

Feedback Procedures

Ryan Jessup, Anthony Bishara, and Jerome Busemeyer

Research Article

Feedback Produces Divergence From Prospect Theory in Descriptive Choice

Ryan Jessup, Anthony Bishara, and Jerome Busemeyer

Department of Psychological and Brain Sciences, Program in Cognitive Science, Indiana
University

Address correspondence to Ryan Jessup, Psychological and Brain Sciences, Indiana University,
1101 E. 10th St., Bloomington, IN 47405, e-mail: rkjessup@indiana.edu.

ABSTRACT—A recent study demonstrated that individuals making experience-based choices underweight small probabilities, in contrast to the overweighting observed in a typical descriptive paradigm. We tested whether trial-by-trial feedback in a repeated descriptive paradigm would engender choices more correspondent with experiential or descriptive paradigms. The results of a repeated gambling task indicated that individuals receiving feedback underweighted small probabilities, relative to their no-feedback counterparts. These results implicate feedback as a critical component during the decision-making process, even in the presence of fully specified descriptive information. A model comparison at the individual-subject level suggested that feedback drove individuals' decision weights toward objective probability weighting.

On June 19, 1986, a college student died in his dorm room from cocaine intoxication. This was no ordinary student: It was Len Bias, who 2 days earlier was drafted second overall by the Boston Celtics basketball team. Bias—whose talent was sometimes compared to Michael Jordan’s—was about to sign a contract worth more than he could have imagined while growing up; unfortunately, a poor decision cut short all of his opportunities.

Even though it was well established by that time that cocaine use, in rare circumstances, could result in death, it might be said that Bias “underweighted” the small probability of this most terrible outcome. He had observed the repeated choices of his friends who consistently obtained the pleasant outcome of drug use, as opposed to the fatal outcome (“Ex-Teammate Tells of Bias Drug Role,” 1987). This is a real-world example of the important difference between descriptive and experiential choice. Although Bias almost certainly knew descriptively that cocaine use had the rare event of death as a possible outcome, he also experienced (through his friends' and his own use) that this tragic outcome was rare. Recent laboratory experiments have confirmed that people make different decisions depending on whether they are based on experiential or descriptive information (Barron & Erev, 2003). We are interested in determining (a) what makes people choose differently across these two paradigms and (b) what theoretical component of the decision process incorporates the dynamics causing this difference.

The descriptive choice paradigm is the basis of most models of decision making. In the descriptive paradigm, participants choose between options, with the possible outcomes and probability of each outcome fully described.¹ For many years, economists assumed that individuals behaved according to a rational theory called expected utility theory (EUT; Bernoulli, 1738/1954; von Neumann & Morgenstern, 1944); Kahneman and Tversky (1979;

Tversky & Kahneman, 1992) later proposed prospect theory (PT) to account for many systematic violations of EUT observed in descriptive choice. PT incorporated into EUT three new concepts (among others): overweighting of small probabilities (small-probability events are considered more likely to occur than they do in reality), the risk reflection effect (people avoid risk when choosing between gains, but seek risk when choosing between losses), and loss aversion (losses have more impact than gains of equal magnitude). The first two concepts allow PT to account for an interesting pattern of preferences among risky choices. For example, consider the following option pairs:

A1: 3 for sure

B1: 4 with probability .8; 0 otherwise

A2: 3 for sure

B2: 32 with probability .1; 0 otherwise

When all outcomes are positive, A1 is preferred over B1, whereas B2 is preferred over A2; but when all outcomes are negative, B1 is preferred over A1, whereas A2 is preferred over B2. Thus, the following four-fold pattern is observed in descriptive choice: risk aversion when the risky gamble has a high probability of resulting in a win or a low probability of resulting in a loss, and risk seeking when the risky gamble has a low probability of resulting in a win or a high probability of resulting in a loss.

Descriptive choice is possible for humans because they can understand and work with symbols representing probabilities and values, but nonhumans are unable to choose descriptively; therefore, choices by animals are experientially based. In contrast to the

predictions of PT, the choices of bees (Real, 1991) are best described as reflecting underweighting of small probabilities, a behavioral pattern also observed in human experiential choice (Benartzi & Thaler, 1995; Hau, Pleskac, Kiefer, & Hertwig, in press; Newell & Rakow, in press; Thaler, Tversky, Kahneman, & Schwartz, 1997; Weber, Shafir, & Blais, 2004)

To gain a better understanding of the difference between experiential and descriptive choice, Barron and Erev (2003) examined several core concepts of PT using an experiential choice paradigm. Specifically, participants learned about the distribution of outcomes through repeated choice and subsequent feedback; no descriptive information about possible outcomes and their likelihoods of occurring was given, beyond the postchoice feedback. Remarkably, in all but one of the cases tested, the patterns of preference were the opposite of what PT predicted. For example, the underweighting of small probabilities, rather than the overweighting posited by PT, was again observed. In other words, the converse of PT best predicted the results!

There are several possible explanations for this discrepancy in behavior between the descriptive and experiential paradigms. First, the experiential paradigm requires individuals to learn about payoff distributions, whereas the descriptive paradigm informs individuals about payoff distributions; thus, learning is a key component of the experiential choice paradigm. Learning is typically not included in theories derived from descriptive choice, but theories derived from experiential choice include learning mechanisms, typically reinforcement learning (e.g., Busemeyer & Stout, 2002; Erev & Barron, 2005; Sutton & Barto, 1998).

Second, the experiential paradigm requires repeated choice, whereas the descriptive paradigm is often conducted using a single choice. Hertwig, Barron, Weber, and Erev (2004) sought to eliminate this methodological difference as an explanation by allowing participants in

the experiential condition to “sample” from the options in order to learn about the distribution of payoffs; the observed samples were not considered choices. After sampling to their satisfaction, participants made a single choice for which they received payment. Thus, participants in both the descriptive and the experiential conditions made one relevant choice; again, the pattern of behavior among participants in the experiential condition was converse to the predictions of PT. It is possible that the underweighting of rare events in this paradigm resulted simply from sampling bias (Fox & Hadar, 2006; Hertwig et al., 2004; Rakow, Demes, & Newell, 2007). That is, in the experiential condition, participants did not know the true payoff distribution: They only observed a small sample. With small samples, rare events are likely to be underrepresented, so resultant behavior would correspond with underweighting of small probabilities. For example, when a sample of seven draws is taken from a distribution paying 64 with a probability of .05, and paying nothing otherwise, the 64 outcome will rarely be sampled, effectively rendering the expected payoff equal to 0. However, even in experiments that use the experiential paradigm and have many trials (and hence large samples; e.g., Barron & Erev, 2003, used more than 100 trials), small probabilities are underweighted. So it cannot be the case that sampling bias accounts for all observed underweighting of small probabilities in the experimental paradigm.

Third, the experiential paradigm requires feedback, whereas the descriptive paradigm does not. Yechiam and Busemeyer (2006) observed that when given descriptive information and feedback about the outcome of their choice, individuals reverse preference as they gain more experience. Specifically, individuals were given an endowment and then chose repeatedly between the following two risky gambles:

A3: Lose 8 with probability .005; lose 2 otherwise

B3: Lose 300 with probability .005; lose 1 otherwise

Individuals initially preferred the riskier option (B3), but after feedback eventually came to prefer the safer option (A3).² However, some unique features of this task prevent one from drawing strong conclusions about PT from it: First, the payoff structure involved only risky gambles, as opposed to completely safe alternatives. Second, individuals were guaranteed to lose money on each trial. Third, the probabilities involved were extreme, and thus the larger losses were very rarely experienced. These features preclude this work from being considered a strong test of PT assumptions.

On the basis of the results just summarized, we hypothesized that feedback is the crucial arbiter dictating whether participants underweight or overweight small probabilities. To test this hypothesis in isolation, we designed a study in which the repeated nature of choice was held constant and the necessity of learning about payoff distributions was eliminated, so that only the feedback component varied. Further, we wanted to focus on the observation that when outcomes are positive, small probabilities are overweighted in descriptive choice but underweighted in experiential choice. Consequently, we tested both the upper and the lower portions of the probability spectrum, ensuring that any observed underweighting of small probabilities was not attributable to adherence to expected value in situations involving repeated choice.

We predicted that when outcome feedback was not given, individuals would conform to PT and overweight small probabilities, thus preferring a small, certain win over a risky, high-probability gamble, but preferring a risky, low-probability gamble over the same small, certain win. Conversely, we predicted that when individuals were given feedback, they would underweight small probabilities and show the opposite choice pattern.

Additionally, we were interested in discerning which theoretical component of the decision process incorporates the dynamics producing the behavioral disparities between the typical descriptive and experiential paradigms. We predicted that the theoretical analysis would reveal that the feedback manipulation affected a probability-weighting parameter, which would be in line with our prediction that people would overweight small probabilities when they did not receive feedback but underweight small probabilities when they did receive feedback.

The model we used for this theoretical analysis was a modified version of decision field theory (Busemeyer & Townsend, 1993), a dynamic and stochastic model of choice that accounts for a wide variety of empirical findings in both descriptive and experiential choice (Busemeyer & Diederich, 2002; Busemeyer, Johnson, & Jessup, 2006). In decision field theory, attention switches between outcomes, and the amount of attention applied to each outcome is a function of the probability of that outcome. This attention-switching activity results in a set of relative preferences for the options at each moment, and these preferences change over time until preference for one option surpasses a threshold, at which point that option is chosen. Other probabilistic choice models could have been applied to this data set; consequently, our analysis was not a test of decision field theory, but rather a tool for examining which parameters changed because of the feedback manipulation.

To address the prediction concerning the weighting of small probabilities, we used a probability-weighting function to modify the attention-switching activity of decision field theory (Prelec, 1998). This probability-weighting function allows small probabilities to be overweighted (i.e., the probability-weighting parameter, λ , is less than 1), underweighted (λ is greater than 1)

or objectively weighted (λ is equal to 1). The appendix presents the computations we used to implement decision field theory and the probability-weighting function.

METHOD

Twenty-nine participants were recruited from the Indiana University community and completed 120 trials in each condition in a repeated measures design; order of the conditions was counterbalanced. The high-probability (HP) condition consisted of a choice between the following two options:

Win 3¢ for sure

Win 4¢ with probability .8; win 0¢ otherwise

The low-probability (LP) condition consisted of these choices:

Win 3¢ for sure

Win 64¢ with probability .05; win 0¢ otherwise

Participants were randomly assigned to no-feedback ($n = 15$) and feedback ($n = 14$) groups. Both groups received all descriptive information about the gambles and made repeated choices. The only difference was that the feedback group received feedback indicating their winnings on the previous trial, whereas the no-feedback group did not.

Because we anticipated that individuals would experience boredom when confronted with the exact same decision repeatedly, especially when feedback was not given, we varied the win amount for the risky gamble slightly across trials, such that the minimum expected value for that gamble was equal to the value of the sure-thing option. In both conditions, five different expected values were shown for the risky gamble (range = 3–3.4), and these five values were equated across the two conditions. The risky gamble with an expected value of 3.2¢ was shown

twice as often as the other four gambles. After completing the task, participants rated the extent to which they were bored by it, using a Likert scale ranging from 1, very slightly or not at all, to 5, extremely.

Participants received instruction and training and completed several practice trials before beginning the experimental trials. They learned that the amount of money they would earn depended on their choices. On each trial, they observed pie charts indicating the amount of money that could be won; the size of the pie pieces indicated the probability of winning the displayed amounts. At the end of the experimental session, participants were paid \$7 for their time, as well as the amount of their cumulative winnings.

Note that four patterns of responding were possible: Individuals could adhere to expected value and consequently prefer the risky gamble in each pair, they could be risk averse and choose the sure thing in each pair, they could choose in accordance with PT and overweight small probabilities (choosing the sure thing in the HP condition and the risky gamble in the LP condition), or they could choose in opposition to PT and underweight small probabilities (choosing the risky gamble in the HP condition and the sure thing in the LP condition).

RESULTS

Statistical Analyses

Figure 1 shows the preference for the sure-thing option relative to the risky gamble in the feedback and no-feedback groups, separately for the two conditions. Results are plotted as a function of blocks of 10 trials. We hypothesized that descriptive choice with no feedback would result in overweighting of small probabilities, whereas descriptive choice with feedback would result in underweighting of small probabilities. This prediction was supported by the data. The

probability of choosing the sure thing was analyzed with a mixed analysis of variance, using condition (LP or HP) and block (12 blocks of 10 trials each) as within-subjects factors, and group (no-feedback or feedback) as a between-subjects factor. There was a significant Group \times Condition interaction, $F(1, 27) = 4.76$, $p < .05$, Cohen's $f = .41$. No other comparison was significant at the .05 level.

Whereas the no-feedback group's behavior generally conformed to theories derived from a descriptive paradigm (e.g., PT), participants in the feedback condition behaved otherwise, underweighting the small probabilities relative to their no-feedback counterparts. Consequently, feedback alone was enough to drive participants to underweight small-probability events. It is possible that the no-feedback group experienced more boredom in the task than the feedback group and that this boredom accounts for any observed behavioral difference. However, this is unlikely because responses to the boredom assessment did not differ significantly between the groups, $t(24) = 1.25$, $p = .22$.

Theoretical Analyses: Which Parameter Did Feedback Alter?

The feedback group may have appeared to underweight small probabilities relative to the no-feedback group, but which theoretical component in decision field theory varied between groups to produce this differential behavior? There are several possibilities, including changes to (a) the weight or attention given to probabilities, (b) the utility function capturing risk-seeking and risk-aversion tendencies, and (c) the threshold controlling the tendency to maximize utility or choose randomly; these concepts are represented by the free parameters λ , α , and θ , respectively.

Decision field theory was fit to the choice probabilities for each individual's data. Table 1 shows the medians of the best-fitting parameters that maximized the likelihood of each subject's choices. The likelihoods in this model were compared with those of a fully saturated baseline model, a model that perfectly reproduced the choice probabilities for each gamble pair. This comparison used the Bayesian information criterion (BIC; Schwarz, 1978), which penalizes models for each additional parameter; thus, adding extra parameters will not necessarily improve the BIC (see the appendix for information concerning the estimation procedure, baseline model, and computation of the BIC). The BIC indicated that decision field theory provided a superior fit for 13 out of 15 no-feedback participants and 11 out of 14 feedback participants, and consequently fit the data quite well.³ Figure 2 presents the observed choice behavior and predictions of the model. The plots were generated by obtaining the predictions from the best-fitting parameter values at the individual level and using the means of these predictions.

We then compared the obtained best-fitting parameters across the two conditions. Recall that we had predicted that the probability-weighting parameter, λ , would account for the different choice behavior across the two groups. As predicted, λ was higher in the feedback group than in the no-feedback group, indicating that individuals who received feedback underweighted small probabilities relative to their no-feedback counterparts. A Mann-Whitney U test of the three free parameters revealed that only the probability-weighting parameter differed significantly between groups, $z = 2.25$, $p < .05$. However, as Table 1 shows, λ was approximately 1 for the feedback group, and thus reflects not underweighting, but rather objective weighting of small probabilities.

To summarize, the data conformed to our predictions of PT-convergent behavior when feedback was not provided (as indicated by overweighting of small probabilities) but PT-

divergent behavior when feedback was provided (as indicated by relative underweighting of small probabilities). A theoretical analysis supported this finding, showing that feedback specifically altered the probability-weighting parameter. But this analysis also revealed that feedback coupled with description drives subjective probabilities toward their objective values.

DISCUSSION

The results of this study implicate feedback as the crucial arbiter between experiential and descriptive choice. Apparently, feedback overwhelms descriptive information and drives individuals toward the objective probabilities (as indicated by the fact that λ was approximately 1 in the feedback group). Given the design of the task, the difference between the feedback and no-feedback conditions cannot be attributed to sampling bias.

The findings are interesting for a variety of reasons. First, finding that feedback alone drives individuals to the objective probabilities bears relevance to many current studies. For example, in most studies testing behavioral models of choice, feedback is not provided. However, neuroimaging studies concerned with substrates and networks of reward and decision mechanisms typically provide feedback after each trial. Consequently, the observed neural substrates of choice may not correspond with the neural mechanisms actually used in the behavioral studies that inspired the neuroimaging studies.

Second, these results suggest that previous observations of small-probability underweighting in experience-based tasks might possibly be described as objective probability weighting coupled with risk aversion. In several reports of underweighting of unlikely events, a utility parameter was not applied to the outcomes; the failure to do so could produce the appearance of underweighting. However, these other studies (e.g., Barron & Erev, 2003; Hertwig

et al., 2004) also differed from our experiment by not including descriptions in the experiential condition. Thus, feedback coupled with description might be needed for approximately objective weighting.

These results do not implicate feedback as the sole factor determining whether or not small-probability events are overweighted, as several other factors could still be involved; the fact that the median best-fitting probability-weighting parameter across individuals in the no-feedback group was higher than usual attests to this possibility. This mildly atypical finding could have been driven by the small outcomes, parametric modulation of the risky outcomes, or repeated nature of the task. However, this finding does not minimize the fact that behavior differed significantly depending only on whether feedback was or was not received.

Though we reported an analysis based on a particular model (decision field theory), we actually considered multiple models, as our objective was to use a model applicable both when feedback was and was not given (and thus no learning occurred). For the feedback participants, we compared 11 other models, including models with delta and Bayesian learning rules, and also a softmax choice rule (Sutton & Barto, 1998). The softmax choice rule was also fit to the no-feedback participants' data. However, these models produced inferior BIC scores relative to the no-learning implementation of decision field theory; the inferiority of the learning models was likely due to the swiftness of learning. Possibly the combination of feedback with descriptions produced rapid learning (cf. Stout, Rock, Campbell, Busemeyer, & Finn, 2005). The softmax choice rule, though having a BIC score slightly inferior to that obtained by decision field theory across all subjects, still produced the same conclusion: Only the probability-weighting parameter differed significantly between groups.

To conclude, these results indicate that individuals respond differently depending on whether or not they receive feedback after their choices, even when they are provided with complete descriptive information. This suggests that feedback plays a crucial role in differentiating between experiential and descriptive choice behavior. The observed behavioral difference has implications for behavioral choice theories in economics and psychology, as well as for neurophysiological studies aimed at uncovering the neural substrates underlying choice behavior.

Acknowledgments—We thank Sarah Queller and Peter Todd for thoughtful comments on the manuscript and Bethany Ward-Blum and Adolfo Rio Blanco for assistance in collecting data. This research was supported by the National Institute on Drug Abuse (R01 DA014119), the National Institute of Mental Health (R01 MH068346), and the National Science Foundation (BCS-0215738).

REFERENCES

- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. Journal of Behavioral Decision Making, 16, 215–233.
- Benartzi, S., & Thaler, R.H. (1995). Myopic loss aversion and the equity premium puzzle. Quarterly Journal of Economics, 110, 73–92.
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. Econometrica, 22, 23–36. (Original work published 1738)
- Busemeyer, J.R., & Diederich, A. (2002). Survey of decision field theory. Mathematical Social Sciences, 43, 345–370.
- Busemeyer, J.R., Johnson, J.G., & Jessup, R.K. (2006). Preferences constructed from dynamic micro-processing mechanisms. In S. Lichtenstein & P. Slovic (Eds.), The construction of preference (pp. 220–234). New York: Cambridge University Press.
- Busemeyer, J.R., & Stout, J.C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. Psychological Assessment, 14, 253–262.
- Busemeyer, J.R., & Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. Psychological Review, 100, 432–459.
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. Psychological Review, 112, 912–931.
- Ex-teammate tells of Bias drug role. (1987, May 28). New York Times. Retrieved March 3, 2008, from <http://www.nytimes.com>

- Fox, C.R., & Hadar, L. (2006). Decisions from experience= sampling error+ prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, 1, 159–161.
- Hau, R.C., Pleskac, T.J., Kiefer, J., & Hertwig, R. (in press). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*.
- Hertwig, R., Barron, G., Weber, E.U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15, 534–539.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- Nelder, J.A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, 7, 308–313.
- Newell, B.R., & Rakow, T. (in press). The role of experience in decisions from description. *Psychonomic Bulletin & Review*.
- Prelec, D. (1998). The probability weighting function. *Econometrica*, 66, 497–527.
- Rakow, T., Demes, K., & Newell, B. (2007). *Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice*. Manuscript submitted for publication.
- Real, L.A. (1991). Animal choice behavior and the evolution of cognitive architecture. *Science*, 253, 980–986.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6, 461–464.

- Stout, J.C., Rock, S.L., Campbell, M.C., Busemeyer, J.R., & Finn, P.R. (2005). Psychological processes underlying risky decisions in drug abusers. *Psychology of Addictive Behaviors*, 19, 148–157.
- Sutton, R.S., & Barto, A.G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Thaler, R.H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, 112, 647–661.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representations of uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- Weber, E.U., Shafir, S., & Blais, A.R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, 111, 430–445.
- Yechiam, E., Barron, G., & Erev, I. (2005). The role of personal experience in contributing to different patterns of response to rare terrorist attacks. *Journal of Conflict Resolution*, 49, 430–439.
- Yechiam, E., & Busemeyer, J.R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, 19, 1–16.

(Received 1/4/08; Revision accepted 3/31/08)

¹Descriptive choice is distinct from descriptive theories of choice, which describe how people choose and contrast with normative theories, which explain what people should choose.

²At first glance, these results seem to contradict the results of other work in which the gambles were completely described and had a similar payoff structure, and in which groups made a single descriptive choice (to be played 100 times) or 100 choices with feedback after every trial (Yechiam, Barron, & Erev, 2005). In that work, individuals who chose descriptively preferred the safe option, but those who chose experientially preferred the risky option. However, that experiment consisted of 100 trials, whereas Yechiam and Busemeyer's (2006) task consisted of 400 trials, analyzed in blocks of 50. Consequently, it is possible that in the latter study, individuals preferred the safe option on the first trial, then quickly moved toward the risky option, and only after 100 trials began the approach toward expected-value maximization, which favored the safe option.

³Median \underline{R}^2 values were .91 and .84 for the no-feedback and feedback groups, respectively. It is possible that these values could be improved upon because we maximized the likelihood rather than \underline{R}^2 , and fit the individual data rather than the group data. We also tested the nested models within this implementation of decision field theory and found that the three-parameter version produced superior BIC scores.

Fig. 1. Probability of choosing the sure thing in the feedback and no-feedback groups, as a function of block of 10 choices. The upper panel shows choice probabilities for the low-probability pair of gambles, and the lower panel shows choice probabilities for the high-

probability pair of gambles. The plots on the far right indicate the overall means collapsed across blocks; the error bars represent standard errors of the means.

Fig. 2. Probability of choosing the sure thing in the feedback and no-feedback groups, as a function of value of the risky gamble. The graph shows both observed results and aggregate model fits obtained by fitting the model to the observed individual data and averaging the predictions across participants. Results for the low-probability gamble are on the left, and results for the high-probability gamble are on the right. There were five different risky gambles in each condition, and the sure thing was always 3¢. The probability of winning the risky amount was .05 in the low-probability gamble and .80 in the high-probability gamble. The error bars represent the standard errors of the means for the observed choice probabilities.

APPENDIX

Decision Field Theory

In decision field theory, preference for an option accumulates at each time step in accordance with a diffusion process until preference for an option exceeds a preset threshold. In this study, we used an implementation of decision field theory that corresponds with a Wiener process (see Busemeyer & Townsend, 1993, Appendix):

$$\Pr(x_i | X) = \frac{e^{-\frac{(\theta - d_i)^2}{2\sigma^2}}}{\sum_j e^{-\frac{(\theta - d_j)^2}{2\sigma^2}}}, [\text{new equation to be substituted}]$$

where the term on the left denotes the probability of choosing option i out of the set of all options X . The drift rate, d , is the difference between the mean valences of the two options; the diffusion rate, σ , is the square root of the sum of the variance of the two options; and the threshold, θ ,

represents the absorbing boundary to which preference accumulates. The drift rate is computed as follows:

$$d = V_A - V_B.$$

V_i represents the mean valence for option i (shown in the equation as options A and B) and is derived according to

$$V_i = \sum_{j=1}^N w_{ij} \cdot x_{ij}^\alpha,$$

where w_{ij} represents the attention weight given to outcome j when option i is considered, and x_{ij}^α represents the utility of outcome j produced by option i . The parameter α is a utility parameter; a value less than 1 denotes decreasing sensitivity to differences in outcome magnitude as magnitude increases. The attention weights were computed from a formula proposed by Prelec (1998), which takes the following form:

$$w = e^{(-(-\log p)^\lambda)},$$

where p is the objective probability of an outcome, and λ is the probability-weighting parameter. The diffusion rate, σ , was computed as follows:

$$\sigma = \sqrt{\sum w_{ij} \cdot (x_{ij}^\alpha - V_i)^2}$$

Consequently, the parameters d and σ were derived from the stimuli, and the parameters λ and α were free to vary. The third free parameter was the standardized threshold, θ/σ .

Maximum Likelihood Estimation

The parameter values that maximized the likelihood of the data were estimated according to

$$LL_M = C + \sum_i^N \ln(\mathcal{O}_M) \cdot y_i,$$

where \underline{LL} is the log likelihood of the data given the parameter values of model \underline{M} , y_i is the number of times option i was chosen, \hat{y}_M represents the predicted choice proportion for model \underline{M} . In our model, \underline{N} was 10 (i.e., 10 different pairs of gambles, 5 in each condition). \underline{C} is a constant that cancels out and is thus ignored. Maximum likelihood estimation was performed in Matlab using the robust combination of a simplex search (Nelder & Mead, 1965) and multiple random starting points.

Saturated Baseline Model

The baseline model we used assumed that a multinomial process generates choices (see Busemeyer & Stout, 2002). This baseline model perfectly reproduces the marginal choice probabilities for each of the gamble pairs and consequently has a sum of squared error equal to 0. However, the model has one free parameter for each of the gamble pairs and is thus punished by the Bayesian information criterion (BIC) because of the large number of free parameters.

Model Comparison

The BIC was used to compare the performance of the baseline model and the performance of decision field theory. The formula for the BIC was

$$BIC = 2 \cdot (LL_{DFT} - LL_B) + k \cdot \ln N,$$

where \underline{k} represents the difference between the two models in the number of degrees of freedom (i.e., 7), \underline{N} is the number of data points (i.e., 240 trials, because the models were compared at the individual level), \underline{LL}_{DFT} is the log likelihood for decision field theory, and \underline{LL}_B is the log likelihood for the baseline model. In this formulation, positive BIC values indicate better performance for decision field theory than for the saturated model.

Table 1

Median Bayesian Information Criterion (BIC) and Best-Fitting Parameter Values in the No-Feedback and Feedback Groups

Group	BIC	λ	$1/\theta$	α
No-feedback	19.78	0.90	0.04	0.92
Feedback	19.94	0.98	0.13	0.98

Note. All data were fit at the individual level. A positive BIC value indicates that decision field theory outperformed the baseline model. The parameters are as follows: λ is a probability-weighting parameter from Prelec (1998), θ controls the extent to which a participant behaves deterministically, and α is a utility parameter.