

Utility functions predict variance and skewness risk preferences in monkeys

Wilfried Genest^{a,1}, William R. Stauffer^{a,b}, and Wolfram Schultz^a

^aDepartment of Physiology, Development, and Neuroscience, University of Cambridge, Cambridge CB2 3DY, United Kingdom; and ^bSystems Neuroscience Institute, Department of Neurobiology, University of Pittsburgh, Pittsburgh, PA 15261

Edited by Charles R. Gallistel, Rutgers University, Piscataway, NJ, and approved May 19, 2016 (received for review February 9, 2016)

Utility is the fundamental variable thought to underlie economic choices. In particular, utility functions are believed to reflect preferences toward risk, a key decision variable in many real-life situations. To assess the validity of utility representations, it is therefore important to examine risk preferences. In turn, this approach requires formal definitions of risk. A standard approach is to focus on the variance of reward distributions (variance-risk). In this study, we also examined a form of risk related to the skewness of reward distributions (skewness-risk). Thus, we tested the extent to which empirically derived utility functions predicted preferences for variance-risk and skewness-risk in macaques. The expected utilities calculated for various symmetrical and skewed gambles served to define formally the direction of stochastic dominance between gambles. In direct choices, the animals' preferences followed both second-order (variance) and third-order (skewness) stochastic dominance. Specifically, for gambles with different variance but identical expected values (EVs), the monkeys preferred high-variance gambles at low EVs and low-variance gambles at high EVs; in gambles with different skewness but identical EVs and variances, the animals preferred positively over symmetrical and negatively skewed gambles in a strongly transitive fashion. Thus, the utility functions predicted the animals' preferences for variance-risk and skewness-risk. Using these well-defined forms of risk, this study shows that monkeys' choices conform to the internal reward valuations suggested by their utility functions. This result implies a representation of utility in monkeys that accounts for both variance-risk and skewness-risk preferences.

experimental economics | stochastic dominance | choice | decision making | neuroeconomics

Future rewards can rarely be predicted with complete accuracy; often, they are inherently risky. A principled approach with well-defined forms of risk is therefore highly desirable for a thorough understanding of economic decisions. Good definitions of risk are available when considering rewards as probability distributions of different values. For instance, rewards can fluctuate symmetrically around a mean value, a form of risk captured by the statistical variance. This simple form of risk is commonly studied by economists, ethologists, and neurophysiologists (1–9). Reward distributions are often asymmetrical, however. This asymmetry can be described formally by considering the statistical “skewness” or “skewness-risk.” For instance, a tree may usually be barren but occasionally be rich with fruit (positive skewness), or a bush may produce a handful of berries most of the time but occasionally yield nothing (negative skewness). Skewness is an important form of risk that is abundant in natural environments, and thus is also likely to be of fundamental importance for the animals' behavior. Despite their prevalence in real-life situations, preferences for these distinct forms of risk are not well characterized.

Classically, risk preferences are derived from formal economic utility functions (1, 10). These functions describe a nonlinear processing of reward magnitude: Doubling an amount of money may not double its utility, for instance. Fig. 1*A* shows how a typical utility function leads to lower utility for gambles with higher mean-preserving variance (blue), which should result in an aversion to

variance-risk. Fig. 1*B* shows how the same utility function yields a higher utility for positively skewed gambles (red), which should result in a preference for skewness-risk, despite such gambles offering a return below the statistical expected value (EV) most of the time. Overall, this one utility function (Fig. 1) predicts an aversion for variance-risk but a liking of positive skewness-risk. This example illustrates how variance-risk and skewness-risk preferences are independent of each other but can still both be accounted for by utility functions (11).

We investigated the predictions of empirical utility functions for both variance-risk and skewness-risk in rhesus monkeys. The use of monkeys allowed us to extend the validity of utility functions to a close evolutionary relative of humans (*i*) without interference by language, (*ii*) with real-world rewards (milliliters of juice) instead of hypothetical outcomes, and (*iii*) in a situation that is suitable for later neuronal recordings [the problem of situation dependency has been highlighted before (12)]. Specifically, we examined how the utility function of each animal predicted the observed choice preferences for gambles with different variance-risk and skewness-risk. In a series of gambles, we manipulated EV, variance, and skewness independently. We estimated with psychometric methods an empirical utility function for each animal and calculated the expected utility (EU) of each gamble. These procedures defined the descending or ascending (13, 14) direction of second-order (variance) and third-order (skewness) stochastic dominance of the gambles. Then, we used direct choices to test whether the monkeys' preferences abided by the dominance relationship between the gambles. The monkeys' preference between gambles consistently matched the predictions of the monkeys' EUs. These results demonstrate, across different forms of risk, the validity of utility as an internal measure of reward value.

Results

Experimental Design and Behavior. Two monkeys made choices between reward-predicting visual stimuli presented on a computer monitor in front of the animal, using a joystick (Fig. 2*A*). Liquid reward (diluted blackcurrant juice) was delivered by means

Significance

Utility, the key decision variable underlying economic choices, should represent risk, which is inherent to real-life decisions. We studied two prevalent forms of risk that are characterized by the spread (variance-risk) and asymmetry (skewness-risk) of rewards. We show that monkeys preferred higher variance and positively skewed gambles. Importantly, empirically estimated utility functions predicted both of these risk preferences. Thus, the abstract concept of utility seemed to explain primates' choices under common forms of risk.

Author contributions: W.G., W.R.S., and W.S., designed research; W.G. performed research; W.G. analyzed data; and W.G. and W.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: wg231@cam.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1602217113/-DCSupplemental.

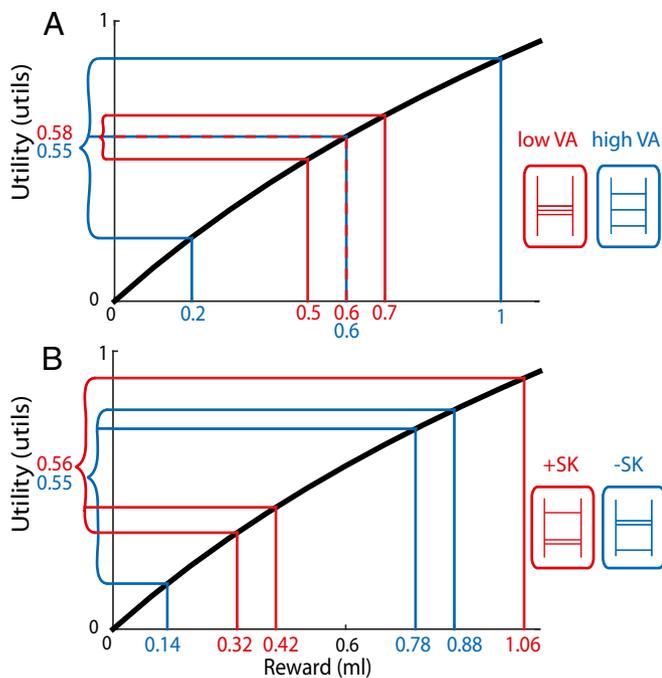


Fig. 1. Independent variance-risk and skewness-risk preferences indicated by a logarithmic utility function [$y = \ln(x + 1)/0.81$]. (A) Variance-risk with symmetrical, three-outcome, equiprobable gambles with a different mean-preserving spread (identical EV = 0.6 mL for both gambles; for each outcome, $P = 1/3$). The gamble with lower variance-risk (red, low VA) has higher EU (numbers to the left) than the gamble with higher variance-risk (blue, high VA). A decision maker with such a utility function should prefer the red gamble. (B) Skewness-risk with asymmetrical, three-outcome, equiprobable gambles with identical EV and variance. The EU of the positively skewed gamble (red, +SK) is higher than the EU of the negatively skewed gamble (blue, -SK). A decision maker with such a utility function should prefer the positively skewed gamble. SK, skewness.

of a computer-controlled solenoid liquid valve (0.004 mL/ms opening time, $R^2 > 0.99$; Fig. S1A). The monkeys were rewarded on every correct trial to maintain motivation, and not based on a randomly selected trial as in human studies. Potential satiety effects were controlled for by pooling choices across and within sessions and by limiting session length. The visual stimuli were horizontal bars; their vertical position indicated the volume of juice the monkey could receive when choosing that option. Stimuli with two, three, or four horizontal bars indicated an equiprobable gamble between the indicated magnitudes (Methods). Importantly, these stimuli enabled independent changes in EV, variance, and skewness of the reward distribution (Fig. 2B). We studied choices between equiprobable gambles and safe rewards, or between equiprobable gambles with the same number of outcomes, thus avoiding differences in probability distortion. In control choices between two safe rewards, both animals chose the higher magnitude option in most trials, suggesting proper understanding of the bar stimuli (Fig. S1B).

Estimating Empirical Utility Functions. To assess the EU of each gamble, we estimated Von Neumann–Morgenstern utility functions under risk (10, 15). We estimated, separately for each animal, utility functions from psychometrically measured certainty equivalent (CE) for binary, two-outcome gambles (8, 16) (Fig. S2A and B), using a fractile procedure that iteratively sections the full reward range according to the obtained CEs (Fig. S2C and SI Methods). We selected a broad reward range (0.1–1.3 mL) to capture both risk-seeking and risk-avoiding behavior. We estimated the underlying function using piecewise polynomial fits

(cubic splines partitioned in three equal segments across the 0.1- to 1.3-mL range).

As observed previously (8), monkeys were risk-seeking for small rewards (convex utility function between 0.1 and 1 mL) and risk-avoiding for larger rewards (concave utility function above 1 mL) (Fig. 3A). To verify the predictive power of the utility functions, we used these utility functions to calculate the EUs of new two-outcome gambles. We then measured psychometrically the animals' CEs for these gambles. The predicted EUs matched the utility of the measured CEs (Deming regression; Fig. S2D). To assess further the predictive power of the curvature of the utility function, we subtracted the utility of the EV of the gambles from the EU and from the utility of the CE. Even after removing the intrinsic correlation between EVs and EUs, the utility functions continued to predict the utility of new binary gambles (Fig. S2E).

Using the fractile procedure on smaller ranges of the utility function (0.1–1.0 mL and 1.0–1.3 mL) led to very similar overall curvature between the compound and the whole-range utility function (Fig. S2F). A Deming regression showed good correlation between the CEs predicted from the compound utility function and the whole-range utility function ($R^2 = 0.99$; Fig. S2G). This result indicates little to no range adaptation of these utility functions in these experiments lasting several months in each animal.

Reaction times are often used as a blunt proxy for subjective reward valuations. In an imperative task without choice, the animals moved the joystick to a single safe or two- or three-outcome gamble. The reaction times correlated better with the EUs of gambles than with their EVs ($R^2 = 0.76$, $P = 0.012$ for EV compared with $R^2 = 0.80$, $P < 0.01$ for utility, single linear regressions; Fig. S2H and I, respectively). These data are compatible with the notion that the animals' behavior was based on nonlinear, subjective utility rather than on linear, objective reward magnitude.

We performed a control experiment to ensure that the small rewards (<0.1 mL) did not have negative net utility due to the effort cost associated with joystick movement. We limited the maximum reward per joystick movement to 0.1 mL and tested monkey B for 2 d in an imperative task for 2.5–3.0 h each day, resulting in 1,500 trials (84.2% correct). This result indicated that the net utility provided by a reward of 0.1 mL minus the effort cost of joystick movement was not negative and that the entire range of 0.1–1.3 mL of juice was within the gain domain.

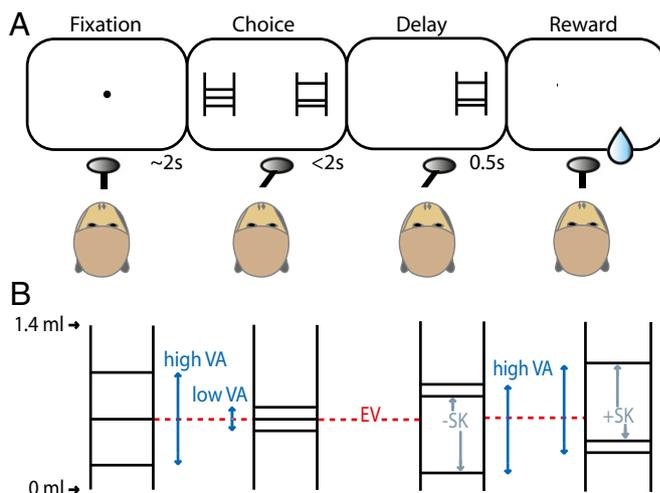


Fig. 2. Behavioral task and gambles. (A) Binary choice task. The animal chose one of two gambles with a joystick and received one of the chosen gamble's rewards. Bar heights in each gamble indicate reward magnitude, and each reward was delivered pseudorandomly with equal probability (here, $P = 1/3$). (B) Typical gambles with three pseudorandomly alternating equiprobable outcomes (each $P = 1/3$). The EV, VA, and SK of each gamble were set independent of each other.

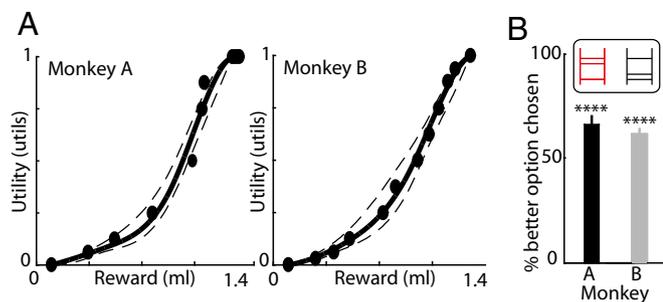


Fig. 3. Utility functions and first-order stochastic dominance. (A) Utility functions obtained from CEs of equiprobable, two-outcome gambles, using the fractile method (*SI Methods*). Data points represent the mean CE ($n = 43$ CEs across 5 d and $n = 106$ CEs across 14 d for monkeys A and B, respectively). Black curves indicate best-fitting polynomials from cubic spline fitting. Dashed lines represent 95% confidence intervals (obtained by bootstrapping). (B) Direct choices follow first-order stochastic dominance ($n = 187$ trials across 5 d and $n = 500$ trials across 13 d for monkeys A and B, respectively). Error bars indicate SEMs across days. **** $P < 10^{-5}$.

First-Order Stochastic Dominance: Preferring “More” to “Less”. First-order stochastic dominance describes the property of a gamble with outcomes at least as good as another gamble and with at least one strictly better outcome, as illustrated by cumulative distributions (Fig. S3A and *SI Methods*). This dominance relationship requires minimal assumptions that are fulfilled by the estimated animals’ utility function (strictly increasing value function) (Fig. 3A). Before using the three-outcome gambles shown in Fig. 2B more extensively, we investigated whether the monkeys’ choice behavior reflected a meaningful understanding of the bar stimuli and their associated probabilities. The animals’ choices followed the first-order stochastic dominance of three-outcome gambles that differed from each other only by one of the outcomes offered (Fig. 3B, *Inset*). In such pairs of gambles, the gamble with the higher outcome is first-order stochastically dominant, and thus should be chosen in more than 50% of trials (17). Indeed, both monkeys preferred the dominant gamble in >60% of trials, with little day-to-day variation ($P < 10^{-5}$ for both animals, two-tailed binomial test; Fig. 3B). Thus, the bar stimuli predicting three-outcome gambles were well understood, and the behavior was meaningful and consistent with utility maximization.

The presence of dominated choices in a minority of trials may reflect the fact that the dominant gamble was negatively skewed and that its visual stimulus resembled the visual stimulus of the dominated gamble in the third-order stochastic dominance test (discussed below). Further factors could be inattention, exploration for possibly improved outcomes, and neuronal noise during decision making. Future experiments with more daily trials of a given test might dissociate between these factors.

Variance-Risk and Second-Order Stochastic Dominance. Second-order stochastic dominance is a property of a gamble that describes how a mean-preserving spread (identical EV but greater variance) of the reward distribution should influence choices, given a specific variance-risk preference (17) (Fig. S3B and *SI Methods*). The ascending or descending direction of second-order stochastic dominance depends on the animal’s valuation of variance-risk (13, 14). Specifically, a higher variance gamble would be (descendingly) second-order stochastically dominant if variance-risk adds to utility. We determined second-order stochastic dominance between the gambles to assess how the animals’ valuations of variance-risk, as expressed by EUs, predicted the animals’ choices.

To assess the influence of variance-risk on the animals’ valuations of the gambles, we placed three-outcome gambles with identical EV but different variance (mean-preserving spread) (Fig. 2B) on convex parts of the utility functions and calculated the gambles’ EU. The gamble with a greater variance was associated with a higher EU than the gamble with a lower variance

(Fig. 4A). Thus, the animals’ utility function defined the gamble with the higher mean-preserving spread as second-order stochastically dominant.

We then estimated directly the CEs of these gambles. The higher variance gamble had significantly greater CEs than the lower variance gamble ($P < 10^{-4}$ for both monkeys, two-sample t test; Fig. 4B). Moreover, the CEs predicted from the gambles’ EUs (Fig. 4B, red dots) did not differ significantly from the psychometrically estimated gambles’ CEs ($P > 0.25$ for both gambles in both monkeys, one-sample t test). Thus, the animals’ utility functions and the CEs provided similar estimation of value for gambles varying only in variance-risk.

To test whether these valuations corresponded to observable variance-risk preferences, we examined direct choices between gambles with different mean-preserving spreads and found that both animals preferred the higher variance gamble, shown above to be second-order stochastically dominant ($P < 0.01$ and $P < 10^{-27}$ for monkeys A and B, respectively; two-tailed binomial tests; Fig. 4C). Further tests refined this finding. When the lower, convex part of the utility function assigned (descending) second-order stochastic dominance to the higher variance-risk gamble, the animal accordingly preferred this riskier gamble ($P < 10^{-10}$ for monkey B, two-tailed binomial test; Fig. 4D, *Left*). By contrast, when the upper, concave part of the utility function assigned (ascending) second-order stochastic dominance to the lower variance-risk

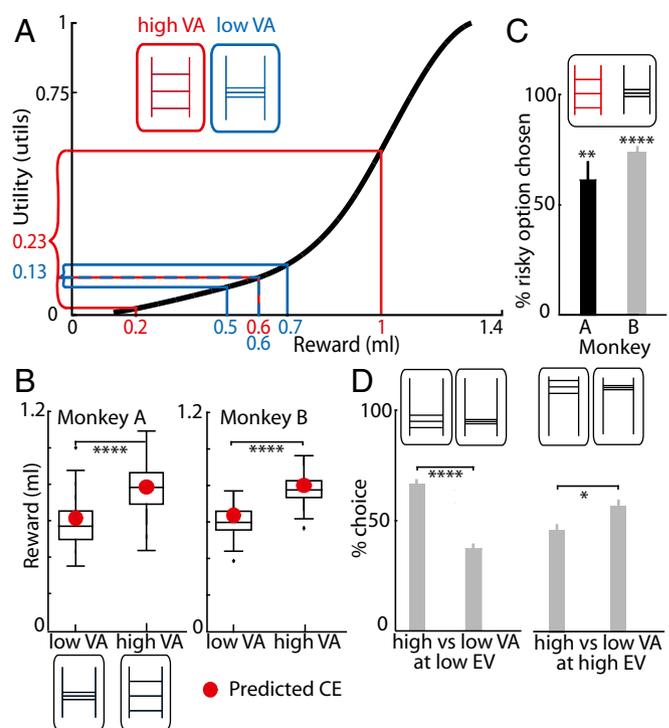


Fig. 4. Second-order stochastic dominance. (A) EUs of symmetrical, three-outcome, variance-risk gambles derived from the fitted utility function of monkey A (same as shown in Fig. 3A, *Left*). The wider mean-preserving spread gamble (red) had a higher EU than the smaller spread gamble (blue). Colored numbers to the left indicate EUs of gambles. (B) Distribution of CEs for mean-preserving spread gambles (same EV but different VA) ($n = 60$ CEs and $n = 80$ CEs for monkeys A and B, respectively). Red dots show CEs predicted from EUs. **** $P < 10^{-4}$. (C) Direct choices follow second-order stochastic dominance for risk-seeking ($n = 199$ trials across 4 d and $n = 500$ trials across 14 d for monkeys A and B, respectively). Error bars indicate SEM across days. ** $P < 0.01$; **** $P < 10^{-27}$. (D) Direct choices follow the changing directions of second-order stochastic dominance across the convex-concave utility function ($n = 500$ trials across 4 d and $n = 500$ trials across 5 d for low EV and high EV, respectively). Monkey B switched from being variance-risk-seeking at low EVs to variance-risk averse with higher EVs. **** $P < 10^{-10}$; * $P = 0.016$.

gamble, the animal accordingly preferred the lower variance-risk gamble ($P = 0.016$ for monkey B, two-tailed binomial test; Fig. 4D, Right). However, as with first-order stochastic dominance, the animals chose the dominated gambles in a nonnegligible fraction of trials.

In many of the gamble pairs, the dominant gamble also had the largest possible outcome, indicated by the highest stimulus bar [Fig. 4C and D (Left), Insets], and the monkeys might have implemented simple “choose gamble with highest bar/highest top outcome” heuristics. However, in the risk-averse range, at high EV, the animal preferred the gamble with the lower maximum outcome [Fig. 4D (Right), Inset]. Therefore, the highest bar’s height would not explain the observed preferences. A similar conclusion had been reached in experiments investigating variance-risk preferences with fractal stimuli (4).

Taken together, the observance of second-order stochastic dominance defined by the valuation of variance-risk suggests that the animals’ behavior was governed by a meaningful representation of nonlinear utility rather than by the EV (which remained constant between gambles).

Skewness-Risk and Third-Order Stochastic Dominance. Third-order stochastic dominance is a property of a gamble that describes how the degree of asymmetry of the reward distribution (with

constant EV and variance) should influence choice preferences (18) (Fig. S3C and SI Methods). The direction of the third-order stochastic dominance between the gambles depends on the animal’s valuation of skewness-risk (14). Specifically, a positively skewed gamble would be third-order stochastically dominant if positive skewness-risk adds to utility. We determined the direction of third-order stochastic dominance to assess how the animals’ valuation of skewness-risk, as expressed by EUs, predicted their choices.

We calculated the EUs of skewed three-outcome gambles with identical EV and variance (Fig. 5A, Inset). In both monkeys, the positively skewed gamble had a higher EU than the negatively skewed gamble (Fig. 5A). This difference defined the positively skewed gamble as third-order stochastically dominant.

We then estimated the CEs and found significantly greater CEs for the positively skewed gamble compared with the CEs of the negatively skewed gamble [$P < 10^{-3}$, $n = 6$ full psychometric curves; $P < 10^{-4}$, $n = 40$ parameter estimation by sequential testing (PEST) procedures; two-sample t tests; Fig. 5B]. Indeed, the CEs increased as the skewness of the gamble went from negative through zero to positive ($\beta = 0.11$, $P < 10^{-3}$ and $\beta = 0.094$, $P < 10^{-4}$ for monkeys A and B, respectively; linear regression; CEs acquired by PEST; Fig. 5C). The CEs predicted from the gambles’ EUs (Fig. 5C, red dots) did not differ significantly

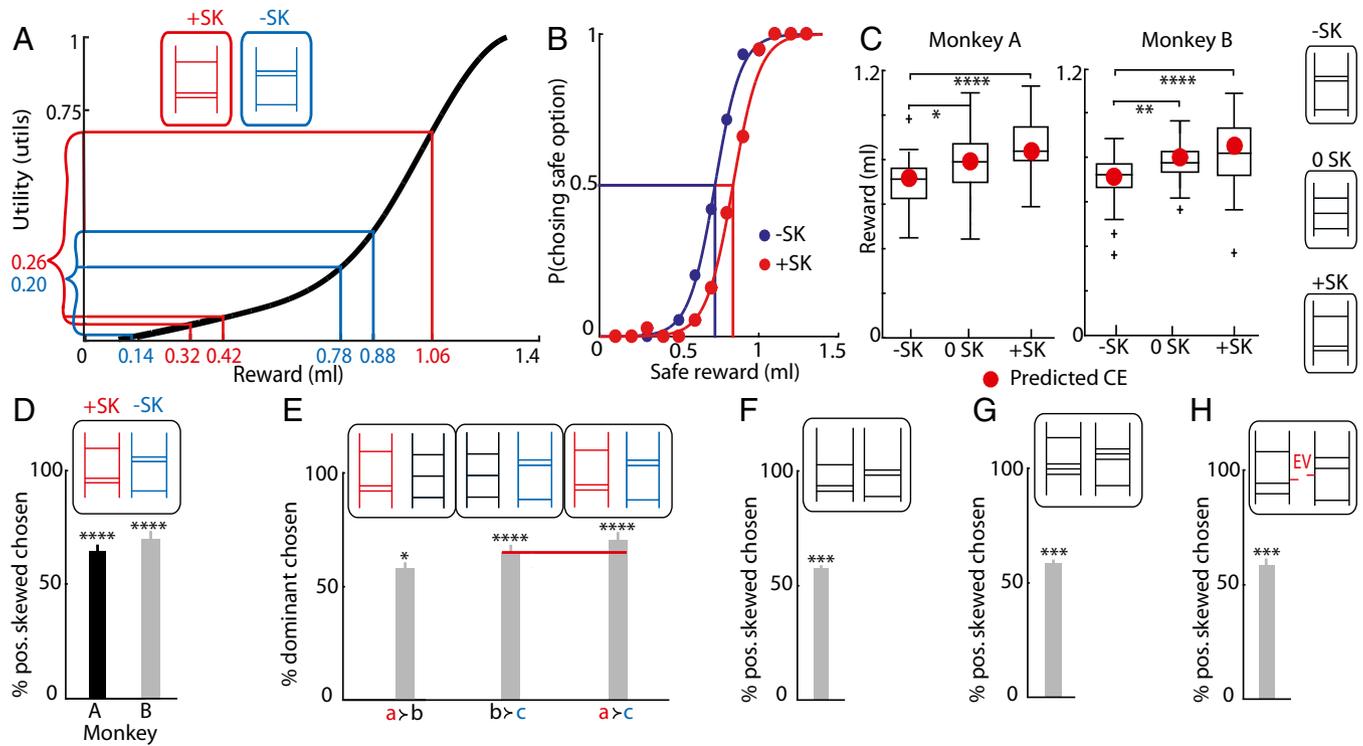


Fig. 5. Third-order stochastic dominance. (A) EUs of skewed gambles derived from the utility function of monkey A. The positively skewed gamble (red, +SK) had greater EU than the negatively skewed gamble (blue, -SK). Gambles were matched for EV and variance. (B) Psychometric functions showing probabilities of choices for safe rewards with different magnitudes over negatively skewed (blue) or positively skewed (red) gambles in monkey B ($n = 716$ and $n = 754$ respectively). Gambles had identical EV and variance, and differed only in skewness. The curves were derived from logistic functions fitted to choice frequencies averaged over 50 trials per data point. (C) Distribution of CEs for gambles with the same EV and same variance, but with different skewness (left to right: -SK, 0 SK, +SK) ($n = 90$ CEs and $n = 120$ CEs for monkeys A and B, respectively). The central “0 SK” box plots are the same as for “high variance” in Fig. 4B. Red dots show the CEs predicted from EUs. $*P = 0.018$; $***P < 0.01$; $****P < 10^{-4}$. (D) Choice frequencies of positively (pos.) skewed over negatively skewed gambles with identical EV and variance ($n = 359$ trials in 7 d and $n = 500$ trials in 13 d). $****P < 10^{-5}$. (E) Satisfaction of transitivity. Monkey B chose, respectively, a positively skewed gamble over a symmetrical gamble ($a > b$), a symmetrical gamble over a negatively skewed gamble ($b > c$), and a positively skewed gamble over a negatively skewed gamble ($a > c$) ($n = 250$ trials in 6 d, $n = 250$ trials in 6 d, and $n = 500$ trials in 13 d, respectively). The horizontal red line defines strong transitivity. $*P = 0.023$; $****P < 10^{-5}$. (F) Skewness preference at low EV. Monkey B chose between a positively and negatively skewed gamble with identical EV and variance ($n = 500$ trials across 4 d). $****P < 10^{-5}$. (G) Skewness preference with four-outcome gambles (identical EV and variance). Monkey B chose a positively skewed gamble over a negatively skewed gamble ($n = 500$ trials across 5 d). $***P < 10^{-3}$. (H) Skewness preference opposite to physical reward magnitude. Monkey B chose between a positively skewed gamble (EV = 0.53 mL) and a negatively skewed gamble with slightly higher EV (0.56 mL) ($n = 500$ trials in 7 d). $***P < 10^{-3}$.

from the psychometrically estimated gambles' CEs ($P > 0.08$ for both skewed gambles in both monkeys, one-sample t test). Thus, the animals' utility functions and the CEs provided similarly valid estimations of value for gambles varying only in skewness-risk.

In direct choices between these gambles, both animals preferred the third-order stochastically dominant gamble ($P < 10^{-6}$ for both monkeys, two-tailed binomial test; Fig. 5D). Further tests confirmed and extended this finding. A positively skewed gamble was preferred to a symmetrical gamble, which was, in turn, preferred to a negatively skewed gamble (all gambles matched for EV and variance) in a strongly transitive (19) manner: Gamble a was chosen over gamble c more often than gamble a was chosen over gamble b , or than gamble b was chosen over gamble c ($P = 0.023$, $P < 10^{-3}$, and $P < 10^{-18}$ for monkey B for $a > b$, $b > c$, and $a > c$, respectively; two-tailed binomial test; Fig. 5E). This transitive preference ranking between the gambles matched the EUs and the third-order stochastic dominance ranking, demonstrating consistent choices under skewness-risk.

Several control tests confirmed the robustness and consistency of the positive skewness-risk preference. Similar positive skewness-risk preference was found with a pair of gambles that were set at a lower EV ($P < 10^{-3}$ for monkey B, two-tailed binomial test; Fig. 5F). With four-outcome gambles of identical EV and variance, the EUs were again higher for positive than negative skewness (0.263 vs. 0.214 utils, respectively), and the monkeys correspondingly preferred positive to negative skew ($P < 10^{-3}$ for monkey B, two-tailed binomial test; Fig. 5G). This result suggests that the observed positive skewness-risk preference was not exclusive to three-outcome gambles. The preference for positive skewness existed even when the positively skewed gamble had a lower EV than the negatively skewed gamble (matched for variance) ($P < 10^{-3}$ for monkey B, two-tailed binomial test; Fig. 5H). Despite the occasional choice of the dominated gamble, the data from these direct choices further demonstrate the power of the positive skewness-risk preference.

Taken together, the observance of third-order stochastic dominance suggests that the prediction of risk preferences by the utility function was not restricted to variance-risk.

EU Captures Variance and Skewness Information. To have predictive power, the EU of a gamble should represent all relevant information that forms the basis for choices under risk. We therefore investigated whether the variance and skewness of the three-outcome gambles added explanatory power to the EUs of the utility functions estimated from two-outcome, variance-risk gambles. We examined choices ($n > 5,500$) between three-outcome and one-outcome gambles. A logistic regression analysis (Eq. 1) showed that the choices of the monkeys depended on the utility of the safe option ($\beta = -6.3$, $P < 10^{-4}$ and $\beta = -10.5$, $P < 10^{-4}$ for monkeys A and B, respectively) and the utility of the gamble ($\beta = 7.8$, $P < 10^{-4}$ and $\beta = 13.9$, $P < 10^{-4}$ for monkeys A and B, respectively), such that a greater utility of the safe option made choosing the gamble less likely and a greater utility of the gamble made choosing it more likely. Adding variance of the gamble (Eq. 2) to Eq. 1 did not significantly improve the model ($P = 0.99$ and $P = 0.88$ for monkeys A and B, respectively; F -test on Eqs. 1 and 2); likewise, adding skewness did not significantly improve the model (Eq. 3) ($P = 0.99$ and $P = 0.88$ for monkeys A and B, respectively; F -test on Eqs. 1 and 3). Thus, the utility functions estimated from two-outcome gambles fully represented the valuation of our three-outcome gambles varying in variance-risk or skewness-risk. EU alone captured all properties of the gambles that are relevant for the observed preferences.

Discussion

This study investigated whether utility functions in rhesus monkeys explain choice preferences for distinct forms of risk: variance-risk and skewness-risk. We used choices between safe rewards and binary gambles to derive empirical utility functions; these utility functions were convex at lower reward values and concave at higher reward values. The EU of each gamble calculated

from these functions was used to determine the direction of second-order and third-order stochastic dominance relationships. We then used direct choices to determine whether the animals' preferences abided by these stochastic dominance relationships. The monkeys preferred variance-risk at low reward values and showed the opposite preference at higher reward values (variance-risk avoidance), as predicted by their utility function. The same utility functions also predicted the animals' preference for positively over negatively skewed gambles (matched for EV and variance-risk). These findings show that monkeys behaved as if they used a representation of utility when making choices under variance-risk and skewness-risk. This result extends the validity of utility as a theoretical internal measure of value from humans to monkeys.

The estimation of utility functions allowed us to define the direction of second-order and third-order stochastic dominance (13, 14). We tested only one statistical moment at a time (variance or skewness) while holding all other moments constant, and used equal probabilities for all possible outcomes [avoiding differences in probability distortion (20–23)]. With these controls, positive skewness-risk preference does not derive from variance-risk preference (11) (Fig. 1). Direct choice data under variance-risk and skewness-risk could therefore be rigorously linked to predictions from utility functions. These results provide a coherent explanation for the independently observed stronger variance-risk preferences and higher CEs and EUs with increased mean-preserving spreads in monkeys (4, 8). The results also account for the positive skewness-risk preference reported here. To our knowledge, this study represents the first time that skewness-risk preference has been predicted by an empirical utility function in monkeys. This result could also help to explain monkeys' choices between nonexplicit pictures associated with complex, skewed reward distributions with various magnitudes and probabilities in a study that did not assess utility (24). Humans also display positive skewness preferences (25, 26).

In this study, we formally and precisely linked preferences under variance-risk and skewness-risk to experimentally estimated utility functions. Importantly, the utility functions estimated from choices between simple two-outcome gambles and safe options accounted for the variance-risk and skewness-risk preferences in direct choices between three-outcome options. This link from the domain of utility function to the domain of direct choices is crucial for understanding internal utility representations, because experimentally measured choices reveal risk preferences in the most direct way. Indeed, direct choices and CE comparisons can sometimes yield contradicting preference rankings (12, 27). We further investigated the predictive power of these empirically derived utility functions with logistic regression analysis and found that our utility measurements seemed to capture all of the effect of variance-risk and skewness-risk on choices. Overall, this study highlights the great predictive power of utility functions for explaining the behavior of monkeys under variance-risk and skewness-risk.

The observed transition from variance-risk preference with small stakes to variance-risk aversion with more substantial outcomes is consistent with human risk tendencies (28) and confirms oculomotor choices in monkeys (3–5, 7, 8) [except for one study that varied reward value by thirst and reported primarily risk aversion with only very mild risk-seeking behavior with low values (6)]. Our observed pattern of risk preference cannot be explained by probability distortion because all outcomes had identical probability. The good task performance with 0.1 mL of maximal reward is evidence against the possibility that the initial convexity of the utility function is due to a net negative utility derived from low utility outcomes in the presence of movement effort. Thus, the convex-concave curvature of the utility function seems to reflect true reward valuations across the whole value range.

It has been suggested that variance-risk attitudes may result from a two-step process: first, the conversion of safe objective value into riskless utility and, then, the incorporation of an intrinsic variance-risk preference into a common utility signal (29). An extension of the model may postulate the existence of a third

step to process an intrinsic skewness-risk preference/aversion. However, in this study, the objective skewness of the gamble did not add any explanatory power beyond the utility of the gamble and the utility of the safe option, as shown using a logistic regression model. This result implies that the effects of skewness-risk on preferences are fully represented by empirical utility functions derived from choices under variance-risk.

The prediction of risk preferences by the EUs suggests a valid representation of the gambles' values by the monkeys' convex-concave utility function. Previous empirical work found utility function with similar curvature in humans (30) and in macaques (8). This type of nonlinearity (convex and then concave) is not uncommon and may originate from known properties of biological systems. Neuronal responses to stimulus strength in the primary visual cortex show similar nonlinearity (31). The initial convexity may represent a "threshold effect": A minimum stimulus strength is needed before a neuronal response is triggered, and doubling stimulus strength may therefore more than double the response intensity (32), resulting in supralinear curvature. By contrast, the ultimate concavity with higher reward magnitudes may represent an adaptive "saturation effect" that accounts for the fact that beyond a certain stimulus intensity, the neuronal response cannot increase further (33). Dopamine neurons' canonical encoding of reward prediction errors seems to reflect this convex-concave shape of the utility function (8). We may therefore conjecture that the nonlinearity of our empirical utility functions could also be due to such neural phenomena, a hypothesis worthy of further investigation.

Methods

Animals and Experimental Setup. Two male rhesus monkeys (*Macaca mulatta*) were used for this study (weighing 7.6 kg and 8.9 kg) during 6 mo and 4 mo of daily testing (2–3 h each day), respectively. The Home Office of the United Kingdom approved all experimental protocols and procedures. During experiments, animals sat in a primate chair (Crist Instruments) positioned 30 cm from a computer monitor. A joystick (Biotronix Workshop), restricted to left and right movements only, was used by the animals to report choices. Joystick position data and digital task event

signals were sampled at 2 kHz and stored at 200 Hz (joystick) or 1 kHz. Custom-made software (MATLAB; The MathWorks) running on a Microsoft Windows 7 computer controlled the behavioral tasks.

Behavioral Testing. We trained each animal to associate visual cues with different reward magnitudes and risk levels to investigate variance-risk; skewness-risk; first-, second-, and third-order stochastic dominance; whole-range psychophysics; and PEST in pseudorandomly interleaved trials. Each trial (Fig. 2A) began with a white spot appearing at the center of the monitor. The joystick had to be kept straight until the central spot disappeared after 1,500–2,500 ms, or else an error sound was played and the trial was aborted. As the central spot disappeared, two cues appeared at the left and the right of the monitor. The animal indicated within 2,000 ms its choice by moving a joystick toward the chosen cue. Then, the unchosen cue disappeared and the chosen cue remained on the monitor for an additional 500 ms with the joystick held in position to confirm the choice, or else an error sound occurred and the trial was aborted. The chosen reward was delivered at offset of the chosen cue, typically as the animal brought the joystick back to its initial position.

Logistic Regression. We analyzed all correct choices between a multiple-outcome gamble and a one-outcome safe option, using the following logistic regression models:

$$Y \sim \beta_0 + \beta_1 u(\text{safe}) + \beta_2 u(\text{gamble}), \quad [1]$$

$$Y \sim \beta_0 + \beta_1 u(\text{safe}) + \beta_2 u(\text{gamble}) + \beta_3 \text{var}(\text{gamble}), \quad [2]$$

$$Y \sim \beta_0 + \beta_1 u(\text{safe}) + \beta_2 u(\text{gamble}) + \beta_3 \text{skew}(\text{gamble}), \quad [3]$$

with u as utility, var as variance, and skew as skewness. Y was set to 1 whenever the gamble was chosen and to zero when the safe option was chosen. Models in Eqs. 2 and 3 were compared with the reduced model in Eq. 1 with the F -test (34) to investigate whether the variance and skewness of the gambles explained any of the variance in the choice data, beyond utility.

ACKNOWLEDGMENTS. We thank Aled David for invaluable help with animal training and Mark Dean (Columbia University) for helpful comments. This work was supported by the Wellcome Trust, European Research Council, and NIH Caltech Conte Center.

- Holt CA, Laury SK (2002) Risk aversion and incentive effects. *Am Econ Rev* 92(5):1644–1655.
- Caraco T, Chasin M (1984) Foraging preferences: Response to reward skew. *Anim Behav* 32:76–85.
- McCoy AN, Platt ML (2005) Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat Neurosci* 8(9):1220–1227.
- O'Neill M, Schultz W (2010) Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron* 68(4):789–800.
- So N, Stuphorn V (2012) Supplementary eye field encodes reward prediction error. *J Neurosci* 32(9):2950–2963.
- Yamada H, Tymula A, Louie K, Glimcher PW (2013) Thirst-dependent risk preferences in monkeys identify a primitive form of wealth. *Proc Natl Acad Sci USA* 110(39):15788–15793.
- Lak A, Stauffer WR, Schultz W (2014) Dopamine prediction error responses integrate subjective value from different reward dimensions. *Proc Natl Acad Sci USA* 111(6):2343–2348.
- Stauffer WR, Lak A, Schultz W (2014) Dopamine reward prediction error responses reflect marginal utility. *Curr Biol* 24(21):2491–2500.
- Rothschild M, Stiglitz JE (1970) Increasing risk: I. A definition. *J Econ Theory* 2(3):225–243.
- Von Neumann J, Morgenstern O (1944) *Theory of Games and Economic Behavior* (Princeton Univ Press, Princeton).
- Menezes C, Geiss C, Tressler J (1980) Increasing downside risk. *Am Econ Rev* 70(5):921–932.
- Grether DM, Plott CR (1979) Economic theory of choice and the preference reversal phenomenon. *Am Econ Rev* 69(4):623–638.
- Meyer J (1977) Second degree stochastic dominance with respect to a function. *Int Econ Rev* 18:477–487.
- Wong WK, Li CK (1999) A note on convex stochastic dominance. *Econ Lett* 62(3):293–300.
- Debreu G (1959) Cardinal utility for even-chance mixtures of pairs of sure prospects. *Rev Econ Stud* 26:174–177.
- Machina M (1987) Choice under uncertainty: Problems solved and unsolved. *J Econ Perspect* 1(1):121–154.
- Mas-Colell A, Whinston MD, Green JR (1995) *Microeconomic Theory* (Oxford Univ Press, Oxford).
- Chiu WH (2010) Skewness preference, risk taking and expected utility maximisation. *Geneva Risk Insur Rev* 35(2):108–129.
- Padoa-Schioppa C (2011) Neurobiology of economic choice: A good-based model. *Annu Rev Neurosci* 34:333–359.
- Allais M (1953) Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica* 21(4):503–546. French.
- Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.
- Prelec D (1998) The probability weighting function. *Econometrica* 66(3):497–527.
- Stauffer WR, Lak A, Bossaerts P, Schultz W (2015) Economic choices reveal probability distortion in macaque monkeys. *J Neurosci* 35(7):3146–3154.
- Strait CE, Hayden BY (2013) Preference patterns for skewed gambles in rhesus monkeys. *Biol Lett* 9(6):20130902.
- Wu CC, Bossaerts P, Knutson B (2011) The affective impact of financial skewness on neural activity and choice. *PLoS One* 6(2):e16838.
- Burke CJ, Tobler PN (2011) Reward skewness coding in the insula independent of probability and loss. *J Neurophysiol* 106(5):2415–2422.
- Slovic S, Lichtenstein S (1983) Preference reversals: A broader perspective. *Am Econ Rev* 73(4):596–605.
- Markowitz H (1952) The utility of wealth. *J Polit Econ* 60(2):151–158.
- Tobler P, Weber E (2014) Valuation for risky and uncertain choices. *Neuroeconomics: Decision Making and the Brain*, eds Glimcher PW, Fehr E (Academic, London), pp 149–172.
- Mosteller F, Nogee P (1951) An experimental measurement of utility. *J Polit Econ* 59(5):371–404.
- Carandini M, et al. (2005) Do we know what the early visual system does? *J Neurosci* 25(46):10577–10597.
- Tolhurst DJ, Dean AF (1987) Spatial summation by simple cells in the striate cortex of the cat. *Exp Brain Res* 66(3):607–620.
- Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: Contrast response function. *J Neurophysiol* 48(1):217–237.
- Snedecor GW, Cochran WG (1989) *Extension of the Analysis of Variance. Statistical Methods* (Iowa State Univ Press, Ames, IA).
- Luce DR (2000) *Utility of Gains and Losses: Measurement-Theoretic and Experimental Approaches* (Lawrence Erlbaum Associates, Mahwah, NJ).

Supporting Information

Genest et al. 10.1073/pnas.1602217113

SI Methods

Psychometric Estimations of CEs for Estimating Utility. We estimated utility functions by measuring the CEs of binary choices using quantitative psychometric procedures. We defined the CE as the amount of blackcurrant juice that was subjectively equivalent to the value associated with each specific gamble. We used two closely related procedures to measure CEs: psychometric testing across the whole reward range and the more efficient PEST procedure (35) that converges on the choice indifference point.

In the full psychometric test (Fig. 24), we randomly varied the safe reward across the whole range of values (flat probability distribution). The safe option value on a given trial was therefore independent of the animal's previous choices. We then estimated the probability with which monkeys were choosing the safe reward over the gamble for a wide range of reward magnitudes. We fitted the logistic function of the following form, weighted for trial numbers, on these choice data:

$$P(\text{choosing safe option}) = \frac{1}{1 + e^{-[\alpha + \beta(\text{safe option magnitude})]}}$$

where α is a measure of choice bias and β reflects sensitivity (slope). The CE of each gamble was then estimated from the psychometric curve by determining the point on the x axis that corresponded to 50% choice (indifference) on the y axis. To obtain a full psychometric curve, 100–300 trials were required.

Each PEST sequence began with the amount of safe reward being chosen randomly from the interval 0–1.4 mL. Based on the animal's choice between the safe reward and gamble, the safe amount was adjusted on the subsequent trial. If the animal chose the gamble on trial t , the safe amount was increased by ε on trial $t + 1$. However, if the animal chose the safe reward on trial t , the safe amount was reduced by ε on trial $t + 1$ (Fig. S2B). Initially, ε was large. After the third trial of a PEST sequence, ε was adjusted according to the doubling rule and the halving rule (31). Specifically, every time two consecutive choices were the same, the size of ε was doubled, and every time the animal switched from one option to the other, the size of ε was halved. Thus, the procedure converged by locating subsequent safe offers on either side of the true indifference value and reducing ε until the interval containing the indifference value was small. The size of this interval is a parameter set by the experimenter, called the exit rule. For our study, the exit rule was 0.1 mL. When ε fell below the exit rule, the PEST procedure terminated and the indifference value was calculated by taking the mean of the final two safe rewards. We collected about 80% of our CEs with this PEST procedure, each of which lasted for 15–20 trials for a given gamble.

We randomly intermingled PEST procedures between different choice sets in a given session, which effectively prevented the animals from pushing the indifference point toward higher rewards. The CEs from PEST procedures and the CEs from full psychometric curves (in which the animal had no control over the reward values, which were exclusively set by the experimenter) should therefore be equivalent. Indeed, direct comparisons of indifference points between PEST and full psychometric curves for two-outcome and three-outcome gambles revealed only insignificant variations (positively skewed gamble: mean CE = 0.83 mL and mean CE = 0.83 mL for psychometric and PEST tests, respectively, $P > 0.9$; negatively skewed gamble: 0.72 mL and 0.71 mL for psychometric and PEST tests, respectively, $P > 0.7$; two-sample t tests, $n = 6$ psychometric tests, $n = 40$ PEST tests).

Constructing Utility Functions with the Fractile Method. We used CEs obtained by our psychometric procedures to construct each monkey's utility function using the iterative, fractile method (2–4) in the range between 0.1 and 1.3 mL. The utility of 0.1 mL was arbitrarily set as 0 util, and the utility of 1.3 mL was set as 1 util. The CE of a binary, equiprobable gamble ($P = 0.5$ for each outcome) between 0.1 and 1.3 mL therefore had a utility of 0.5 util (0 util, $p + 1$ util, $1 - p$) (Fig. S2C, *Left*). Then, we created two gambles, one between the first CE and 0.1 mL, and one between the first CE and 1.3 mL. The CEs of these two new gambles had utilities of 0.25 util (0 util, $p + 0.5$ util, $1 - p$) and 0.75 util (0.5 util, $p + 1$ util, $1 - p$), respectively (Fig. S2C, *Center and Right*). Further iterations resulted in a more fine-grained function with more closely spaced utilities.

We then fitted the data acquired in each fractile procedure using local data interpolation [i.e., splines (MATLAB SLM tool)]. This procedure fits cubic functions on consecutive segments of the data and uses the least-squares method to minimize the difference between empirical data and the fitted curve. The number of polynomial pieces was controlled by the number of knots that were placed so as to partition the x axis in equal segments. We partitioned the x axis in three equal segments. We obtained confidence intervals by randomly selecting with replacement (bootstrapping) one CE per utility level tested and fitting a curve to these CEs. We could then find the 95% confidence intervals among the CEs predicted by these curves for each of 1,000 values along the utility axis (partitioning the utility axis in 0.001-util intervals).

Stochastic Dominance. Stochastic dominance is a relationship between a pair of gambles based on the gambles' reward probability distribution and the utility function of the chooser (13, 14).

For any chooser with a strictly increasing value function, a gamble is first-order stochastically dominant over another gamble when its outcomes are at least as good as the other gamble and at least one of its outcomes is strictly better. For instance, the gamble 0.2 mL, $P = 0.33$; 0.8 mL, $P = 0.33$; 1.0 mL, $P = 0.33$ is dominant over the gamble 0.2 mL, $P = 0.33$; 0.4 mL, $P = 0.33$; 1.0 mL, $P = 0.33$. Indeed, the 0.2 mL, $P = 0.33$ and 1.0 mL, $P = 0.33$ outcomes are common to both gambles, whereas the 0.8 mL, $P = 0.33$ outcome is strictly better than the 0.4 mL, $P = 0.33$ outcome. Although first-order stochastic dominance implies a greater EV, the reverse implication is not true. For instance, the gamble 0.1 mL, $P = 0.33$; 0.8 mL, $P = 0.33$; 1.0 mL, $P = 0.33$ has a greater EV than the gamble 0.2 mL, $P = 0.33$; 0.4 mL, $P = 0.33$; 1.0 mL, $P = 0.33$; however, it is not stochastically dominant because the smallest outcome (0.1 mL, $P = 0.33$) is not at least as good as 0.2 mL, $P = 0.33$. Graphically, a first-order stochastically dominant gamble's cumulative distribution will be "shifted to the right" for at least one value and will never cross the cumulative distribution of the dominated gamble (17) (Fig. S34). Any chooser with a strictly increasing value function should prefer first-order stochastic dominant options (but will not necessarily always prefer options with a greater EV because of risk preferences).

For any decision maker with a convex utility function, a gamble is (descendingly) second-order stochastically dominant over another gamble when it is a mean-preserving spread (14). For instance, for variance-risk-seeking choosers, the gamble 0.2 mL, $P = 0.33$; 0.6 mL, $P = 0.33$; 1.0 mL, $P = 0.33$ is dominant over the gamble 0.5 mL, $P = 0.33$; 0.6 mL, $P = 0.33$; 0.7 mL, $P = 0.33$. Note that both gambles share the same EV but the dominant gamble has a greater variance; thus, it is a mean-preserving spread of the dominated gamble. Graphically, (descending) second-order

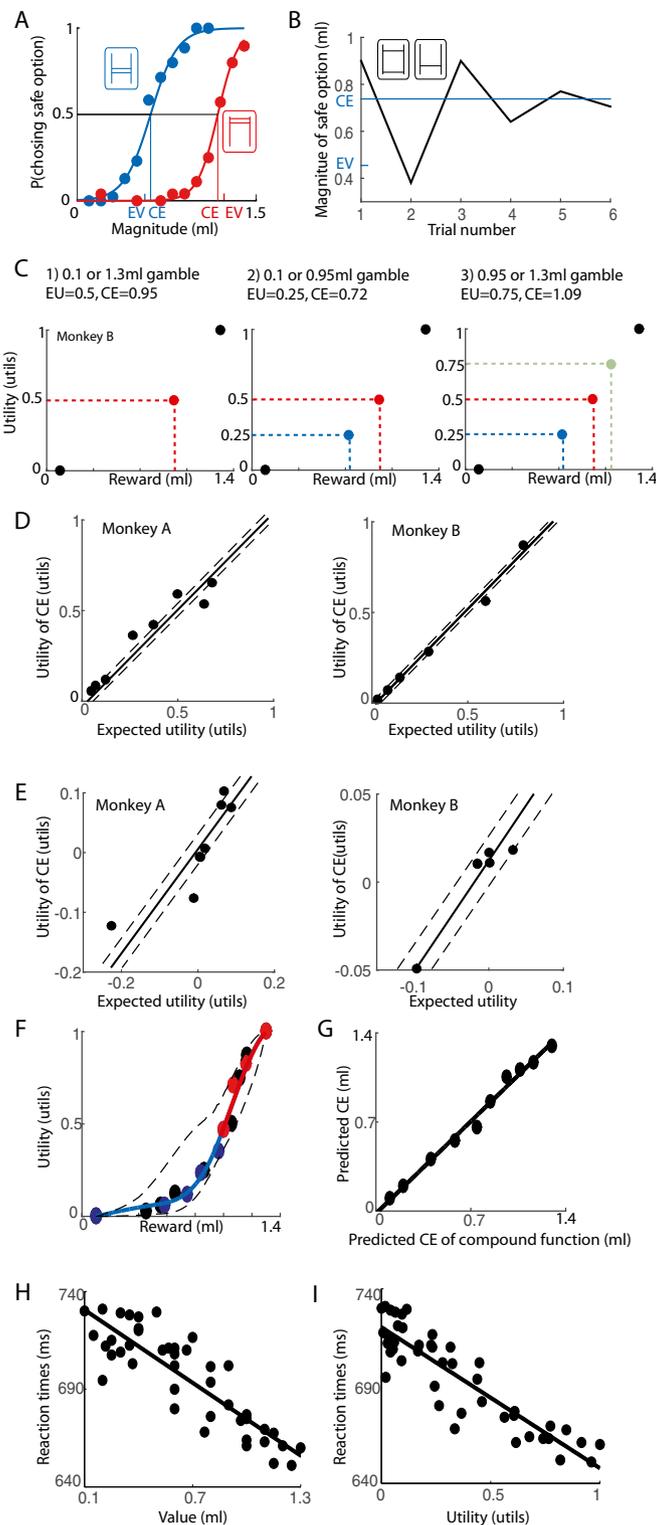


Fig. S2. Utility functions. (A) Full psychometric assessment of CEs. The monkey chose between a safe, adjustable reward and a constant gamble. The safe reward amount varied randomly on each trial between 0.1 and 1.4 mL, independent of the animal's previous choice. The curves were derived from logistic functions fitted to choice frequencies averaged over 30 trials per data point ($n = 300$ trials for each of the low-EV and high-EV gambles). For each gamble, the vertical line indicates choice indifference (which defines the CE) and the small marker indicates the gamble's EV. For the high-magnitude gamble (1.0 or 1.3 mL) in red, the CE is lower than the EV, indicating risk aversion. For the low-magnitude gamble (0.5 or 0.7 mL) in blue, the CE is higher than the EV, indicating risk-seeking. (B) PEST procedure. Black traces show tests for a binary gamble (0.1 mL, $P = 0.5$; 1 mL, $P = 0.5$). The gamble remains unchanged throughout a PEST sequence, whereas the safe amount is adjusted based on the previous choice following the iterative PEST protocol (*SI Methods*). Each data point shows the safe value offered on that trial (trial 2 is shown in the *Inset*). The CE of each gamble (blue line) is estimated by averaging across the final two safe rewards of each PEST sequence. The short blue marker indicates the gamble's EV. For this low-magnitude gamble, the CE is higher than the EV, indicating risk-seeking.

Legend continued on following page

(C) Fractile method for estimating utility using two-outcome gambles. In step 1, the CE of the gamble between 0.1 and 1.3 mL (each $P = 0.5$) was measured using PEST (*SI Methods*) (here, CE = 0.95 mL), which corresponds to 0.5 util. In step 2, the CE of the gamble between 0.1 and 0.95 mL was measured (CE = 0.72 mL), corresponding to 0.25 util. In step 3, the CE of the gamble between 0.95 and 1.3 mL (CE = 1.09 mL) corresponds to 0.75 util. (D) Out-of-sample predictions for eight and six new gambles not used for constructing the utility functions ($n = 38$ and $n = 30$ CEs for monkeys A and B, respectively). All expected utilities were derived from the respective utility functions shown in Fig. 3A. The black line represents the fit of a Deming regression on the mean utility, and dashed lines indicate the 95% confidence interval from the regression. (E) Same data as in D, but with the EVs removed from the predicted and measured values. The solid line represents the fit of a Deming regression, and dashed lines indicate the 95% confidence interval from the regression. (F) Compounded smaller range utility functions for monkey A, measured from CEs of binary equiprobable gambles using the fractile method. The black data points represent the mean CEs for a fractile method over the full range of 0.1–1.3 mL ($n = 40$ CEs), the red data points represent the mean CEs for a fractile method over the restricted range of 1.0–1.3 mL ($n = 6$ CEs), and the blue data points represent the mean CEs for a fractile method over the restricted range of 0.1–1 mL ($n = 19$ CEs). The blue-red curve indicates the best-fitting function for the black data points obtained from cubic spline fitting. The dashed lines represent 95% confidence intervals (obtained by bootstrapping). (G) Good fits between compound, partial-range, and full-range utility functions. For the out-of-sample predictions, 10 gambles not used for constructing any of these utility functions were placed on both the compound utility function shown in F (0.1–1.0 mL, blue; 1.0–1.3 mL, red) and the full-range 0.1- to 1.3-mL utility function shown in F (black dots, same as shown in Fig. 3A, *Left*). Then, the utilities of the outcomes of the 10 gambles were read on each utility function (y axis), their EUs were calculated, and their CEs were predicted on the x axis. The line shows the fitted Deming regression between the CEs predicted from the compounded utility function (blue and red in F) and the CEs predicted from the full-range utility function (black dots in F) ($R^2 = 0.99$). (H) Reaction times of monkey B for choosing one-outcome, two-outcome, or three-outcome options in an imperative task plotted against the value of the options ($n = 2,240$ trials across 10 d). (I) Same as in H, but with reaction times plotted against the utility of the options. Both the value and the utility of the options could explain the variance in the observed reactions times, but the utility explained more variance.

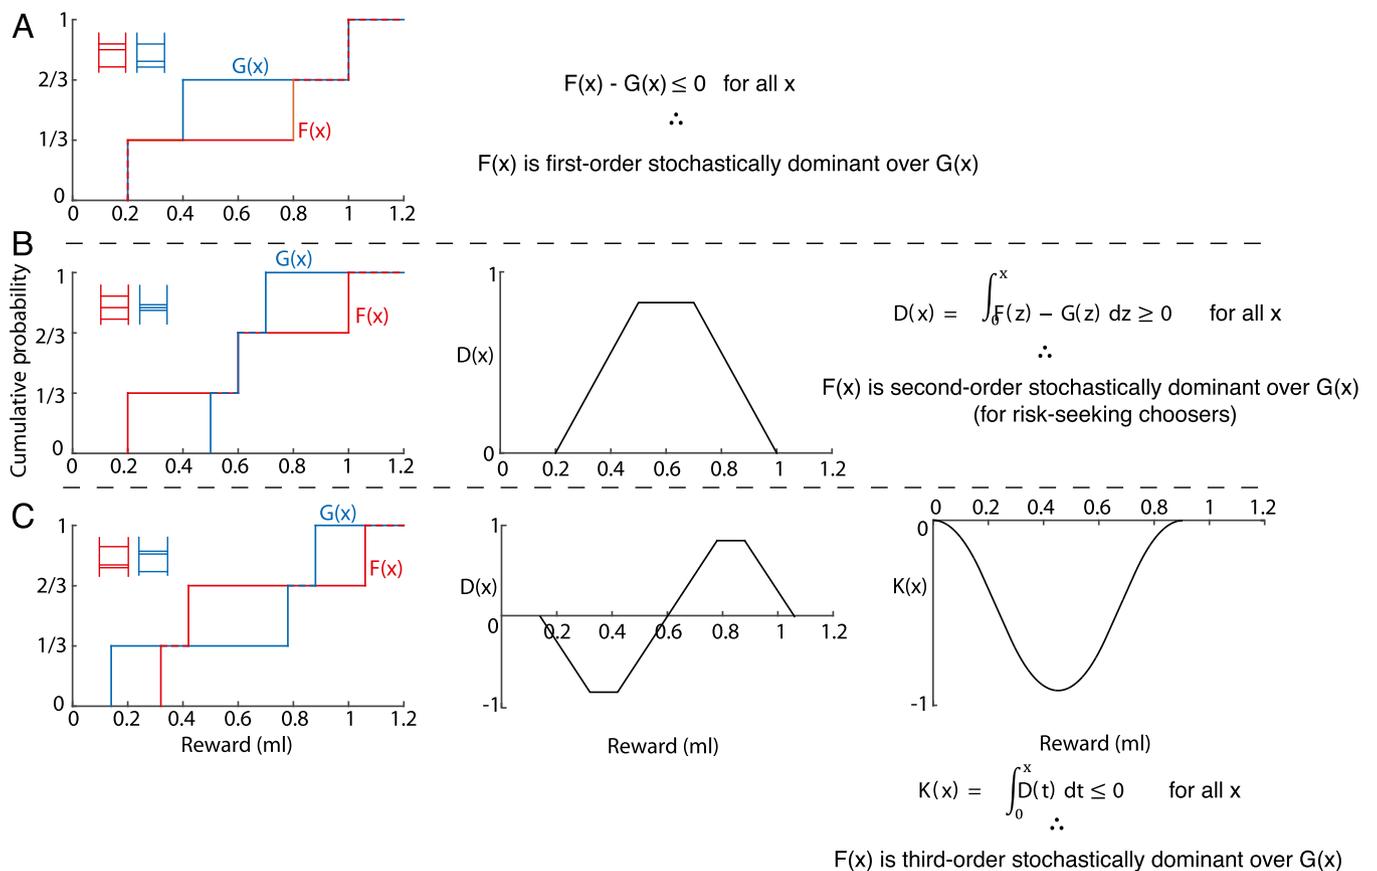


Fig. S3. Cumulative distributions of gambles defining stochastic dominance. (A) First-order stochastic dominance. Cumulative distribution of two gambles with distribution $F(x)$ and $G(x)$. The red gamble [$F(x)$] dominates the blue gamble [$G(x)$] because its values are always equal to or to the right of the values of $G(x)$. (B) Second-order stochastic dominance for risk-seeking. (Left) Cumulative distribution of two gambles with distribution $F(x)$ and $G(x)$. (Center) Integral of the difference between $F(x)$ and $G(x)$ over a range of x . The red gamble [$F(x)$] is a mean-preserving spread of the blue gamble [$G(x)$] and dominates the blue gamble in risk-seekers. (C) Third-order stochastic dominance. (Left) Cumulative distribution of two gambles with distribution $F(x)$ and $G(x)$. (Center) Integral of the difference between $F(x)$ and $G(x)$ over a range of x . (Right) Integral of the integral of the difference between $F(x)$ and $G(x)$ over a range of x .