Compliance with the continuity axiom of Expected Utility Theory supports utility maximization in monkeys

Simone Ferrari-Toniolo*[1], Philipe M. Bujold[1], Fabian Grabenhorst[1], Raymundo Báez-Mendoza[1,2], Wolfram Schultz[1]

Affiliations:

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

**Corresponding author*

Email addresses

Simone Ferrari-Toniolo: simone.ferraritoniolo@gmail.com

Wolfram Schultz: Wolfram.Schultz@protonmail.com

Short title: Testing the continuity axiom of Expected Utility Theory in monkeys

Abstract

Expected utility theory (EUT), the first axiomatic theory of risky choice, describes choices as a utility maximization process: decision makers assign a subjective value to the choice options, and choose the option with the highest subjective value. This description can be obtained for every subject that complies with the four axioms of EUT. The continuity axiom, central to EUT and to its modifications, requires decision makers to be indifferent between a gamble and a specific probabilistic combination of a more preferred and a less preferred gamble. Compliance with the axiom is necessary for the definition of numerical subjective values. We experimentally tested the continuity axiom for a broad class of gamble types in four monkeys, showing that their choice behavior complied with the existence of numerical subjective values. We used the numerical quantity defined by 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 numerical quantities.

INTRODUCTION

In risky choices, we should select the option with the highest expected return in order to maximize our profits in the long term. Since human and animal subjects consistently violate this principle, to better describe real-world choices economists introduced the concept of utility, defining an internal, subjective evaluation of the choice options. Expected utility theory (EUT), the first axiomatic theory of risky choice (1), defined an option's subjective value as its expected utility ($\mathrm{EU} = \sum_i U(m_i) \cdot p_i$), i.e. the utility *U(m)* of possible outcome magnitudes (m_i) weighted by their respective probabilities of occurrence (p_i). This subjectively defined quantity, EU, replaced the objective expected return $\mathrm{EV} = \sum_i m_i \cdot p_i$ as the key quantity driving decisions. Although EUT was unable to explain choices in particular situations, its basic ideas remained central to the generalized EU theories that were later developed. Most notably, prospect theory (2–4) explained deviations from EUT introducing the concepts of subjective probability weighting and reference point.

In EUT, decisions are modeled "as if" subjects had an internal utility representation, making no assumptions about the brain processes underlying choice (5). 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 subjective 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, see below). If the assumptions are satisfied, the utility measure can be elicited following econometric methods based on EUT. 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 starting point: testing the basic assumptions of EUT in order to infer the existence of a utility measure.

The EU theorem demonstrated mathematically that if a subject's behavior followed a simple set of rules, or axioms, their choices could be described by the maximization of EU (SI methods: *EU theorem*), a general and basic process determining the subject's survival. All four axioms (completeness, transitivity, continuity and independence) contribute to the EU theorem, with each axiom based on the preceding ones. The first two axioms (completeness, transitivity) define a "weak order", a fundamental requirement for consistent choice behavior: subjects with complete and transitive preferences are able to order all offered options. The third axiom (continuity) introduces the contribution of reward probability: given three subjectively ranked gambles, it requires the existence of an indifference point (IP) between the intermediate gamble and a probabilistic combination of the two other gambles. The continuity axiom ensures that no option is considered infinitely better (or worse) than any other option, making it possible to define a finite, numerical value for each gamble (6–8). Finally, the independence axiom mathematically defines such numerical value as the option's EU. Violations of the independence axiom have been reported in human and animal experiments, highlighting the limits of EUT (3, 9–11). 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, i.e. the definition of a scale of numerical subjective values (2, 6, 12).

Together with completeness and transitivity, continuity constitutes the foundation for establishing well-behaved preference functions that can be used to subjectively order the choice options. Different additional constraints produced a broad spectrum of choice theories, including prospect theory, subjectively weighted utility, disappointment theory, rank-dependent and lottery-dependent utility theories (13). A form of continuity was also defined for non-risky choice theories, most notably revealed preference theory (14). The continuity axiom thus emerges as a fundamental construct in all economic schemes that imply some form of value computation.

The axiom was originally defined as a "plausible continuity assumption" (1). Thought experiments intuitively clarified how the continuity axiom could be violated (7, 15, 16), for example when options had infinitely different values. Nevertheless, the axiom was considered a reasonable condition and not experimentally tested.

Here, we explicitly tested compliance with the continuity axiom of EUT in four monkeys. The test was performed within the confines of EUT, assuming no adaptation in reference point. We tested the axiom in choices between a single gamble and a probabilistic combination of a more preferred and a less preferred

gamble. In each trial, the animal revealed its preference by selecting one of the two options. Compliance with the continuity axiom requires the existence of indifference between the middle gamble and one of the probabilistic combinations. To comply with the general definition of the axiom, 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.

RESULTS

Design. To test the continuity axiom of EUT in non-human primates, we trained four monkeys to perform a binary choice task. In each trial, the animal chose between two options, presented simultaneously on a computer monitor (Fig. 1a), offering liquid rewards varying in amount and probability (Fig. 1b). 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,

$$\forall\, A \succ B \succ C,\ \ \exists!\, \alpha \in (0,1)\ such\ that\ \ \alpha A + (1-\alpha)C \sim B \qquad (1)$$

where "$\succ$" defines a preference relation and "$\sim$" indifference; α is the specific probability associated to gamble A for which indifference occurs. Note that the axiom should be satisfied for any arbitrary set of gambles A, B and C.

To experimentally test the axiom, we first defined three gambles (Fig. 1c) 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. 1d). 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. 1e). Compliance with the continuity axiom would thus be demonstrated by the existence of a unique α between 0 and 1, while violations would occur if α were not identifiable or when multiple α existed (Fig. S1).

The continuity axiom must be satisfied for any arbitrary set of gambles. 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 *Indifference curves in the magnitude-probability space*); (*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 four monkeys, while the further two, more specific variations (*2* and *3*) were tested in only two of the animals (monkeys A and B).

We used pseudo-random repetitions of all presented choice pairs in order to account for the stochasticity of choice behavior and as a basis for future recordings of neuronal activity. Because the EUT axioms were defined as deterministic rules, we extended their definitions to the stochastic domain (4, 17).

Basic choice behavior. We investigated the consistency of choice behavior to make sure that the four 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 (SI methods: *Logistic regression*). Standardized regression coefficients (β) corresponding to reward magnitude and reward probability were significantly different from zero (one-sample *t* test, $p<0.05$, FDR corrected) in all four animals (Fig. 1f), indicating that both variables were choice-driving factors. Compared to 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: $p=1.7\cdot10^{-10}$ (A), $p=5.1\cdot10^{-3}$ (B), $p=7.3\cdot10^{-5}$ (C),

p=4.1·10^{-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 analogue of "more is better", and represents a basic requirement of EUT and of the continuity axiom in particular (SI methods: *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. Due to choice stochasticity, a number of dominated choices are naturally expected, but to comply with FSD their proportion must be significantly below 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 four animals complied with FSD by preferring the dominant option in more than 50% of trials across all FSD tests (Fig. 1g; binomial 95% CI above 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 (18), in every trial our choice set had a finite number of offered alternatives (two) with mutually exclusive (only one option could be selected) and exhaustive (the set included all possible alternatives) 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 (SI methods: *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 four animals complied with both WST and SST (Fig. 1h). 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 gambles unequivocally.

Compliance with the continuity axiom. Following the formal definition of the continuity axiom (equation 1), 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 four 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 below or above 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 (SI methods: *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 non-preferred 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. 2a,b) as well as when using a non-zero C gamble (monkeys C and D , Fig. 2c) in a different task (SI methods: *Experimental setup*) thus confirming the robustness of our results.

Importantly, different initial 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 ml to 0.35 ml; Fig. 2a,b) produced larger α values; decreasing the magnitude of the most preferred gamble (A) (e.g. from 0.50 ml to 0.40 ml (Fig. 2b)) resulted in higher α values (non-overlapping 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. S1b). 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 (SI methods: *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 both magnitude and probability information.

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.

Indifference curves in the magnitude-probability 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 in a two-dimensional diagram with reward magnitude and probability (MP) as variables. Such MP space was used to represent the continuity axiom tests, carried out in monkeys A and B, in which the B gamble is either a degenerate gamble or a true (non-degenerate) gamble, with degenerate A and C gambles and C = 0 ml. Each gamble used in a continuity test (B and AC combinations) corresponded to a single point in the MP space (Fig. 3a). 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∼AC in Fig. 3a).

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 indifference curve (IC). Importantly, there were no discontinuities ('jumps') in the IC while varying the A magnitude in 0.01 ml steps, thus fulfilling a requirement of the continuity axiom: as the magnitude of the A gamble increased, IPs gradually decreased without any apparent discontinuity (Fig. 3b).

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 to linear and hyperbolic ones (Table S1). 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. 3c).

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. 3d).

Altogether, 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 indifference curves. 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 magnitude-probability trade-off relation (Fig. S2). Note that 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. S2a), as a special case.

We estimated the utility function using single-trial choices from each session, through a maximum likelihood estimation (MLE) method. We defined a discrete choice model in standard fashion (18), with the probability of choosing one option described by a softmax function, 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 (SI methods: *Economic 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 Bayesian Information Criterion, BIC: 51.7±39.1 SD, $p=2.9\cdot10^{-18}$, Monkey A; 54.9±29.5 SD, $p=1.7\cdot10^{-27}$, Monkey B; one-sample *t* test), while the s-shaped utility function outperformed the power-shaped one (BIC difference: 12.2±12.3 SD, $p=2.1\cdot10^{-13}$, Monkey A; 8.7±12.2 SD, $p=1.4\cdot10^{-8}$, Monkey B). The two recovered parameters for the s-shaped utility functions (Fig. 4a, histograms) were both significantly different from one ($p<10^{-15}$ in both monkeys; one-sample *t* test), confirming that utility functions were non-linear and had a significant inflection point (i.e. a change in curvature), resulting in an s-shaped curve.

We used the recovered s-shaped utility functions (Fig. 4a) 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. 4b), 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 (red lines in Fig. 4b) was smaller for the EU model than for the objective-EV model (dashed curves in Fig. 4b) (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 to 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 (vertical dotted lines in Fig. 4c). A lower average variance for the EU model indicated that it was better at describing preferences compared to the EV model (Table S2). This was also confirmed using a standard model comparison analysis (BIC and AIC scores, Table S2). These results indicate that the non-linearity in the utility function was able to capture the subjective quality of the IPs and of the revealed preferences (Fig. 4b,c).

Non-linear probability weighting is a further subjective factor explaining economic choices, as proposed in prospect theory and other generalized EU theories (4, 12). 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 to the EV model. Therefore, the PW model had similar descriptive power compared to the EU model. Overall, the PW model outperformed the EU model in 3 out of 4 goodness-of-fit measures, in both monkeys (Table S2), suggesting that in the tested choice situation adding the subjective weighting of probabilities marginally improved the description of preferences compared to EUT, representing a possible refinement to our EU model for describing preferences (Fig. 4c).

The mean-variance approach, an alternative economic model which approximates EUT without relying on the concept of utility (19), defines a gamble's value as the sum of the corresponding EV and risk components (SI methods: *Economic 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 (Table S2, Fig. S3), with an improvement in the ICs 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 pairs of IPs. A significant ρ, both within each IC (one-sample *t* test, per animal: $p=1.8\cdot10^{-5}$ (Monkey A); $p=5.4\cdot10^{-4}$ (Monkey B)) and across different ICs ($p=8.2\cdot10^{-5}$ (Monkey A); $p=1.5\cdot10^{-3}$ (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 modelling of ICs and the correlation among IPs support the idea of choices resulting from a utility maximization process: the combination of a subjectively defined utility function with reward probabilities (possibly subjectively weighted) was able to describe the choice behavior and in particular the smooth trade-off relation between reward magnitudes and probabilities.

Continuity axiom test in the Marschak-Machina triangle. The Marschak-Machina triangle (20, 21) has been extensively used in economic studies of human behavior for evaluating and comparing different generalized EU theories (22, 23). 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. 5a, inset). The Marschak-Machina triangle represents gambles with three fixed outcome magnitudes (defined in our experiment as 0 ml, 0.25 ml 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. 5a).

We defined gamble A with 0.25 ml and 0.5 ml as possible outcomes and gambles B and C each with 0 ml and 0.25 ml outcomes; AC combinations then corresponded to gambles with possible reward magnitudes of 0 ml, 0.25 ml and 0.5 ml (Fig. 5b). In the Marschak-Machina triangle, gamble A lay on the y 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. 5b, 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 (labelled B~AC in Fig. 5b, bottom).

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. 5c). Results showed the existence of IPs in all tested cases (Fig. 5d), 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) lead 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. 5c).

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 demonstrated compliance of monkey behavior with the continuity axiom of EUT, implying a magnitude-probability trade-off relation and determining a numerical utility measure able to describe choices.

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 two 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, while 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. Though it is still unknown whether non-human 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 monkey decisions reflect both the magnitude and probability information of the choice options (24, 25), 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 multi-attribute choices: reward probability and magnitude were both considered and combined when evaluating the options. However, we did not test further 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 contributing to long-term learning of values and probabilities (26). Thus, further tests should delimit the viability of continuity satisfaction in different choice situations.

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 indifference curve. 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 indifference curves, while remaining constant within each indifference curve.

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 non-linear weighting of reward probabilities resulted in a marginally better description of choice behavior compared to EUT. Therefore, the present study points to non-linear probability weighting as a possible refinement to EUT in monkeys, compatible with the human experimental results that led to the development of prospect theory (3). 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.

METHODS

Animals and Experimental Setup. Four male rhesus macaques (*Macaca mulatta*) were used in this study. During the experiment, the monkey sat in a primate chair (Crist Instruments) and made choices between two rewarding stimuli presented on a computer monitor. The animals reported their preferences with a left-right motion joystick (monkeys A, B) or through arm movements toward a touch screen (monkeys C, D). Task event-times were sampled and stored at 1 kHz on a Windows 7 computer running custom MATLAB (The MathWorks) code, using Psychtoolbox 3. All experimental protocols were assessed and approved by the Home Office of the United Kingdom.

Axioms of EUT. The axioms of EUT are necessary and sufficient conditions for choices to be described by the maximization of EU. The continuity axiom requires the existence of an indifference in choices between a fixed gamble and a combination of two other gambles and 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 larger than 0.5 (binomial test, $p<0.05$). For multiple comparisons, we applied a false discovery rate (FDR) correction (Benjamini-Hochberg procedure) (27). More details about behavioral training, task and data analysis are available in the SI Appendix, Methods section.

ACKNOWLEDGEMENTS

This work was supported by Wellcome Grants WT 095495 and WT 204811 and European Research Council Advanced Grant 293549.

REFERENCES

1. J. von Neumann, O. Morgenstern, Theory of Games and Economic Behavior. *Princet. Univ. Press*, 625 (1944).
2. C. Starmer, Developments in nonexpected-utility theory: The hunt for a descriptive theory of choice under risk. *Adv. Behav. Econ.* **XXXVIII**, 332–382 (2000).

3. D. Kahneman, A. Tversky, Prospect Theory: An Analysis of Decision under Risk. *Econometrica* **47**, 263–292 (1979).

4. J. D. Hey, C. Orme, Investigating Generalizations of Expected Utility Theory Using Experimental Data. *Econometrica* **62**, 1291 (1994).

5. C. F. Camerer, Neuroeconomics: Opening the Gray Box. *Neuron* **60**, 416–419 (2008).

6. G. a Jehle, P. J. Reny, *Advanced Microeconomic Theory* (Boston, MA: Addison-Wesley, 2001).

7. A. Chateauneuf, M. Cohen, J.-M. Tallon, "Decision under Risk: The Classical Expected Utility Model" in *Decision-Making Process*, (ISTE, 2008), pp. 365–384.

8. E. J. Johnson, R. Ratcliff, *Neuroeconomics* (Elsevier, 2014).

9. M. Allais, Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* **21**, 503 (1953).

10. B. J. Conlisk, Three Variants on the Allais Example. *Am. Econ. Rev.* **79**, 392–407 (1989).

11. R. C. Battalio, J. H. Kagel, D. N. MacDonald, Animals' Choices over Uncertain Outcomes: Some Initial Experimental Results. *Am. Econ. Rev.* **75**, 597–613 (1985).

12. B. Y. D. W. Harless, C. F. Camerer, The Predictive Utility of Generalized Expected Utility Theories. *Econometrica* **62**, 1251–1289 (1994).

13. M. Weber, C. Camerer, Recent developments in modelling preferences under risk. *OR Spektrum* **9**, 129–151 (1987).

14. P. A. Samuelson, Consumption Theory in Terms of Revealed Preference. *Economica* **15**, 243–253 (1948).

15. J. Levin, Choice under Uncertainty. *Lect. notes* (2006).

16. N. Georgescu-Roegen, Choice, Expectations and Measurability. *Q. J. Econ.* **68**, 503 (1954).

17. G. Loomes, Modelling the stochastic component of behaviour in experiments: Some issues for the interpretation of data in *Experimental Economics*, (2005), pp. 301–323.

18. D. McFadden, Economic Choices. *Am. Econ. Rev.* **91**, 351–378 (2001).

19. H. Levy, H. M. Markowitz, Approximating expected utility by a function of mean and variance. *Am. Econ. Rev.* **69**, 308–317 (1979).

20. J. Marschak, Rational Behavior, Uncertain Prospects, and Measurable Utility. *Econometrica* **18**, 111 (1950).

21. M. J. Machina, "Expected Utility" Analysis without the Independence Axiom. *Econometrica* **50**, 277 (1982).

22. C. F. Camerer, An experimental test of several generalized utility theories. *J. Risk Uncertain.* **2**, 61–104 (1989).

23. B. Sopher, G. Gigliotti, A test of generalized expected utility theory. *Theory Decis.* **35**, 75–106 (1993).

24. M. L. Platt, P. W. Glimcher, Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).

25. A. Lak, W. R. Stauffer, W. Schultz, Dopamine prediction error responses integrate subjective value from different reward dimensions. *Proc. Natl. Acad. Sci. U. S. A.* **111**, 2343–2348 (2014).

26. S. Ferrari-Toniolo, P. M. Bujold, W. Schultz, Probability distortion depends on choice sequence in rhesus monkeys. *J. Neurosci.* **39**, 2915–2929 (2019).

27. Y. Benjamini, Y. Hochberg, Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *J. R. Stat. Soc. Ser. B* **57**, 289–300 (1995).

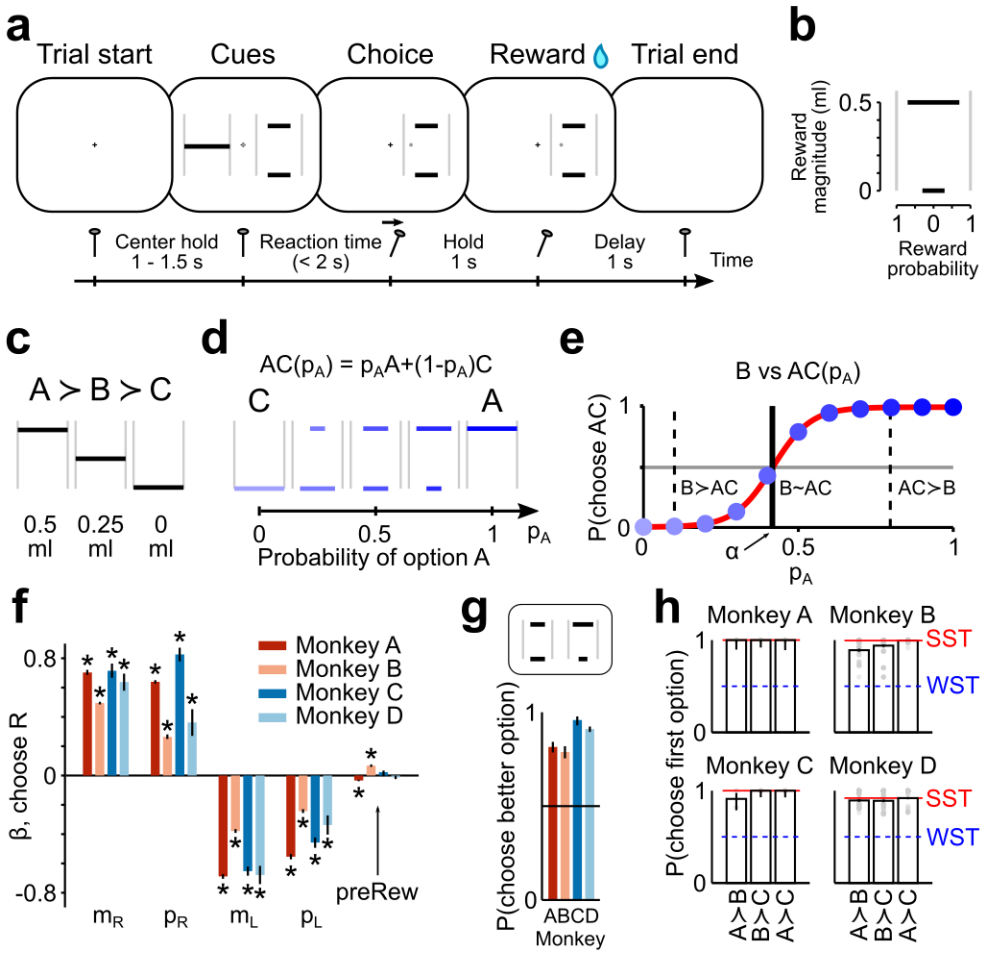

Figure 1. *Experimental design and consistency of choice behavior*. **(a) Trial sequence**. Monkeys chose between two options by moving a cursor (gray dot) to one side of the screen. After a delay, the reward corresponding to the selected cue was delivered. **(b) Visual cues** indicated magnitude and probability of possible outcomes through horizontal bars' vertical position and width, respectively. **(c,d,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 correspond to 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 choice indifference between the two options (B~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 sigmoid function (red line) to the proportion of AC choices (blue dots). **(f,g,h) Consistency of choice behavior.** The standardized beta coefficients from logistic regressions of single trials' behavior (**f**) 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: 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 (**g**) (error bars: 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: average across all sessions; gray dots: single sessions; error bars: binomial 95% CI) animals preferred A to B, B to C and A to C (**h**), complying with both weak and strong stochastic transitivity (WST: proportion of choices of the better option >0.5 (blue dashed line); SST: proportion of A over C choices (red line) ≥ other choice proportions).

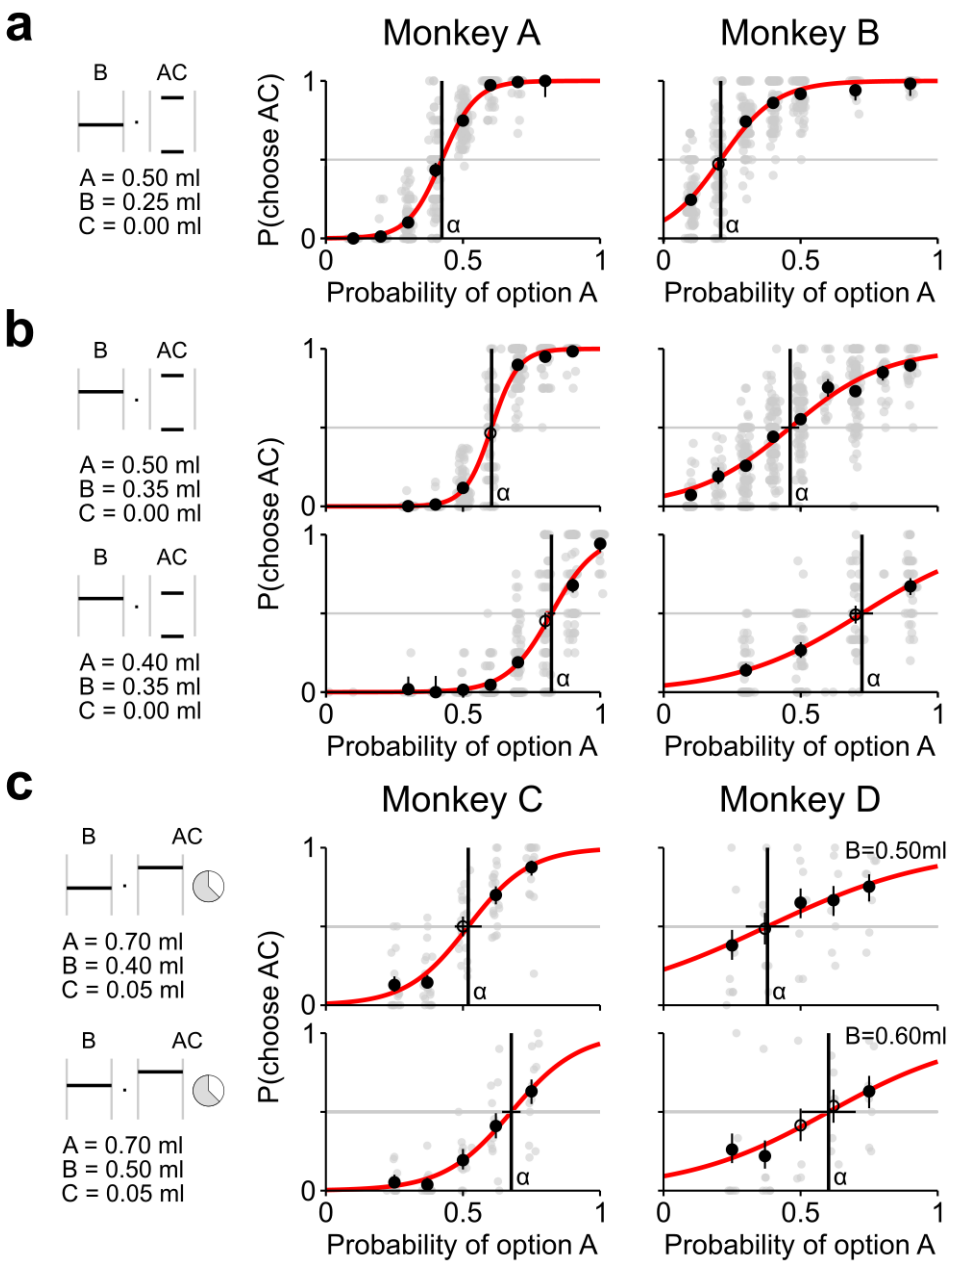

Figure 2. *Experimental test of the continuity axiom*. **(a,b,c) Compliance with the continuity axiom**. The axiom was tested through choices between a gamble B and a varying AC combination (left: visual stimuli for an example choice pair with p_A=0.5 (a,b) or p_A=0.375 (c)); increasing p_A values resulted in gradually increasing preferences for the AC option. In each plot, gray dots represent the proportion of AC choices in single sessions, black circles the proportions across all tested sessions with vertical bars indicating the binomial 95% confidence intervals (filled circles indicate 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 four 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 non-preferred ($p_{choose\ AC}$<0.5) to preferred ($p_{choose\ AC}$>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 had a value of 0.5 (horizontal bars: 95% CI); α values shifted coherently with changes in A and B values in all four animals, indicating a continuous magnitude-probability trade-off relation.

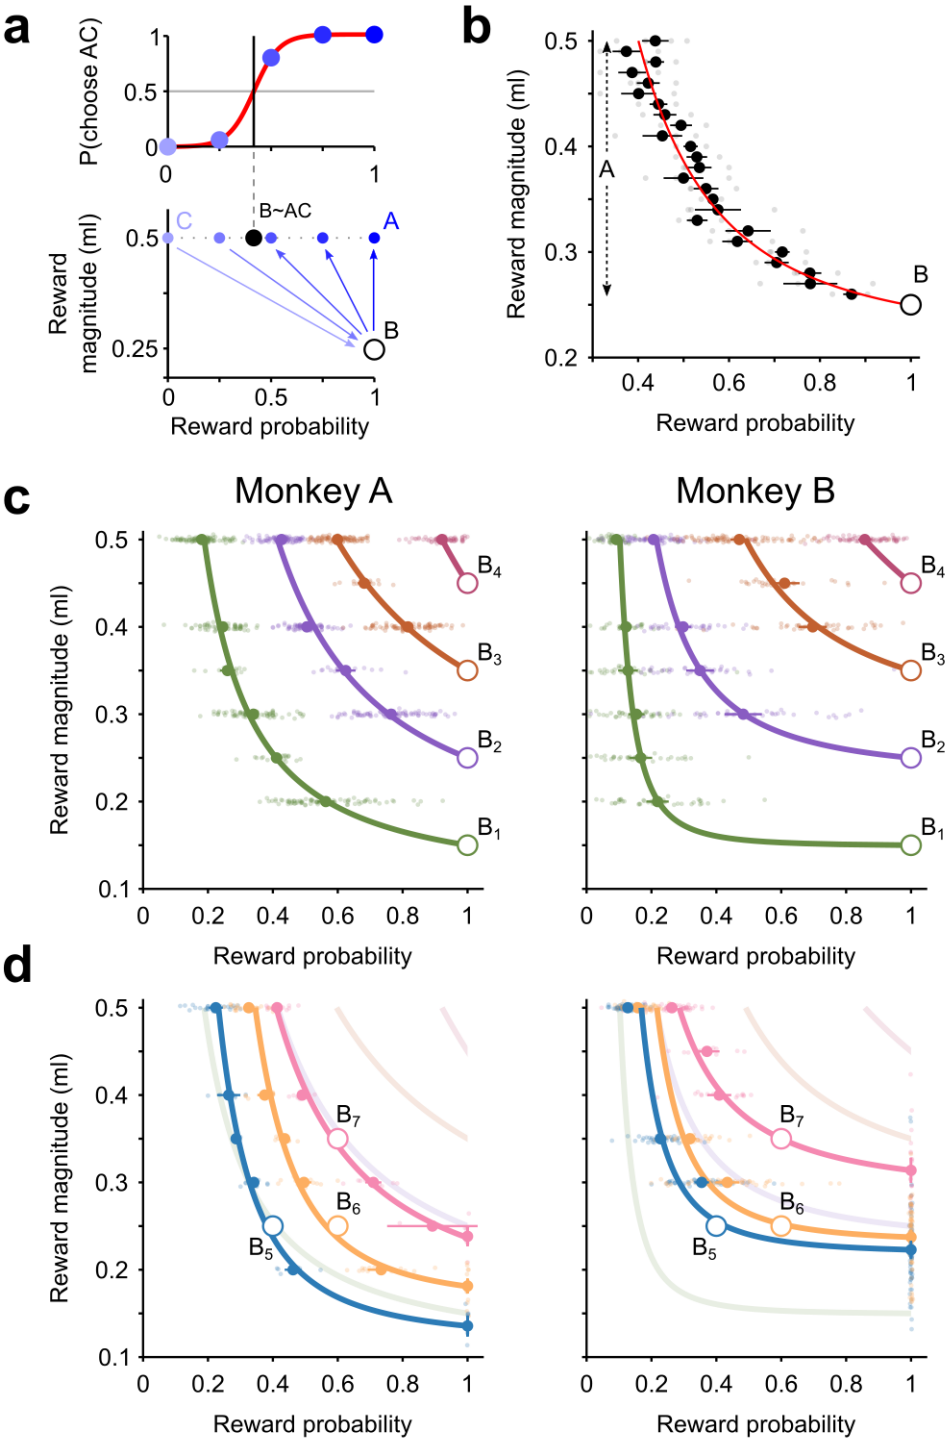

Figure 3. *Indifference curves in the MP space*. (**a**) **Representation of the continuity axiom test in the MP space**. The gambles used for testing the axiom can be mapped into the magnitude-probability diagram. Preference in choices between B (circle) and combinations of A and C (graded blue dots) 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**) **Indifference curve**. IPs (gray dots: single sessions; black dots: averages; bars: SE) obtained using different A values (step 0.01 ml) shifted continuously, producing an IC in the MP space. Curve: best fitting power function. Data from monkey A (5 sessions, 1781 trials). **(c,d) Indifference map**. ICs for different B values (colored curves), described the gradual variation of the average IPs (colored dots, with SE bars) for each B. Small dots represent 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.

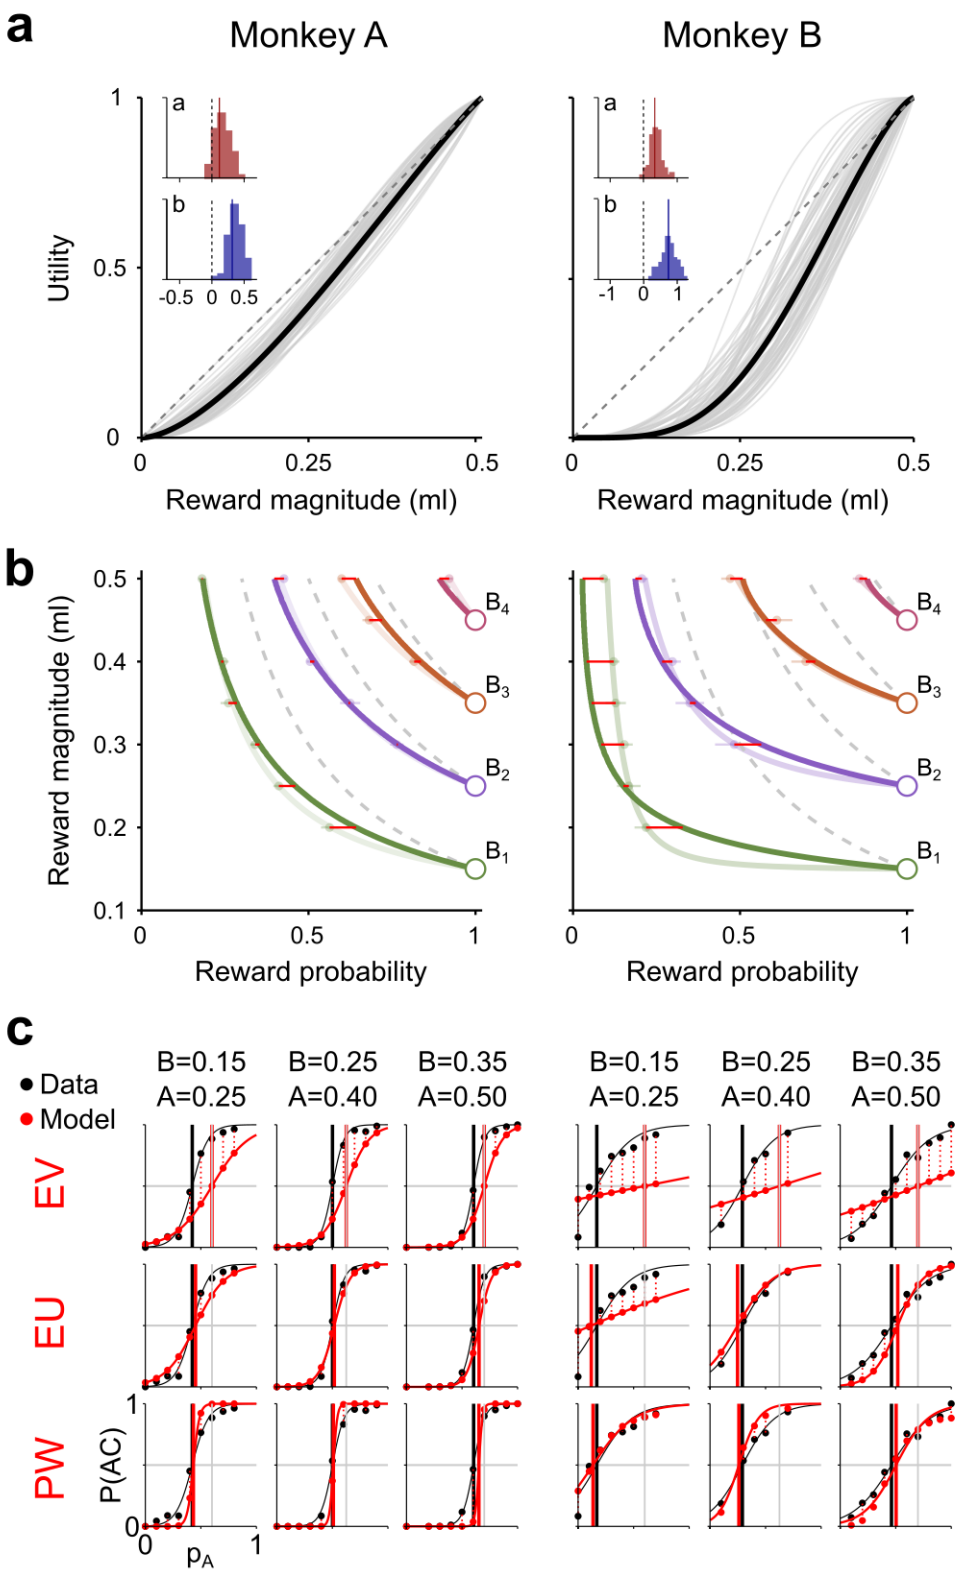

Figure 4. *Indifference curves are compatible with economic models.* **(a) Utility functions.** Single sessions' utility functions (gray) and averages (black) estimated through MLE using single-trial choice data. The two estimated parameters (*a* and *b* in inset, histograms of log-values) were both significantly positive, indicating s-shaped utility functions. **(b) EUT predicted indifference curves**. Indifference map reconstructed using the estimated utility functions. The light-colored curves represent the measured indifference map (Fig. 3c); the red horizontal lines identify the distance between measured and modeled IPs. Dashed gray curves define points with equal EV, corresponding to a linear utility model. **(c) Comparison of modeled and revealed preferences**. Percentage of choices for the AC gamble (P(AC)) measured in the axiom test (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 (grey vertical lines), with a larger error in the prediction of the P(AC) (vertical dotted red lines) compared the EU model. The PW model, which included a subjective probability weighting, was better at capturing the revealed preferences only in specific cases (e.g. B=0.15, A=0.25 in Monkey B).

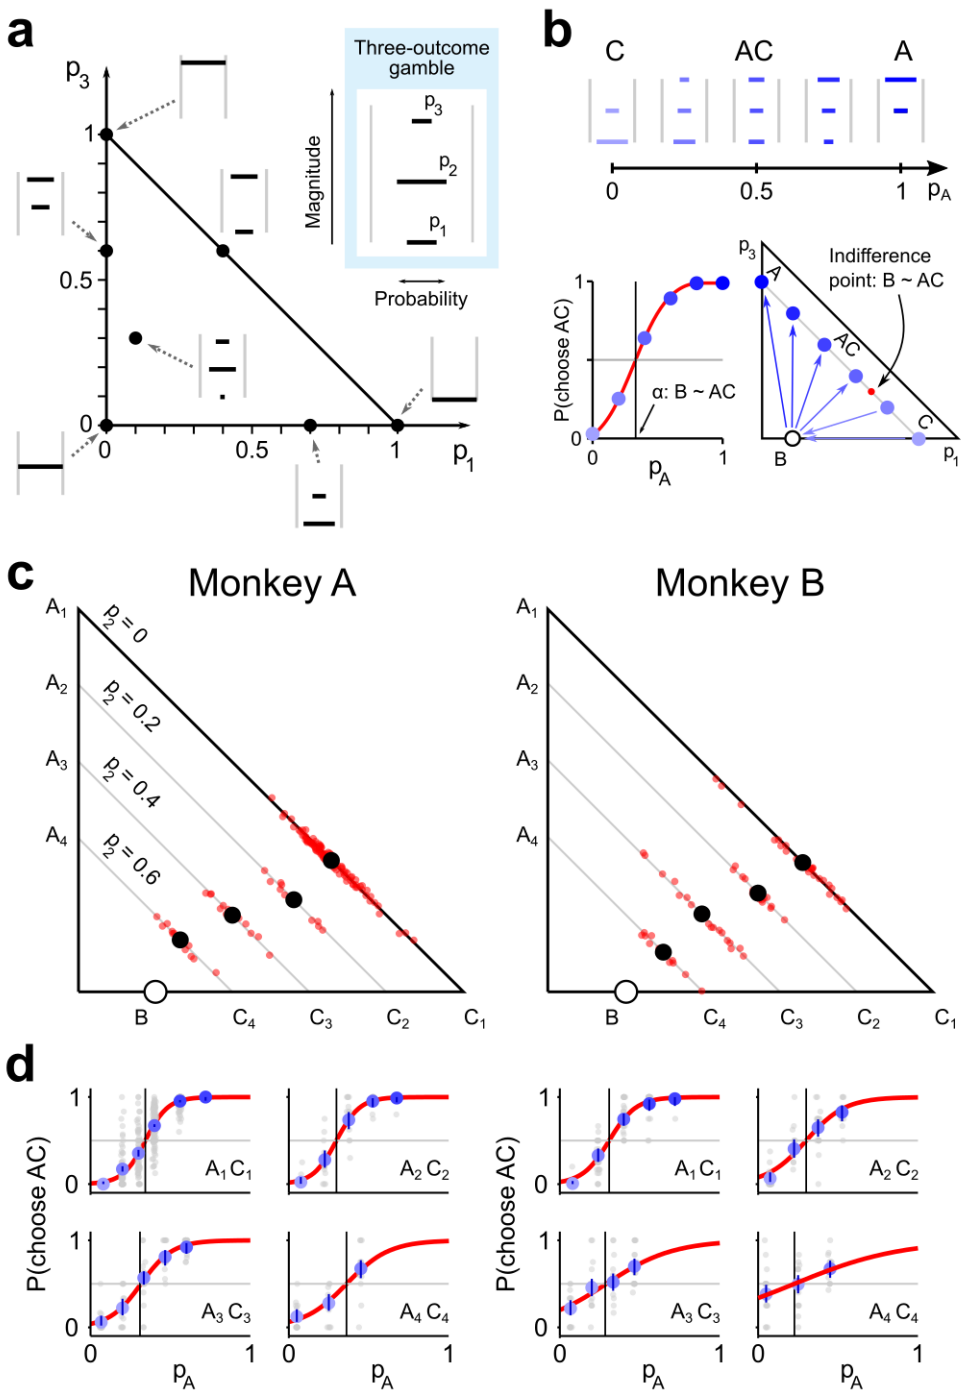

Figure 5. *Continuity axiom test in the Marschak-Machina triangle.* (**a**) **Three-outcome gambles**. Gambles with three fixed outcome magnitudes and any combination of outcome probabilities can be represented in the Marschak-Machina triangle. The visual cue (inset) for three-outcome gambles included three horizonal lines representing 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). The arrows point toward the preferred option, consistently with the proportion of AC choices (bottom left). **(c) Continuity axiom 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) Continuity axiom test 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). Other symbols as in Fig. 2.

SUPPLEMENTARY INFORMATION

METHODS

Experimental setup

We trained (>10,000 trials) four male rhesus macaque monkeys (weight, per animal: 12.7 kg (Monkey A), 13.8 kg (Monkey B), 10.3 kg (Monkey C) and 12.5 kg (Monkey D)) to express their preferences between pairs of probabilistic rewards, represented as visual cues on a computer monitor. Monkeys sat in primate chairs (Crist Instruments) and expressed choices through arm movements. Monkeys A and B moved a joystick (Biotronix workshop, Cambridge) restricted to left/right movements, to control a cursor on a computer monitor vertically positioned 30 cm in front of them; monkeys C and D made arm movements towards 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 0.00-0.50 ml (Monkeys A and B) or 0.05-0.90 ml (Monkeys C and D). A computer-controlled solenoid valve delivered juice reward from a spout in front of the animal's mouth.

Task 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 non-safe 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 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 (Fig. 1a). Visual cues were presented on a blank screen, indicating the amount (magnitude) and probability of receiving a reward (fruit juice) though 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 (Fig. 1b).

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 = cued probability and a fixed small reward (0.05 ml) could be obtained with P = 1 – 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 1,500 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 1,000 ms which ended with extinction of the gray background.

Logistic regression

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

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

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 $preCh \cdot preRew$ thus increases for larger rewards obtained when choosing the right-side option in the previous trial; ε is the error term. Regression coefficients (β) were standardized by multiplying each coefficient with the ratio of the corresponding independent variable's SD over the SD of the predicted variable. Standardized regression coefficients were tested for statistical significance through one-sample t test.

Axioms of expected utility theory

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 (1).

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, 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*: $\forall\, A \succ B \succ C, \exists!\, p_1, p_2 \in (0, 1)$: $p_1A+(1-p_1)C \succ B \succ p_2A+(1-p_2)$

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

$A \succ B \Leftrightarrow U_A > U_B$

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):

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

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:

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

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 of the objective (mathematical) expected value $EV = \sum_i m_i \cdot p_i$.

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 (Fig. S1, right).

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 to 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 (binomial test, $p<0.05$).

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_AA+(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 if 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 indifference point (α), its numerical value was determined by fitting a softmax function to the choice data though non-linear least squares fit. The softmax function was defined as follows:

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

where τ (softmax "temperature" parameter) represents the steepness of the preference function.

Data fitting of indifference points

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:

$$Linear: f(x) = ax + b$$

$$Power: f(x) = a(x - c)^b$$

$$Hyperbolic: f(x) = a + b/(x - c)$$

where x represents the reward magnitude, $f(x)$ the IC, i.e. the reward probability as a function of reward magnitude. A non-linear least squares method was used to minimize the error in the probability domain (x-axis in Fig. 3 and Fig. 4b)

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:

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

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.

$$LL(param|choices) = \sum_{trials} A_chosen \cdot \log\big(P(A, param)\big) + (1 - A_chosen) \cdot \log(1 - P(A, param))$$

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 $P(A, param)$ 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:

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

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

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

with $m_0 = 0.5\ ml$, 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 expected value of the options: it would perform as the best model only if monkeys were choosing by comparing the objective, mathematical expected value 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 probability weighting 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 the gamble value is defined as $V = \sum_i U(m_i) \cdot w(p_i)$; we used the two-parameter *Prelec* function as the probability weighting function $w(p)$. According to the mean-variance approach, the value definition does not rely on the utility concept: $V = EV + \beta \cdot Risk$, 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 expected value of the squared deviation from the mean: $Risk = \sum_i p_i \cdot (m_i - EV)^2$.

In order to construct the full indifference map predicted by a model, for each IC we numerically computed the indifference points 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 indifference point. 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 five tested models (EUT with three possible utility functions, PT and mean-variance) we defined four goodness-of-fit measures: 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, both introducing a penalty term when increasing the number of model parameters; the variance in the differences of modeled vs 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 to a higher value.

TABLES

Monkey	*Linear*	*Hyperbolic*	*Power*
A	$1.1 \cdot 10^{-2}$	$9.9 \cdot 10^{-5}$	$9.6 \cdot 10^{-5}$
B	$3.4 \cdot 10^{-2}$	$2.2 \cdot 10^{-4}$	$1.7 \cdot 10^{-4}$

Table S1. Comparison of ICs fit models. Mean squared error (MSE) resulting from fitting IPs to three different models. Bold face indicates the lowest MSE value for each animal.

	Goodness-of-fit measure	*Model 1: EV*	*Model 2: EU (power)*	*Model 3: EU (s-shaped)*	*Model 4: PT*	*Model 5: Mean-variance*
Monkey A	$\sqrt{MSE}$	0.108	0.046	0.028	0.024	0.066
	BIC	$2.64 \cdot 10^4$	$2.25 \cdot 10^4$	$2.22 \cdot 10^4$	$2.24 \cdot 10^4$	$2.34 \cdot 10^4$
	AIC	$2.61 \cdot 10^4$	$2.19 \cdot 10^4$	$2.14 \cdot 10^4$	$2.09 \cdot 10^4$	$2.28 \cdot 10^4$
	Var	0.154	0.122	0.117	0.113	0.129
Monkey B	$\sqrt{MSE}$	0.274	0.093	0.052	0.062	0.193
	BIC	$3.03 \cdot 10^4$	$2.59 \cdot 10^4$	$2.53 \cdot 10^4$	$2.48 \cdot 10^4$	$2.80 \cdot 10^4$
	AIC	$3.00 \cdot 10^4$	$2.54 \cdot 10^4$	$2.44 \cdot 10^4$	$2.34 \cdot 10^4$	$2.75 \cdot 10^4$
	Var	0.225	0.176	0.165	0.155	0.197

Table S2. Comparison of economic models. Each row of values is a comparison across models using one goodness-of-fit measure (averaged across all tests and sessions), with bold face indicating the best fitting model according to that measure. EV corresponds to the EU model assuming a linear utility function. In prospect theory (PT) the gamble value was computed as $V=U(m) \cdot w(p)$ (valid, according to PT, for all gambles with one non-zero outcome), with $w(p)$ being the probability weighting function (2-parameter *Prelec* function). The square root of the MSE represent the average distance between model and IP in probability units. *Var* represents the variance in the differences of modeled vs measured preferences (the proportion of AC vs B choices across all continuity tests).

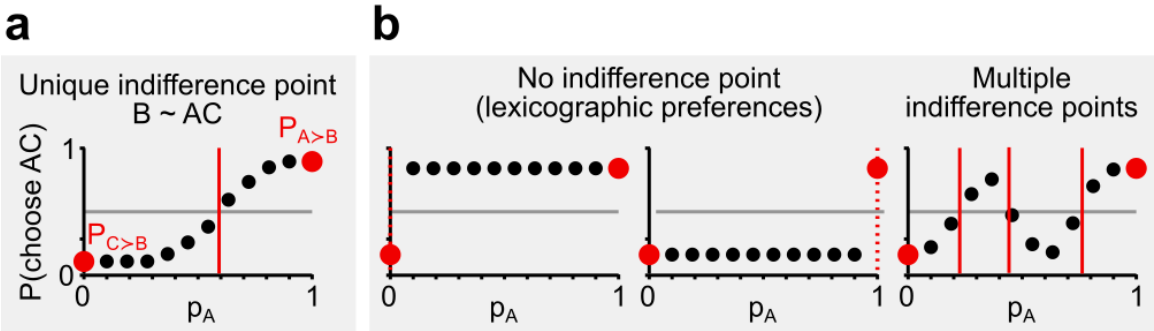

Figure S1. Continuity axiom compliance and violation. Choice pattern compatible with the continuity axiom (a) and possible axiom violations (b). Red dots represent the proportion of AC choices when p_A=0 or 1, corresponding to the axiom's initial requirement (A≻B and B≻C, implying P(A≻B)>0.5 and P(C≻B)<0.5).

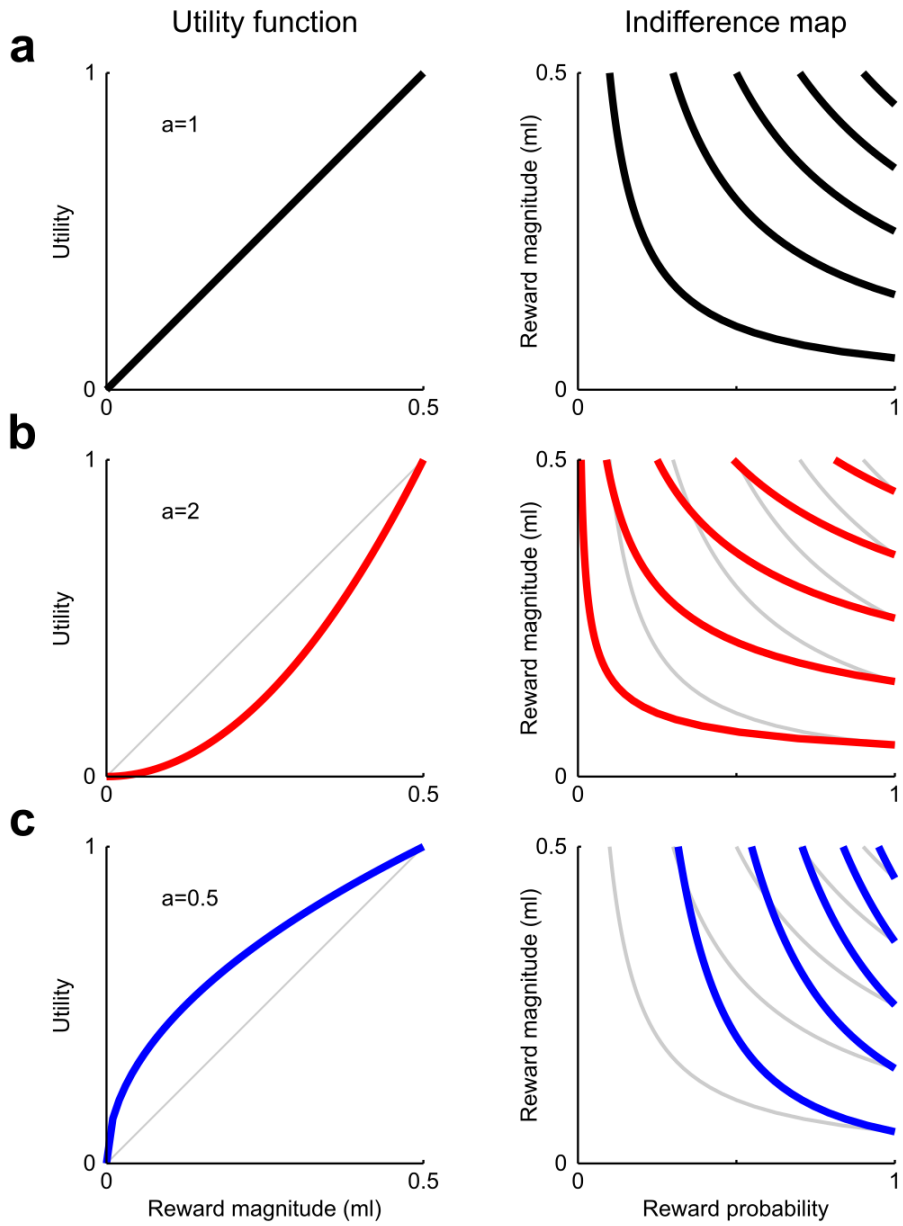

Figure S2. Relation between the utility function and the corresponding indifference map. Sample indifference maps obtained from different utility functions, with a representing the single parameter of a power function ($U(m) = (m/m_0)^a$, with $m_0 = 0.5\ ml$). The indifference curve for a degenerate 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 indifference curve equation can be obtained: $p(m) = (m_B/m)^a$, 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), risk aversion for concave utility (c). In all plots, risk neutrality (i.e. choices based on the objective EV of gambles) is represented by grey curves. The indifference map globally warps according to the risk attitude.

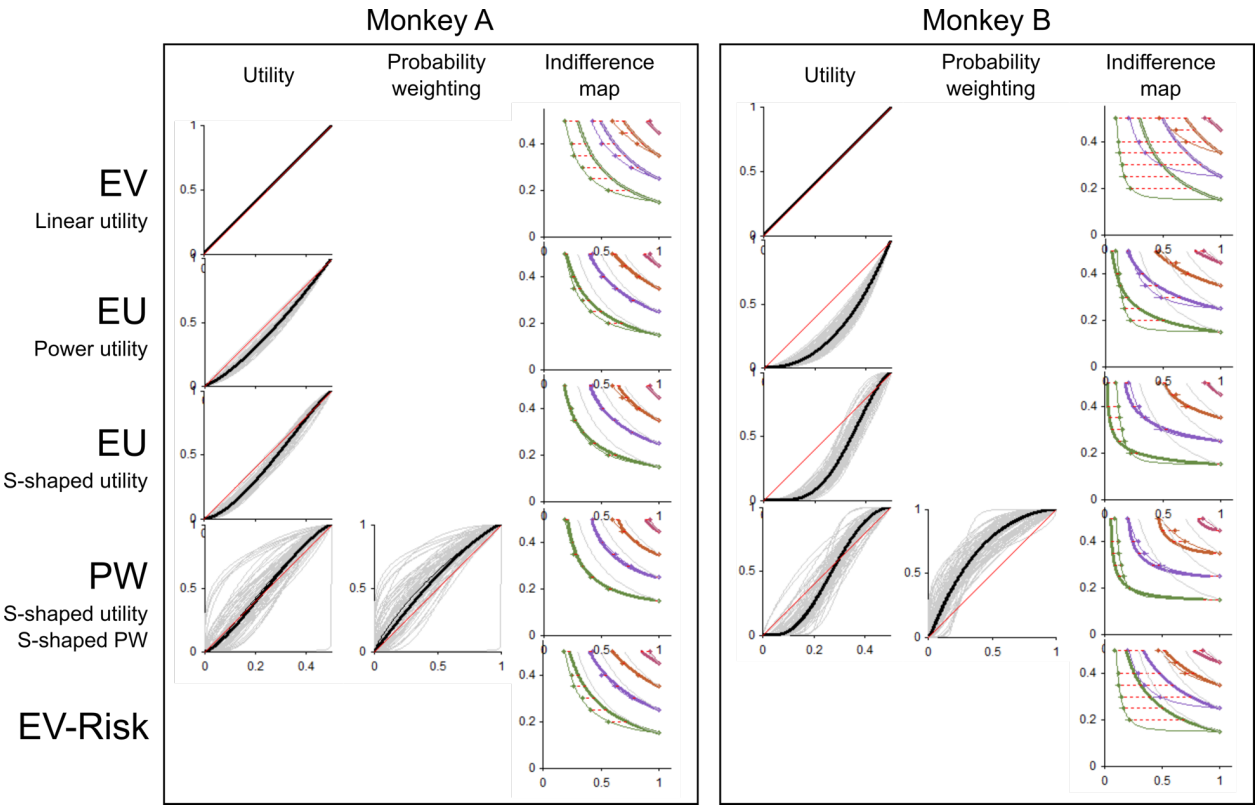

Figure S3. Comparison of economic models. Recovered utility function, probability weighting function and corresponding indifference map for each economic model (rows). Grey curves represent single session estimates, black curves are means, plotted by averaging the recovered parameters across all sessions. Red lines represent linear utility and probability weighting (PW) functions, for comparison. The EV and EU models assume linear probability weighting. Note that the mean-variance model (EV-Risk) does not have a utility representation. Other conventions and symbols as in Fig. 4b.