantities, transformed into subjective quantities? And how are such quantities combined into a single value? These questions naturally extend to the neurophysiological domain. The axiomatic approach allows to investigate such points, clarifying with a robust procedure if the gambles' dimensions are subjectively combined as mathematically defined by modern economic decision theories.

Lexicographic preferences and other classes of choice heuristics represent known continuity axiom violations. By showing compliance with the continuity axiom, we could exclude an important class of heuristics (the lexicographic rules) as the driving mechanism for choices in the tested situation. This ensured that all presented information were used to make actual multiattribute choices: reward probability and magnitude were both considered and combined when evaluating the options. However, we did not test other heuristic decision strategies. For example, a recent study observed a win-stay/lose-switch strategy, which only contributed marginally to single-trial choices while possibly influencing the long-term learning of values and probabilities (Ferrari-Toniolo et al., 2019). Thus, further tests should delimit the viability of continuity satisfaction in different choice situations, while exploring alternative heuristic-based models able to describe the observed choice patterns.

An important consequence of continuity axiom violations is that subjective preferences cannot be described by assigning numerical values to the options. IPs, and hence indifference maps,

would be undetermined if the axiom is violated, as would the concept of utility as a subjective measure of an option's value. An EUT model could still be fitted to the data to recover a utility function; yet, such a function would not have the intended meaning of expressing the options' subjective values. When the axiom is fulfilled, instead, we showed how the IPs could be expressed numerically as utility, and the resulting indifference map could be generated through S-shaped utility functions. Having a utility representation of values allows for the assignment of a specific numerical value to each IC. The activity of a neuron encoding the options' subjective values should comply with the indifference map: it should be proportional to the elicited numerical utility levels across ICs, while remaining constant within each IC (see Fig. 8).

The utility function is not considered to be fixed over time: its shape and range have been theorized to adapt and to change based on internal and external factors. Multiple factors might contribute to the shaping of each subject's utility function in different ways. Some of the observed variability in our utility measure might be explained as an adaptation to changes in the task or in the external environment, as suggested by the gradual changes observed in IPs across sessions. This calls for more sophisticated models integrating longer term learning effects on the formation and the evolution of utility functions with experience.

According to theories relying on the continuity axiom, utilities are combined with probability information to give a gamble's EU. The exact form of such combination remains to be tested: the independence axiom is required to define exactly how utilities and probabilities combine into EU. Although we did not yet explicitly test compliance with the independence axiom in monkeys, we observed that nonlinear weighting of reward probabilities resulted in a marginally better description of choice behavior compared with EUT. Therefore, the present study points to nonlinear PW as a possible refinement to EUT in monkeys, compatible with the human experimental results that led to the development of prospect theory (Kahneman and Tversky, 1979). The Marschak-Machina triangle framework could be used to directly investigate compliance with the independence axiom in monkeys, allowing for the quantitative investigation of the neural underpinnings of several generalized EU choice theories.

In conclusion, by explicitly testing the continuity axiom, we verified that, in the tested situation, the choice mechanism was compatible with the computation of finite, numerical utilities, gaining crucial information on the plausible mechanisms guiding choices toward the maximization of utility.

References

- Allais M (1953) Le Comportement de l'Homme Rationnel devant le Risque: critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* 21:503.
- Averbeck BB (2015) Theory of choice in bandit, information sampling and foraging tasks. *PLoS Comput Biol* 11:e1004164.
- Battalio RC, Kagel JH, MacDonald DN (1985) Animals' choices over uncertain outcomes: some initial experimental results. *Am Econ Rev* 75:597–613.
- Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Ser B* 57:289–300.
- Camerer CF (1989) An experimental test of several generalized utility theories. *J Risk Uncertainty* 2:61–104.
- Camerer CF (2008) Neuroeconomics: opening the gray box. *Neuron* 60:416–419.
- Chateaufneuf A, Cohen M, Tallon JM (2008) Decision under risk: the classical expected utility model. In: *Decision-making process*, pp 365–384. London: ISTE Ltd.

- Conlisk BJ (1989) Three variants on the Allais example. *Am Econ Rev* 79:392–407.
- Farashahi S, Azab H, Hayden B, Soltani A (2018) On the flexibility of basic risk attitudes in monkeys. *J Neurosci* 38:4383–4398.
- Farashahi S, Donahue CH, Hayden BY, Lee D, Soltani A (2019) Flexible combination of reward information across primates. *Nat Hum Behav* 3:1215–1224.
- Ferrari-Toniolo S, Bujold PM, Schultz W (2019) Probability distortion depends on choice sequence in rhesus monkeys. *J Neurosci* 39:2915–2929.
- Georgescu-Roegen N (1954) Choice, expectations and measurability. *Q J Econ* 68:503.
- Harless BY, Camerer CF (1994) The predictive utility of generalized expected utility theories. *Econometrica* 62:1251–1289.
- Hey JD, Orme C (1994) Investigating generalizations of expected utility theory using experimental data. *Econometrica* 62:1291–1326.
- Jehle GA, Reny PJ (2001) *Advanced microeconomic theory*. Boston: Addison-Wesley.
- Johnson EJ, Ratcliff R (2014) *Neuroeconomics*. Amsterdam: Elsevier.
- Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47:263–292.
- Kobayashi S, Schultz W (2008) Influence of reward delays on responses of dopamine neurons. *J Neurosci* 28:7837–7846.
- Lak A, Stauffer WR, Schultz W (2014) Dopamine prediction error responses integrate subjective value from different reward dimensions. *Proc Natl Acad Sci USA* 111:2343–2348.
- Lau B, Glimcher PW (2005) Dynamic response-by-response models of matching behavior in rhesus monkeys. *J Exp Anal Behav* 84:555–579.
- Levin J (2006) Choice under uncertainty. Lecture Notes. Stanford University. Available at <https://web.stanford.edu/~jdlevin/Econ%20202/Uncertainty.pdf>.
- Levy H, Markowitz HM (1979) Approximating expected utility by a function of mean and variance. *Am Econ Rev* 69:308–317.
- Loomes G (2005) Modelling the stochastic component of behaviour in experiments: some issues for the interpretation of data. *Exp Econ* 8:301–323.
- Luce RD (1959) *Individual choice behavior: a theoretical analysis*. New York: Wiley.
- Machina MJ (1982) Expected utility analysis without the independence axiom. *Econometrica* 50:277.
- Marschak J (1950) Rational behavior, uncertain prospects, and measurable utility. *Econometrica* 18:111–141.
- McFadden DL (2001) Economic choices. *Am Econ Rev* 91:351–378.
- McFadden DL (2005) Revealed stochastic preference: a synthesis. *Econ Theory* 26:245–264.
- McFadden DL, Richter MK (1990) Stochastic rationality and revealed preference. In: *Preferences, uncertainty, and optimality: essays in honor of Leonid Hurwicz* (Chipman JS, McFadden DL, Richter MK, eds), pp 163–186. Boulder, CO: Westview.
- Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA (2004) Cognitive control signals for neural prosthetics. *Science* 305:258–262.
- Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223–226.
- Pascal B (2004) *Pensées* [translated by Ariew R], pp 1658–1662. Indianapolis: Hackett.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.
- Samuelson PA (1948) Consumption theory in terms of revealed preference. *Economica* 15:243–253.
- Sopher B, Gigliotti G (1993) A test of generalized expected utility theory. *Theor Decis* 35:75–106.
- Starmer C (2000) Developments in nonexpected-utility theory: the hunt for a descriptive theory of choice under risk. *Adv Behav Econ* 38:332–382.
- Stauffer WR, Lak A, Schultz W (2014) Dopamine reward prediction error responses reflect marginal utility. *Curr Biol* 24:2491–2500.
- Tremblay L, Schultz W (1999) Relative reward preference in primate orbitofrontal cortex. *Nature* 398:704–708.
- Tversky A (1969) Intransitivity of preferences. *Psychol Rev* 76:31–48.
- von Neumann J, Morgenstern O (1944) *Theory of games and economic behavior*. Princeton NJ: Princeton UP.
- Weber M, Camerer C (1987) Recent developments in modelling preferences under risk. *OR Spektrum* 9:129–151.