Journal of Experimental Psychology:
Learning, Memory, and Cognition
2008, Vol. 34, No. 6, 1446–1465

0278-7393/08/$12.00 DOI: 10.1037/a0013646

The Probabilistic Nature of Preferential Choice

Jörg Rieskamp
Max Planck Institute for Human Development, Berlin, and University of Basel

Previous research has developed a variety of theories explaining when and why people's decisions under risk deviate from the standard economic view of expected utility maximization. These theories are limited in their predictive accuracy in that they do not explain the probabilistic nature of preferential choice, that is, why an individual makes different choices in nearly identical situations, or why the magnitude of these inconsistencies varies in different situations. To illustrate the advantage of probabilistic theories, three probabilistic theories of decision making under risk are compared with their deterministic counterparts. The probabilistic theories are (a) a probabilistic version of a simple choice heuristic, (b) a probabilistic version of cumulative prospect theory, and (c) decision field theory. By testing the theories with the data from three experimental studies, the superiority of the probabilistic models over their deterministic counterparts in predicting people's decisions under risk become evident. When testing the probabilistic theories against each other, decision field theory provides the best account of the observed behavior.

Keywords: expected utility theory, cumulative prospect theory, priority heuristic, decision field theory, probabilistic choice models

Previous research has developed a variety of theories explaining when and why people's decisions under risk deviate from the standard economic view of expected utility maximization. The goal of this article is to show that these theories are limited in their predictive accuracy when they ignore the probabilistic nature of preferential choice.

The standard economic view of expected utility maximization assumes that when choosing between risky options, people always select the one with the maximum expected utility (see, e.g., von Neumann & Morgenstern, 1947; see also Savage, 1954). This idea has been challenged by a substantial body of psychological research. The most influential alternative view put forward so far is that of prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), which asserts that people evaluate the outcomes of risky options as gains and losses resulting from a comparison with a context-dependent reference point. Other alternative theories include rank-dependent expected utility theories (e.g., Quiggin, 1982), regret and disappointment theory (Bell, 1982, 1985; Loomes & Sugden, 1982), weighted utility theory (e.g., Fishburn, 1983), decision field theory (Busemeyer & Townsend, 1993; Roe, Busemeyer, & Townsend, 2001), the transfer-of-attention exchange model (Birnbaum & Chavez, 1997), decision affect theory (Mellers, 2000), and the proportional difference model (González-Vallejo, 2002). For overviews of alternative approaches challenging the expected utility perspective, see Camerer (1989), Machina (1989), Starmer (2000), or Rieskamp, Busemeyer, and Mellers (2006).

These alternative accounts differ with respect to whether they are *deterministic,* always predicting the same choices, or *probabilistic,* predicting different choices with different probabilities. I will illustrate the importance of taking the probabilistic nature of preferential choice into account by comparing probabilistic theories of decision making under risk against their deterministic counterparts. Specifically, I will first extend a recent deterministic model of decision making under risk, the priority heuristic (Brandstätter, Gigerenzer, & Hertwig, 2006), into a generalized probabilistic choice model. Second, I will consider the most prominent model of decisions under risk, namely, cumulative prospect theory (Tversky & Kahneman, 1992), and test it against a probabilistic generalization. Finally, I will consider decision field theory (Busemeyer & Townsend, 1993), which was developed specifically to explain the probabilistic nature of preferences. Through a reanalysis of one previous and two new experimental studies, I will illustrate the advantages of the probabilistic models in comparison with the deterministic models.

The Probabilistic Nature of Preferences

The probabilistic nature of preferential choice is manifested in two phenomena. First, the same individual does not always make the same choices in the same (or nearly the same) situations; that is, people make inconsistent choices. Second, when repeatedly encountering nearly identical choice problems, people show varying magnitudes of inconsistency. In other words, for one choice problem, a person might prefer a particular option 99% of the time, whereas for another choice problem, she might prefer a particular option just 60% of the time.

The probabilistic nature of preferential choice has been investigated in numerous studies. In a seminal study, Mosteller and Nogee (1951) argued that when first offering a bet with a particular

Jörg Rieskamp, Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Berlin, Germany, and Department of Psychology, University of Basel, Basel, Switzerland.

I would like to thank Ido Erev for providing me with the opportunity to analyze the data of Erev et al. (2002). Special thanks are due to Joseph Johnson for helpful discussions and comments that initiated this project. I gratefully acknowledge Bryan Bergert and Gerd Gigerenzer for helpful comments on previous versions of this article.

Correspondence concerning this article should be addressed to Jörg Rieskamp, University of Basel, Department of Psychology, Missionsstrasse 62A, 4055 Basel, Switzerland. E-mail: joerg.rieskamp@unibas.ch

probability of winning and then increasing that probability, “there is not a sudden jump from no acceptances to all acceptances at a particular offer, just as in a hearing experiment there is not a critical loudness below which nothing is heard and above which all loudnesses are heard” (p. 374). Instead, when the probability of winning increases, “the bet is taken occasionally, then more and more often, until, finally, the bet is taken nearly all the time” (p. 374). To test their argument, Mosteller and Nogee conducted an experiment over a period of 10 weeks with three sessions per week. The participants had to decide repeatedly whether to accept bets with a particular probability of winning a specific amount. If they accepted a bet, it was played out in front of them; if they won they received their winnings immediately, and if they lost they paid a fixed amount. The authors did not find a sudden all-or-nothing jump from no acceptance to complete acceptance when the potential winnings increased. Instead, low potential winnings were always rejected, but when the winnings increased, the proportion of acceptance increased as well, and at some point the participants always accepted the bet, which could be best described by a sigmoid function like the one shown in Figure 1. Figure 1 clearly illustrates that for intermediate winnings the bets were neither always rejected nor always accepted.

Hey (2001) provided another demonstration of the probabilistic nature of preferential choice. In this study, participants took part in five experimental sessions. In each session, they made choices for the same set of 100 pairs of gambles. The participants showed numerous inconsistencies. When comparing the first and second sessions, the participants made different decisions for an average of 11% of the pairs of gambles. Even in the fifth session, after having gathered extensive experience, participants changed their choices for 9% of the pairs in comparison with those in the fourth session. No single participant made the same choices every time, with the most consistent of the 53 participants making the same decisions in the last three sessions.

These studies advocate—as has been shown in the area of psychophysics (e.g., Falmagne, 1976)—the development of error theories for describing decisions under risk. In a recent review, Rieskamp et al. (2006) concluded that when confronted twice with a choice between two gambles with similar expected values but varying risks, people’s decisions differ in approximately 25% of the cases. Thus, decision theories need to acknowledge the probabilistic nature of preferential choice. For researchers working in the area of psychophysics, this claim appears to go without saying, but when considering the dominance of deterministic, axiomatic theories of decision making, it requires additional support. This point has been made before: Block and Marschak (1960) stated that absolute consistency of choice is “usually assumed in economic theory, but is not well supported by experience” (p. 90); others have made similar arguments since then (e.g., Busemeyer, 1980; Luce, & Suppes, 1965; Marley, 2002). In the words of Duncan Luce, “progress on this front would be most welcome” (Luce, 2000, p. 281).

A Deterministic Heuristic of Decision Making: The Priority Heuristic

Recently, Brandstätter et al. (2006) proposed a new challenge to the expected utility approach. The authors asserted that people’s choices among risky options can be predicted by using a simple

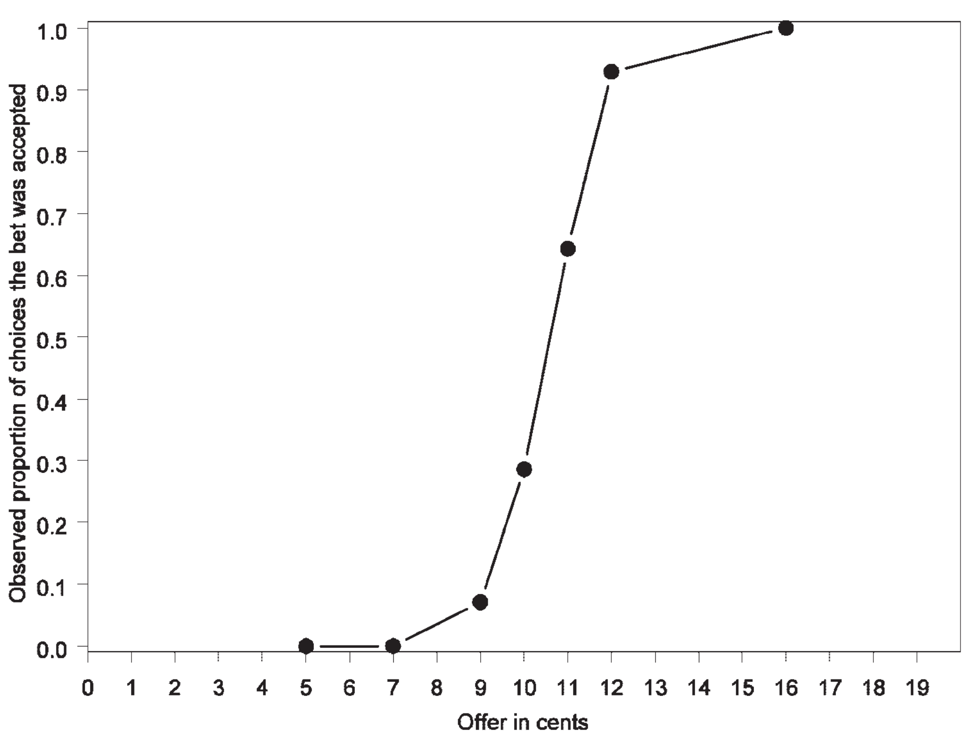

Figure 1. The choice proportions of participant B-I in Mosteller and Nogee (1951) of accepting a bet that led with a probability of .33 to a specific gain represented on the abscissa and with a probability of .67 to a loss of 5 cents. Each bet was offered 14 times, and when the participant rejected the bet, a payoff of zero resulted.

deterministic heuristic that compares the options' aspects sequentially and makes a decision on the basis of one single aspect alone. The authors reported that this *priority heuristic* predicts people's decisions under risk better than alternative theories, including prospect theory.

According to the priority heuristic, when making a choice between gambles, each gamble can be described in terms of several aspects: its minimum outcome, its maximum outcome, and the probability with which the minimum outcome occurs. The gambles' minimum outcomes are compared first. If the difference between the two outcomes exceeds 10% of the maximum payoff that can be obtained from either of the two, the gamble with the larger minimum outcome is selected. For gambles with only positive payoffs, the minimum outcome refers to the minimum gain (i.e., the lowest payoff); for gambles with only negative payoffs, the minimum outcome refers to the minimum loss (i.e., the highest payoff).[1] (Thus, unconventionally, for only negative payoffs the minimum outcome refers to the mathematical maximum.) If the difference does not exceed the threshold, the second aspect, the probability of the minimum outcome, is considered. If the alternatives differ on this aspect by more than .10, the gamble with the lower probability is selected. Otherwise, the last aspect is considered by choosing the gamble with the largest maximum outcome. The heuristic has similarities to other (deterministic) process models, in particular to semi-lexicographic heuristics that assume a sequential decision process (e.g., Leland, 1994; Payne, Bettman, & Johnson, 1988; Rubinstein, 1988; Thorngate, 1980).

To illustrate the priority heuristic's predictions, consider the choices between three gambles: Gamble A leads with a probability of 10% to a payoff of $0 and with a probability of 90% to a payoff of $100. Gamble B leads with a probability of 10% to $40 and with a probability of 90% to $90. Gamble C leads with a probability of 90% to $40 and with a probability of 10% to $90. When comparing Gamble A with Gamble B, the priority heuristic chooses B due to its higher minimum outcome of $40. When comparing Gamble A with Gamble C, the priority heuristic chooses C, again due to its higher minimum outcome. Thus, the priority heuristic reacts very insensitively to Gamble B's and Gamble C's substantially different expected values of $85 and $45, respectively. Likewise it does not choose Gamble A, the gamble with the highest expected value of $90.

A Probabilistic Generalization—The Priority Model

Due to its deterministic character, the priority heuristic cannot explain why people make inconsistent choices or why the magnitude of inconsistency varies. To overcome these shortcomings, I developed a probabilistic generalization, called the *priority model*. Any potential predictive advantage of the priority model can be easily evaluated by incorporating the heuristic as a special case of the model. The model's underlying probabilistic mechanisms assume a stochastic component of decisions under risk that can be traced back to the early work of Thurstone (1927, 1959). Since then, the Thurstonian approach to constructing probabilistic models has been successfully applied in research on judgment and decision making (e.g., Busemeyer & Townsend, 1993; Erev, Wallsten, & Budescu, 1994; González-Vallejo, 2002; Mellers & Biagini, 1994; Tversky, 1972).

The priority model shares the main basic assumptions of the priority heuristic. First, it assumes that when people make choices between options, they do not evaluate each option independently. Instead, it asserts that people compare options with each other by determining differences between the options with respect to several aspects. Second, the model assumes that people do not integrate information but instead consider the different aspects sequentially. Third, the priority model assumes a fixed order in which aspects are considered, but in contrast to the priority heuristic, this order can vary across individuals.

To explain the probabilistic nature of preferential choice, the priority model assumes that when an individual compares one aspect of each of two options, the outcome is a subjective difference, best represented by a random variable. Specifically, when the model starts to compare two gambles by their minimum outcomes, the resulting subjective difference is normally distributed with a mean of the actual difference and the standard deviation sigma, which is a free parameter of the model. The difference is then compared with the threshold $\delta_{out} = \delta^*_{out} \times \max(\max_A, \max_B)$, where $0 < \delta^*_{out} < 1$ is a second free outcome threshold parameter and $\max(\max_A, \max_B)$ is the greater of the two maximum outcomes being compared.[2] The threshold varies across individuals but is fixed for each individual. When the priority model compares the probabilities with which the minimum outcomes occur, the result is again a normally distributed difference with a mean of the difference and the standard deviation sigma. This difference is then compared with the probability threshold δ_{pr}, which is the third free parameter varying between 0 and 1 (where the subscript refers to probability). Finally, the model assumes that the possible six orders in which the three aspects are considered can vary across individuals, parameterized with two free parameters. The mathematical details of the priority model are specified in Appendix A.

The priority model has the property of approximating the priority heuristic's prediction in the case of a small standard deviation sigma (e.g., $\sigma = 0.001$), an outcome threshold of $\delta^*_{out} = .1$, a probability threshold of $\delta_{pr} = .1$, and implementing the priority heuristic's order of aspects. Henceforth these values will be called the *default parameter values* of the priority model. Moreover, the priority model has the property of being able to represent various deterministic heuristics, which become probabilistic when the standard deviation parameter sigma increases. In sum, whereas the priority heuristic represents one specific process that people could follow to make their choices, the priority model is able to represent a variety of decision processes.

In conceptual terms, the priority model assumes that several comparison processes lead to subjective differences that can vary from context to context as well as over time, explaining the probabilistic nature of choices. On the basis of the subjective

[1] Brandstätter et al. (2006) suggested that the priority heuristic could in principle also be applied to gambles that involve both gains and losses (i.e., mixed gambles). However, they concede that due to a lack of experimental evidence, they cannot state conclusively whether the priority heuristic can be generalized to mixed gambles. I will fill in this gap with the conducted experimental studies by testing the priority heuristic on mixed gambles.

[2] For mathematical simplicity, the priority model does not assume that δ_{out} is rounded to a prominent number (cf. Albers, 2001) as the priority heuristic does. However, the results and the conclusions drawn in this article do not depend on this difference. All analyses presented in this article were repeated by including rounding to prominent numbers, which gave very similar results.

difference, the decision maker applies a deterministic decision rule, for instance, the rule of choosing the option with the largest minimum outcome. In this respect, the model is similar to random utility models, which also assume that the utility of options varies over time and context, but that at each time point the option with the largest utility is always selected. In contrast, the priority model differs from fixed utility theories, which assume that each option can be assigned a fixed value but which apply a probabilistic choice rule to make probabilistic predictions (see also Rieskamp et al., 2006).[3]

The priority model will make predictions similar to the priority heuristic when it uses the same order of aspects and the same thresholds as the priority heuristic. However, the model's predictions can vary when assuming that the subjective differences vary substantially. For instance, when comparing the three gambles described above—A = {10%, \$0; 90%, \$100}, B = {10%, \$40; 90%, \$90}, and C = {90%, \$40; 10%, \$90}—the difference of the minimum outcomes might not decisively determine the predicted choice probability; instead the probabilities with which the minimum payoffs occur will be of importance. Therefore the priority model can predict a lower choice probability for choosing Gamble A when comparing A and B than when comparing A and C, because for Gamble C the high probability with which the minimum outcome occurs is at a large disadvantage in comparison with Gamble A.

When testing the priority heuristic with a maximum likelihood statistic, a disproportionately bad fit will be given to the heuristic for an incorrect prediction. Therefore I additionally define a *priority heuristic* implementing a naïve error theory, which assumes that a person applying the heuristic deviates from the prediction with a constant probability of α, representing an application error, with $0 < \alpha < .5$ (for a similar approach, see Rieskamp, 2006; Rieskamp & Otto, 2006). With an application error of $\alpha = .10$, for instance, the heuristic predicts its choice with a probability of .90. In the following, the optimal application error will be estimated for each individual. The priority heuristic with this naïve error theory can describe constant inconsistencies. However, in contrast to the priority model, it is not possible to explain why larger inconsistencies occur for some choice situations compared with others. Table 1 provides an overview of the three versions of the priority heuristic and its free parameters.

Table 1
Summary of the Models' Estimated Parameters

Model	Parameter
Priority heuristic[a] (deterministic) with fixed parameters	No free parameter
Priority heuristic (probabilistic) with a naïve error theory	1. Application error parameter α
Priority model (probabilistic) with five free parameters	1. Standard deviation of error component σ 2. Outcome threshold δ_{out} 3. Probability threshold δ_{pr} 4.–5. Importance weights for minimum outcomes w_{min} and for probability of minimum outcomes w_{pr}
Cumulative prospect theory[b] (deterministic) with five free parameters	1. Subjective value function's exponent for gains α 2. Subjective value function's exponent for losses β 3. Weight for losses λ 4. Probability function's exponent for gains γ 5. Probability function's exponent for losses δ
Cumulative prospect theory (probabilistic) with an additional sixth parameter	6. Sensitivity parameter ϕ
Decision field theory (deterministic) with a high decision threshold	No free parameter
Decision field theory[c] (probabilistic) with one free parameter	1. Decision threshold θ_{DFT}

[a] According to Brandstätter et al. (2006). [b] According to Tversky and Kahneman (1992). [c] Representing a simplified version of decision field theory of Busemeyer and Townsend (1993).

Cumulative Prospect Theory

Among all challengers of the expected utility approach, prospect theory (see Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) represents the most prominent one, making it an obvious choice as a competing model. According to prospect theory, when choosing between two gambles, an overall subjective value V can be determined for each gamble. The decision maker will prefer the gamble with the larger value. To determine a gamble's overall subjective value, each payoff is represented by a subjective value multiplied by a decision weight determined by the probabilities with which the payoffs occur (for details, see Appendix B). Thus, cumulative prospect theory assumes that each option is evaluated independently of other options. Cumulative prospect theory can be assigned to the class of rank-dependent utility models, which assume that the objective probabilities with which outcomes occur are represented by subjective probabilities that depend on the rank of the outcome relative to the other outcomes (see Birnbaum & Chavez, 1997; Luce & Fishburn, 1991, 1995; Quiggin, 1993).

Tversky and Kahneman (1992, p. 300) provided an axiomatic definition of cumulative prospect theory that yields a deterministic model. Following my main argument of probabilistic preferences, I also examine a probabilistic version of prospect theory. Prospect theory assigns a fixed value to each choice option, so that it seems reasonable to construct a probabilistic version in line with fixed utility theories (see Luce & Suppes, 1965). These theories assume

[3] To what extent random utility models and fixed utility models make different predictions depends on the specific assumption made. According to random utility models, the utility of an option varies due to an error component, and depending on the assumed distribution of the error component different predictions follow. In contrast, the predictions of a fixed utility model depend on the choice rule. For instance, Luce's fixed utility model (Luce, 1959) can lead to very similar predictions as a random utility theory (see Luce, 1977). McFadden (1974) showed that Luce's choice model predicts the same choice probabilities as a random utility model if and only if the error component of the random utility model follows an independent identical extreme value distribution. In contrast, when assuming, for instance, correlated error components, random utility models can explain violations of strong stochastic transitivity that fixed utility theories cannot explain.

that each option has a deterministic subjective value for the decision maker but that the final choice process is probabilistic. I will apply the standard exponential choice rule to specify this choice process (consistent with Birnbaum & Chavez, 1997; Lopes & Oden, 1999; see also Luce, 1959).

Accordingly, the probability of choosing Gamble A over Gamble B can be determined by

$$p(\mathrm{A,B}) = \frac{e^{\varphi \cdot V(\mathrm{A})}}{e^{\varphi \cdot V(\mathrm{A})} + e^{\varphi \cdot V(\mathrm{B})}}, \qquad (1)$$

where $\varphi > 0$ is a sensitivity parameter, specifying how sensitively the model reacts to the differences among the gambles' subjective values. A very large sensitivity parameter value makes the probabilistic version essentially identical to the deterministic version. This probabilistic version does not explain why people's preferences are probabilistic; instead the choice probabilities are simply assumed to be a monotonic function of the differences of the gambles' subjective values. Overall, the probabilistic prospect theory has six free parameters specifying a nonlinear utility function and a probability weighting function (see also Table 1).

To illustrate prospect theory's predictions consider again the choices between Gamble A = {10%, $0; 90%, $100}; Gamble B = {10%, $40; 90%, $90}; and Gamble C = {90%, $40; 10%, $90}. By using the parameter estimates of Tversky and Kahneman (1992), prospect theory predicts that Gamble B is preferred to Gamble A, but A is preferred to C. Although Gamble A has a higher expected value than Gamble B, Gamble B is preferred because the low probability of 10% will be overweighed (i.e., transformed to a value of .19), and the high probability of 90% will be underweighed (i.e., transformed to a value of .71). In contrast, Gamble A is preferred to Gamble C because here the transformed probabilities do not compensate Gamble A's larger expected value advantage.

Decision Field Theory

Decision field theory (DFT, Busemeyer & Townsend, 1993) is another prominent theory of decision making under risk, which was specifically developed to account for the probabilistic nature of preferences, providing a promising third theory for comparison in this article. DFT aims to provide a description of the cognitive processes underlying decisions. DFT assumes that people do not evaluate options independently of each other; rather, the theory assumes that the options' aspects are compared with each other (similar to the priority heuristic). In contrast to the priority heuristic, however, the result of each comparison is summed up in an overall difference score. Once this score passes a decision threshold, a choice is made. According to DFT, the attention an individual devotes to the options' aspects fluctuates over time. Depending on which aspects receive the most attention during a decision process, the probability of choosing a specific option varies. Additionally, DFT assumes that the decision threshold varies across individuals and decision situations. A high threshold leads to a long decision time where all aspects are numerously compared, so that the option with the larger expected value will be selected with a high probability. In contrast, a low decision threshold leads to a quick decision after a few comparisons have been performed.

In the following, I specify a simplified version of DFT.[4] When comparing two gambles, a difference score d_{DFT} results from computing and then summing the options' outcome differences. A decision is made whenever this difference score passes the decision threshold θ_{DFT}. Because the difference score is the result of accumulated differences, the probability of passing the threshold is a function of the standard deviation of the differences (i.e., σ_{DFT}). Busemeyer and Townsend (1993) have shown that the probability of choosing Gamble A over Gamble B is defined as the probability that the accumulated difference passes the decision threshold favoring Gamble A, which can mathematically be defined as

$$p(\mathrm{A,B}) = G\left(2 \times \frac{d_{\mathrm{DFT}}}{\sigma_{\mathrm{DFT}}}\, \theta_{\mathrm{DFT}}\right), \qquad (2)$$

where G is the standard cumulative logistic distribution function. Appendix C explains how the difference score d_{DFT} and the standard deviation of the differences σ_{DFT} can be computed. The decision threshold θ_{DFT} is the only free parameter of the model (see also Table 1). The threshold θ_{DFT} corresponds to $\theta^*_{\mathrm{DFT}}/\sigma_{\mathrm{DFT}}$ of the original model of Busemeyer and Townsend (1993, p. 440), so it is assumed that the decision maker always chooses a threshold θ^*_{DFT} that is proportional to σ_{DFT}.

The choice probability is a decreasing function of the variance of the difference score, so that the lower the variance the easier it is to discriminate between the gambles and the larger the resulting choice probability. The choice probability is also a function of the decision threshold; increasing the threshold increases the probability of the most likely option. The simplified DFT has the decision threshold as its only free parameter. With this simplification, DFT will always predict that the gamble with the larger expected value will be chosen with a probability larger than 50%. With an extremely high decision threshold, DFT's predictions become deterministic, so that DFT will predict that the gamble with the larger expected value will always be chosen. In the following section, this deterministic version of DFT will be tested against the probabilistic version.

To illustrate DFT's prediction, consider again the choices between Gamble A = {10%, $0; 90%, $100}; Gamble B = {10%, $40; 90%, $90}; and Gamble C = {90%, $40; 10%, $90}. When using a decision threshold of $\theta_{\mathrm{DFT}} = 1$, DFT predicts that Gamble A will be preferred to Gamble B with a choice probability of .57, whereas Gamble A will be preferred to Gamble C with a choice probability of .94. Thus, contrary to the priority model and prospect theory, DFT always predicts that the gamble with the larger expected value will be chosen with a probability above 50%, and DFT can react very sensitively to the difference between the gambles' expected values.

Study 1: Reanalyzing Previous Findings

The different theories were first tested with the data from a previous study by Erev, Roth, Slonim, and Barron (2002). In this

[4] The full version of DFT also assumes a forgetting process, so that during the decision process the current comparison of the options receives more weight compared with previous comparisons. The full model also entails additional parameters to account for approach–avoidance conflicts and to allow continuous-time course predictions (cf., Busemeyer & Townsend, 1993). In its later extensions, DFT can also account for choices among more than two alternatives (Roe et al., 2001).

study, the participants made 200 choices between pairs of gambles. Each gamble offered a particular probability of a positive payoff ranging from 1 to 100, and the remaining probability of a zero payoff. All gambles were created randomly by selecting a probability and a payoff with equal probability from the corresponding interval. The study by Erev et al. (2002) appears highly appropriate for testing the three theories for various reasons. First, the study was conducted independently of the theories examined in the present article. Second, Erev et al. (2002) selected gambles randomly, so that no a priori advantage for one theory over another should be expected. Finally, Brandstätter et al. (2006) illustrated the relatively good performance of the priority heuristic by using 100 randomly selected pairs of gambles from Erev et al. (2002). Thus, by using these 100 pairs of gambles again, the evaluation of the priority model, cumulative prospect theory, and DFT can be easily connected to the conclusions drawn by Brandstätter et al.

Method and Results

To evaluate the models, I applied a maximum likelihood approach.[5] Accordingly, a model's goodness of fit was defined by the G^2 measurement (e.g., Sokal & Rohlf, 1994):

$$G^2_{\text{Model}} = -2\sum\nolimits_{i=1}^{N} \ln[f_i(y|\theta)], \tag{3}$$

for which i refers to the choice situation with N as the total number of choices, and where $f_i(y|\theta)$ refers to the probability with which the model given its parameters θ predicts an individual's choice y; that is, if Gamble A was chosen, then $f_i(y|\theta) = p_i(\text{A,B})$ and if Gamble B was chosen then $f_i(y|\theta) = 1 - p_i(\text{A,B})$. How can G^2 be interpreted? A very accurate model predicts choices with a high probability, leading to low values for G^2 (with a minimum value of zero), whereas large values reveal a model's poor fit. The precise value of G^2 depends on the number of choices that are evaluated. It is helpful to note that when predicting a choice between two options with an equal probability of .50, then G^2 has a value of $-2 \times \ln(.50) = 1.39$, so that for Study 1 with 100 choices $G^2 = 139$ results. A model with a G^2 value above 139 performs more poorly than does an equal chance prediction. With a likelihood method, any model that incorrectly predicts even a single choice with a probability of 0 would be infinitely inaccurate. Therefore for all models considered, the choice probabilities in all three studies were always truncated to a minimum and maximum of .01 and .99, respectively. Thus, for one single choice the worst G^2 value is $-2 \times \ln(.01) = 9.21$, yielding $G^2 = 921$ for 100 choices. In Studies 2 and 3 with 180 choices, predicting choices with an equal probability of .50 leads to $G^2 = 249.5$, and the worst possible fit is $G^2 = 1{,}658$. For all reported fitting processes, a grid search (with varying densities depending on the parameters) was used to find the best parameter values to minimize G^2, and if an exhaustive search was computationally intractable, the best-fitting grid values were used as a starting point for subsequent optimization using the simplex method (Nelder & Mead, 1965) as implemented in *MATLAB*.

Priority model. The priority heuristic was able to predict 85% of the choices correctly and had a fit of $G^2 = 206$. For the priority heuristic with a naïve error theory, an application error of $\alpha = .24$ was estimated, leading to a substantially improved fit of $G^2 = 109$ (for details, see Table 2). For the priority model, the optimal parameters were first estimated sequentially by fitting one parameter after the other and leaving the other parameters at their default values. Erev et al.'s (2002) pairs of gambles always had a payoff of zero as the minimum outcome, so that their difference was always zero. Therefore a fixed error parameter of $\sigma_{\min} = 0.001$ was employed for the priority model when comparing the minimum outcomes, and the error parameter σ was effective only for the other comparisons. All four components of the priority model were able to produce a substantially improved fit. When fitting all four parameters of the full priority model, the following parameters were obtained: $\sigma = 0.2$, $\delta_{\text{out}} = .71$, $\delta_{\text{pr}} = .10$; with the order of aspects: minimum outcome, maximum outcome, and minimum probability. The full priority model had a significantly better fit of $G^2 = 93$ compared with that of the priority heuristic, $\chi^2(5) = 113$, $p < .01$, according to a likelihood ratio test, which was applied in the present article for all model comparison tests when comparing nested models (cf., Wong, 1994).

The priority model's assumption of a stochastic component at each comparison accounted for the observed behavior better than did the priority heuristic, assuming a naïve constant application error. To illustrate this point, consider those gambles out of the 100 pairs of gambles that the priority heuristic predicts will be chosen. Taken literally, the heuristic predicts that these gambles should be chosen in 100% of all cases, as opposed to the average choice proportion of 76% ($SD = 24$) actually observed. Figure 2 shows a scatter plot of the full priority model's average choice probability and the actual choice proportions, which were strongly correlated ($r = .83$, $p < .01$). Figure 2 illustrates that the choice proportions for the different gambles varied substantially, speaking against a constant, unsystematic application error.

Cumulative prospect theory. Does the priority model also describe the data better than cumulative prospect theory? First, Tversky and Kahneman's (1992) estimated parameter values relevant for positive payoffs (i.e., $\alpha = .88$, $\gamma = .61$) and a sensitivity parameter of $\phi = 100$ were used to determine prospect theory's predictions. With these estimates, prospect theory achieved a fit of $G^2 = 310$. To examine whether the probabilistic version improved the fit, I estimated the optimal sensitivity parameter, leading to a sensitivity parameter of $\phi = 0.13$, which led to a substantially improved fit of $G^2 = 112$, $\chi^2(1) = 198$, $p < .01$. Finally, I estimated all relevant parameters. To avoid over-fitting and to retain prospect theory's main assumption, I restricted the parameters of cumulative prospect theory for all studies in this article to a reasonable range: $0 < \alpha \leq 1$; $0 < \beta \leq 1$; $1 \leq \lambda \leq 10$; $.40 \leq \gamma \leq 1$; $.40 \leq \delta \leq 1$; and $0 < \phi \leq 10$. For instance, λ needs to be larger than 1 to retain prospect theory's main assumption of loss

[5] Thus, unlike Brandstätter et al. (2006), I do not use the majority responses as a goodness-of-fit criterion because this can lead to mistaken conclusions. For example, consider the case of 100 individuals choosing between two gambles, A and B. Two theories X and Y are tested against each other for predicting the choices. Theory X predicts the choice of A with a probability of 1.0, while Theory Y predicts the choice of A with a probability of .49 and the choice of B with a probability of .51. Suppose 51 of the 100 individuals chose Gamble A and 49 chose Gamble B. In this case, Theory Y provides a more accurate description of the results than does Theory X. However, when using the majority response as a goodness-of-fit criterion, the more accurate theory, Y, had a fit of 0%, and the inaccurate theory, X, had a fit of 100%, leading to incorrect conclusions.

Table 2
Results of Estimating Each Parameter of the Priority Model Separately

Model	Parameter estimation: Study 1: Erev et al. (2002)	Parameter estimation: Study 2
Priority heuristic with fixed parameters	$G^2 = 206$	*Mdn* $G^2 = 697$
Priority heuristic with a naïve error theory	$\alpha = .24$ $G^2 = 109$ $\chi^2(1) = 97$ $p < .01$	*Mdn* $\alpha = .42$ *Mdn* $G^2 = 245$ $\chi^2(30) = 13{,}398$ $p < .01$
Priority model with only sigma estimated	$\sigma = 0.1$ $G^2 = 102$ $\chi^2(1) = 104$ $p < .01$	*Mdn* $\sigma = 1.7$ *Mdn* $G^2 = 247$ $\chi^2(30) = 13{,}471$ $p < .01$
Priority model with only outcome threshold estimated	$\delta_{out} = .1$ $G^2 = 206$	*Mdn* $\delta_{out} = .43$ *Mdn* $G^2 = 585$ $\chi^2(30) = 2{,}852$ $p < .01$
Priority model with only probability threshold estimated	$\delta_{pr} = .1$ $G^2 = 206$	*Mdn* $\delta_{pr} = .25$ *Mdn* $G^2 = 661$ $\chi^2(30) = 1{,}575$ $p < .01$
Priority model with only order of aspects estimated[a]	PrM, Max, Min $G^2 = 200$ $\chi^2(1) = 6.7$ $p < .01$	Max, Min, PrM *Mdn* $G^2 = 550$ $\chi^2(60) = 3{,}684$ $p < .01$

Note. Min = minimum outcome; PrM = probability of minimum outcomes; Max = maximum outcome.
[a] The different orders were implemented with two free parameters as described in Appendix A.

aversion. With these restrictions, the best parameter values were $\alpha = .31$, $\gamma = .76$, and a sensitivity parameter of $\theta = 2.20$, providing the best fit of $G^2 = 86$, $\chi^2(2) = 26$, $p < .01$.

How did cumulative prospect theory perform against the priority model? To compare the models, I determined the Akaike Information Criterion (AIC = $G^2 + 2 \times$ number of free parameters) for each model, which is a model selection criterion that takes a model's complexity into account (see Burnham & Anderson, 2002). With an AIC of 92, cumulative prospect theory outperformed the full priority model with an AIC of 103, yielding an AIC difference of ΔAIC = 11. Burnham and Anderson (2002) suggested that a ΔAIC greater than 10 provides strong empirical support that the model with the better fit provides a substantially better explanation of the data than does the competing model. Likewise, when considering that the participants choose the gamble most likely predicted by prospect theory in 98% of all cases and choose the gamble most likely predicted by the full priority model in 88% of all cases, prospect theory does perform better than the full priority model.

Decision field theory. The probabilistic prediction of DFT depends on the decision threshold and the variance of the differences when comparing the options. When using an extremely high decision threshold, the deterministic DFT reached a fit of $G^2 = 213$. To test this deterministic version against DFT, the decision threshold of DFT was estimated, resulting in a threshold of $\theta_{DFT} = .96$. With this estimate, DFT reached a fit of $G^2 = 109$, which is substantially better than its deterministic counterpart, $\chi^2(1) = 104$, $p < .01$. However, when comparing DFT with the priority model or prospect theory, it did not perform as well as either, with an AIC of 111.

Discussion

The reanalysis of the Erev et al. data (2002) illustrates that the probabilistic versions of the three theories describe the data better than do their deterministic counterparts. The estimated parameter values for the full priority model differ substantially from the parameter values implicitly assumed by the priority heuristic. These results suggest that the decision making process is better captured by the priority model than by the priority heuristic, illustrating the superiority of the priority model relative to Brandstätter et al.'s (2006) results. However, the different estimated orders of aspects obtained for the priority model should be interpreted carefully, because Erev et al. (2002) always used the same minimum outcomes with a payoff of zero for both gambles. A further limitation to interpreting the results arises from the use of aggregated data. Figure 2 shows aggregated choice proportions that indicate the probabilistic nature of people's preferences. However, individuals could use different deterministic strategies, which would also lead to varying choice proportions (for different conclusions based on aggregate rather than individual data, see also Estes & Maddox, 2005; Haider & Frensch, 2002).

When comparing the priority model with cumulative prospect theory and DFT, prospect theory performed best. However, this

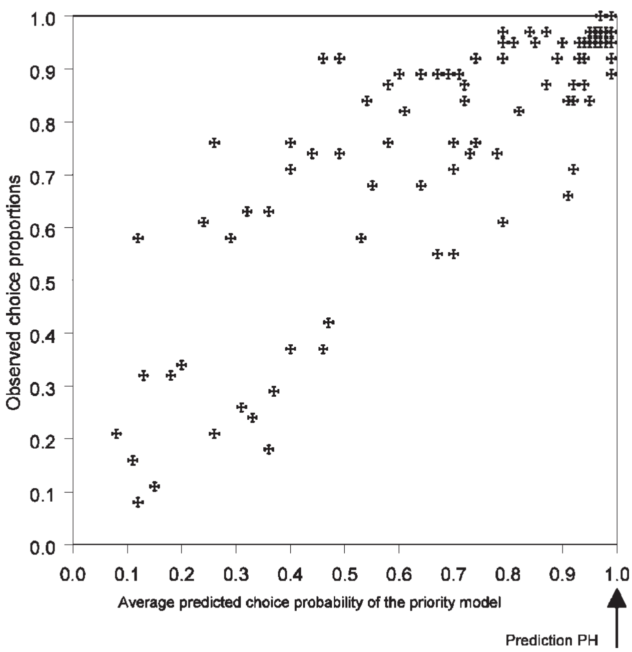

Figure 2. Scatter plot of the average choice probability determined for the full priority model for those gambles the priority heuristic (PH) chooses and the proportion of participants actually choosing the gamble. The choice proportions are shown for all 100 pairs of gambles of Erev et al. (2002).

conclusion is drawn from a comparison of the models' AIC values. Unfortunately for non-nested models such as prospect theory, the priority model, and DFT, it has been argued that AIC is a less appropriate model selection criterion because it neglects the parameters' functional relationships (Pitt, Myung, & Zhang, 2002, but see also Burnham & Anderson, 2002). Moreover, as described above in Erev et al.'s (2002) study, only specific types of gambles were considered where only one outcome led to a gain. Therefore the main purpose of Study 1 was to demonstrate the advantage of probabilistic models in comparison to their deterministic counterparts. To address the limitations of the reanalysis and specifically to test the non-nested models against each other, I conducted two new experimental studies.

Study 2

For Study 2, a new experiment was conducted with a broad set of gambles, covering losses, gains, and mixed outcomes. The set of gambles was created randomly to avoid using a specific test environment. The set was restricted to pairs of gambles with similar expected values and with no stochastically dominated gambles (because people seldom choose dominated gambles, however, see Birnbaum, 2005). The main purpose of Study 2 was to estimate the parameters of the different theories by examining preferential choices across a wide range of situations. With these estimations, the models' predictions can be derived and tested. This crucial generalization test of the models (without fitting the parameters again) was performed in Study 3.

Method

Materials. First, 10,000 pairs of gambles were produced for each of three payoff domains. Each gamble had two outcomes: both positive for the gain domain, both negative for the loss domain, and one positive and the other negative for the mixed domain. The possible outcomes were drawn with equal probability from integer numbers ranging from 0 to 100 for gains, from –100 to 0 for losses, and one payoff from each of the two intervals for the mixed domain. The probability of obtaining one of the two outcomes was drawn with equal probability from the interval 0 to 1 (rounded to two decimals). A payoff occurring with a probability of zero was replaced with a payoff of zero. Those 43% of pairs where one gamble stochastically dominated the other gamble were deleted. The absolute difference of the gambles' expected values was determined and divided by the minimum of the two absolute expected values. All gambles for which this ratio was larger than 1 were eliminated, so that 32% (i.e., 9,598) of the original 30,000 pairs of gambles remained, from which 60 pairs were selected randomly for each domain.

Participants. Thirty people (18 women and 12 men) with an average age of 25 years participated in the experiment. The computerized task lasted approximately 1 hr. Participants were students from various departments at the Free University of Berlin. They received a payment of €15 for their participation (the equivalent of $18 at the time of the study). In addition, after completing the task, one pair of the 180 gambles was selected randomly, and the participant's chosen gamble was played out in front of him or her; 5% of the gamble's payoff was added or subtracted from the €15 payment, resulting in possible payments of between €10 and €20.

Procedure. The instructions explained that participants were to choose repeatedly between pairs of gambles. The two gambles would offer two possible payoffs with specific probabilities, and each gamble would be represented by a pie chart. For illustration, two pie charts were shown to the participants. Participants were told that they would first make three decisions to familiarize themselves with the task and then make another 180 decisions. They were also told that there would be no correct or incorrect decisions but that each decision would depend on the participants' subjective preferences. Finally, they were informed that after they made all 180 decisions, one pair of gambles would be randomly selected, and the chosen gamble would be played in front of them. Five percent of the gain or loss would be added to or subtracted from their payment of €15 for their participation. After reading the instructions and making their three initial choices, participants decided between the 180 pairs of gambles presented in a completely random order. Each gamble was presented separately, on either the left or the right side of the computer screen as shown in Figure 3. The gambles, labeled *Gamble A* and *Gamble B,* were represented by a pie chart and participants made their choices by selecting a button beneath them. The task number was presented at the bottom of the screen.

Results

First, I examined whether participants' choices were systematic. Next, I analyzed whether the probabilistic models described the data better than did their deterministic counterparts, and finally I compared the probabilistic models against each other.

Choice proportions. Considering that the participants made a large number of choices, that the pairs of gambles were presented in completely random order, and that the incentives were not very high, one might wonder whether the participants' behavior was unsystem-

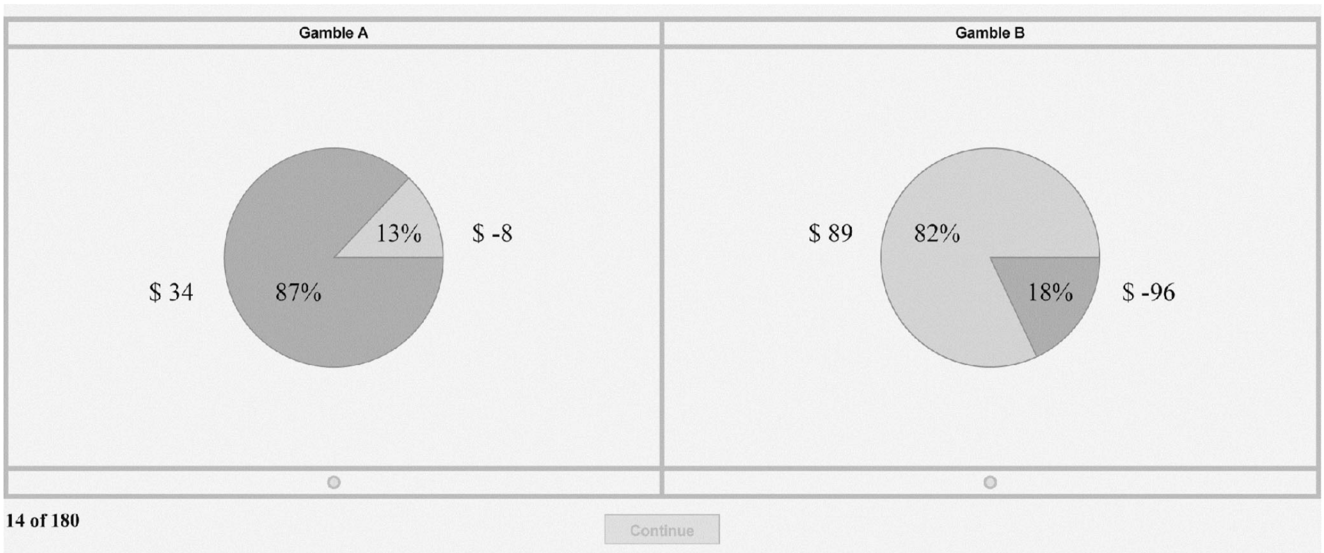

Figure 3. A screenshot of the experimental software used in Studies 2 and 3. Participants made their decisions by selecting one of the buttons beneath the gambles.

atic. I first determined the correlation between the difference of the gambles' expected values and the choice proportions. Participants chose the gamble with the largest expected value 70% of the time. Figure 4 shows a scatter plot illustrating that the larger the expected value of one gamble compared with that of the other, the larger the proportion of participants choosing that gamble: Here, the correlation was $r = .80$ ($p < .01$) as opposed to the zero correlation expected for unsystematic behavior. Nevertheless, Figure 4 also shows the large variance of the choice proportions. If the participants had always chosen the gamble with the higher expected value, then for the left side of the abscissa, for which Gamble A had a lower expected value than Gamble B, the proportion of participants choosing Gamble A should have been 0%, and for the right side of the abscissa, for which Gamble A had a higher expected value than Gamble B, the proportion of choosing Gamble A should have been 100%, contrary to what was observed.

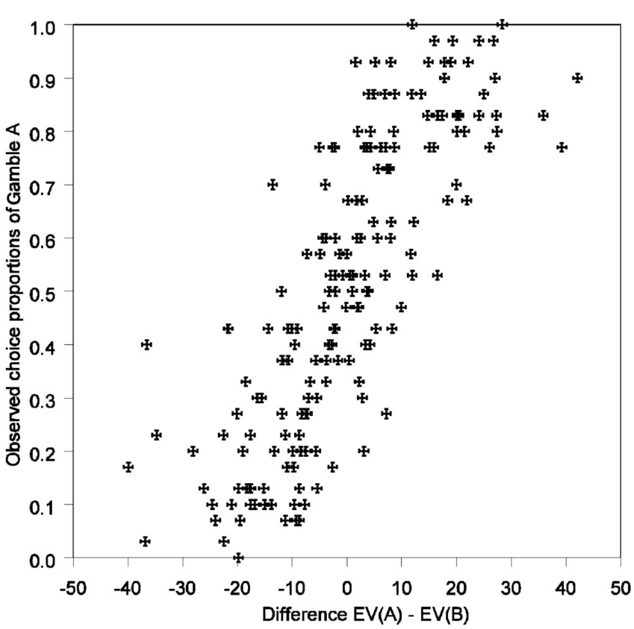

Figure 4. Scatter plot of the difference of the expected value of Gamble A (EV[A]) and of Gamble B (EV[B]), and of the proportion of participants choosing Gamble A from the pair of gambles in Study 2.

Priority heuristic and priority model. First, the priority heuristic's fit was determined by implementing the heuristic with the default parameter values within the priority model. The priority heuristic was able to predict 59% of all choices and obtained a median fit of *Mdn* $G^2 = 697$. The priority heuristic was the worst at predicting the choices for the mixed gambles, with *Mdn* $G^2 = 249$, as compared with *Mdn* $G^2 = 222$ and *Mdn* $G^2 = 231$ for the gain and loss domains, respectively. Although the heuristic's fit differs across the three domains, it still appears appropriate to apply the heuristic to mixed gambles as well, as suggested by Brandstätter et al. (2006). For the priority heuristic with a naïve error theory, an application error of *Mdn* $\alpha = .42$ was estimated, leading to a substantially improved fit of *Mdn* $G^2 = 245$ (for details, see Table 2).

The priority model's free parameters were estimated sequentially and separately for each individual, fitting one parameter after the other and leaving the other parameters at their default values. Table 2 shows that each component of the priority model increased the fit of the model significantly in comparison with that of the priority heuristic. When comparing the four components, the largest improvement in fit was obtained when fitting the standard deviation of the error component. In contrast to the results reported above for the study by Erev et al. (2002), the priority model that only incorporated an error component did not lead to a higher fit in comparison with the priority heuristic with a naïve error theory.

To test whether the combination of parameters could increase the priority model's fit even further, a stepwise procedure for testing the priority model was applied. At each step, it was tested whether an additional parameter could improve the priority model's fit significantly. It turned out that each of the four components was able to improve the fit significantly, so that ultimately, the full priority model with all five parameters achieved the best fit with *Mdn* G^2 = 202.[6] To test whether the full priority model was preferable to the priority heuristic with a naïve error theory, their AICs were determined. For 87% of the participants, the full priority model had a better AIC than did the priority heuristic with a naïve error theory: The average AIC difference was ΔAIC = 30, again providing strong empirical evidence for the priority model. Thus, of all the model variations, the full priority model provided the best account of the observed behavior.[7] The participants chose the most likely gamble predicted by the full priority model in 70% (*SD* = 8%) of all cases. The estimated parameters for the full priority model were as follows: a median standard deviation for the error component of $\sigma = 0.25$, a median probability threshold of $\delta = .49$, and a median outcome threshold of $\delta = .38$. The model's best order of aspects was minimum probability, minimum outcome, and maximum outcome.

How does the full priority model differ from the priority heuristic? Figure 5 shows the frequency distributions of the estimated parameters separately for the priority model's four components, and it illustrates that the distribution differs substantially from the priority heuristic's assumptions. In the four panels of Figure 5, the priority heuristic's assumed parameters are indicated by an arrow. If these assumptions were the best one could make, then the corresponding parameter values should have been frequently estimated for the full priority model. This was not the case. First, the standard deviation of the error component exceeds the very low value implicitly assumed by the priority heuristic: For 90% of the participants, a standard deviation of between .20 and .40 was estimated. Second, the priority model's probability and outcome thresholds were, for most participants, larger than those assumed by the priority heuristic. For the probability threshold, for 90% of the participants, a value of between .29 and .70 was estimated. For the outcome threshold, for 90% of the participants, a value of between .10 and .60 was estimated. Finally, the best order of aspects to describe participants' decisions with the priority model (implemented with two free parameters) differed from the order assumed by the priority heuristic. For most of the participants (60%), the best order of aspects was to consider the minimum probability first, the minimum outcome second, and the maximum outcome last.

Cumulative prospect theory. To determine prospect theory's predictions, I first used the parameter values estimated by Tversky and Kahneman (1992) with a large sensitivity parameter of θ = 100. With these estimates, cumulative prospect theory achieved a fit of *Mdn* G^2 = 437. The probabilistic version of prospect theory with an estimated sensitivity parameter of $\phi = 0.08$ led to a substantially improved fit of *Mdn* G^2 = 207, $\chi^2(1) = 8{,}358$, $p < .01$. Second, all six parameters of prospect theory were estimated by using the constraints described above. The optimal estimated parameters were $\alpha = .91$, $\beta = .91$, $\lambda = 1$, $\gamma = .69$, $\delta = .76$, and $\phi = 0.13$; again improving the fit substantially with $G^2 = 190$, $\chi^2(5) = 322$, $p < .01$. Finally, to take individual differences into account, the parameter values were estimated for each participant separately, leading to median parameter values of $\alpha = .93$ (*SD* = .24), $\beta = .89$ (*SD* = .28), $\lambda = 1$ (*SD* = 2.5), $\gamma = .77$ (*SD* = .21), $\delta = .76$ (*SD* = .20), and $\phi = 0.18$ (*SD* = 0.26). With these parameter values, prospect theory achieved the best fit of *Mdn* G^2 = 172, which is better than when using only one set of parameters for all participants; $\chi^2(174) = 770$, $p < .01$. In 74% of all cases, the participants chose the gamble that was predicted as most likely by prospect theory. Thus, prospect theory was best at describing individuals' choices when taking the probabilistic nature of preferences and the heterogeneity of people's decision processes into account. When examining prospect theory's parameters, it is surprising that for many participants a loss aversion parameter $\lambda = 1$ was estimated, indicating no loss aversion. I speculate that this is due to the large number of choices participants made in which pairs of gambles with losses, gains, and mixed payoffs occurred in a random order, such that the differing attention they usually pay to losses versus gains diminished.

To compare prospect theory with the probabilistic priority model, the models' AICs were compared for each participant. Here, for 25 of the 30 participants, cumulative prospect theory achieved a better AIC than did the priority model ($p < .01$, according to a sign test). The average AIC difference was ΔAIC = 23, again with ΔAIC greater than 10 providing strong empirical evidence for the better model. These results provide evidence that cumulative prospect theory is more appropriate for describing preferential choices than is the priority model.

Decision field theory. The deterministic version of DFT using a very high decision threshold reached a median fit of *Mdn* G^2 = 501, whereas the probabilistic version of DFT, for which a median decision threshold of $\theta_{DFT} = 1.19$ (*SD* = 0.84) was estimated, reached a much better fit of *Mdn* G^2 = 201, $\chi^2(30) = 9122$, $p < .01$. The DFT fitted to each individual is also much better than DFT using only one single decision threshold for all participants, which led to an average fit of $G^2 = 211$; $\chi^2(29) = 425$, $p < .01$.

[6] In the first step the parameters for three versions of the priority model implementing one additional mechanism were estimated, including the standard deviation of the error component as one free parameter and one of the three remaining mechanisms as additional parameter(s). Of these three versions, the one with the outcome threshold as an additional free parameter improved the fit most when compared with the one-parameter priority model that included only the error component: *Mdn* G^2 = 229, $\chi^2(30) = 390$, $p < .01$. As a second step, the parameters of two versions of the priority model were estimated that included the standard deviation of the error component and the outcome threshold as two free parameters and one of the two remaining mechanisms as additional parameter(s). By adding the probability threshold as a third free parameter, the priority model's fit improved most in comparison with that of the two-parameter model: *Mdn* G^2 = 207, $\chi^2(30) = 529$, $p < .01$. Finally the full priority model with five free parameters achieved a better fit compared with the three-parameter model: *Mdn* G^2 = 201, $\chi^2(60) = 166$, $p < .01$.

[7] The priority model was fitted separately to each individual in order to show the heterogeneity across individuals. The variance observed in the estimated parameters (see Figure 5) indicates that the cognitive process of preferential choice differs substantially across individuals. To justify this interpretation, a priority model was estimated using the same set of parameters for all participants. For this model, the estimated parameter values were a standard deviation for the error component of $\sigma = 0.33$, a probability threshold of $\delta = .44$, an outcome threshold of $\delta = .56$, and the following best order of aspects: minimum outcome, minimum probability, and maximum outcome. This model showed a fit of *Mdn* G^2 = 215, which was significantly worse than the priority model when fitted to each individual separately: $\chi^2(145) = 569$, $p < .01$. Thus, the full priority model with parameter values for each individual proved to be the best model to describe participants' choices.

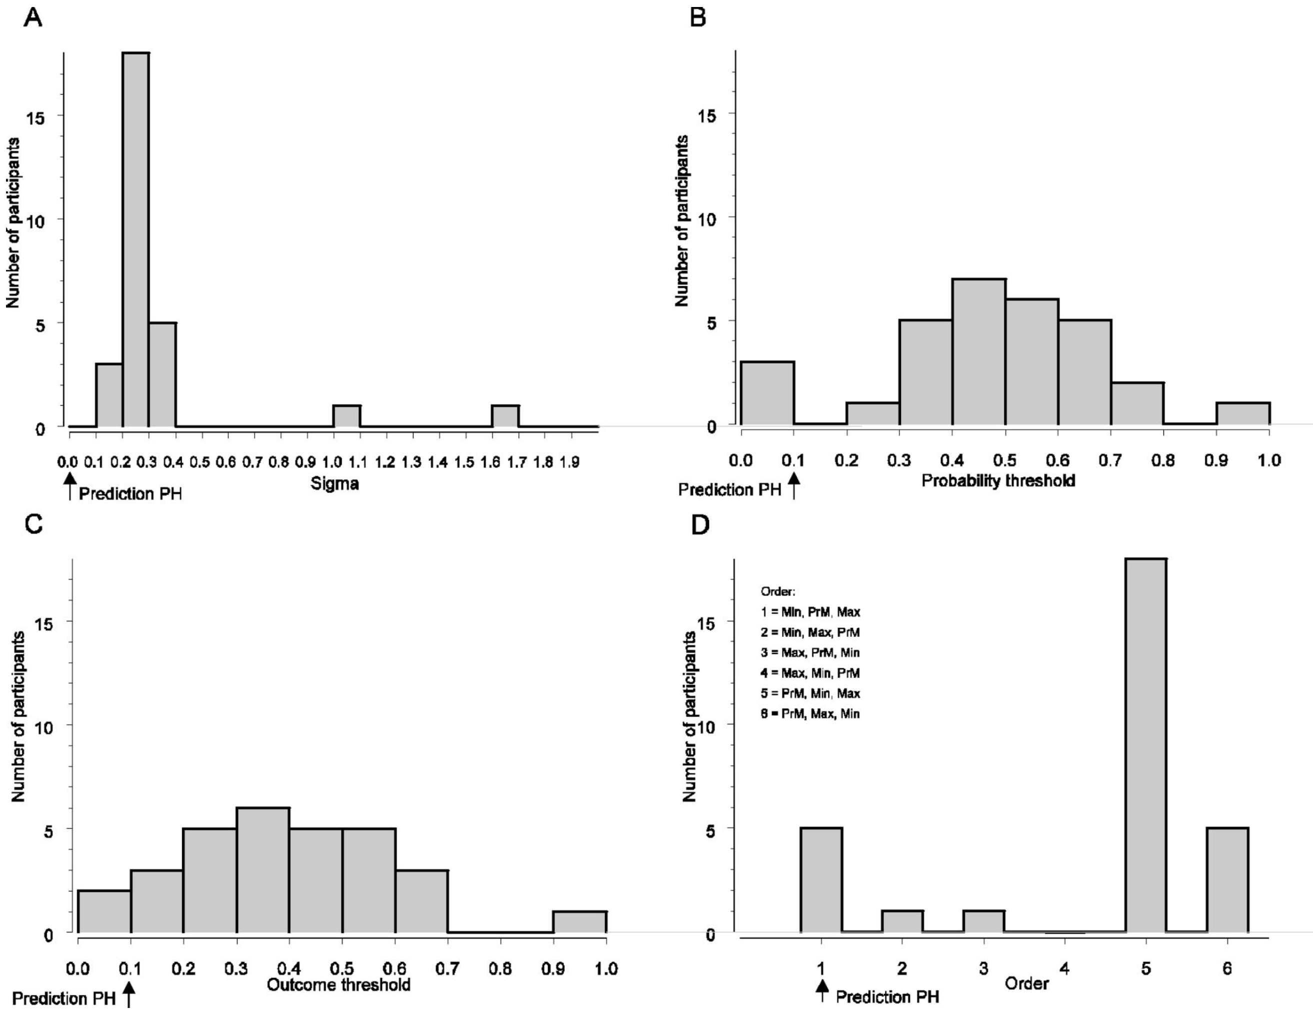

Figure 5. Distributions of the estimated values of the priority model's parameters in Study 2. For all parameters, the estimated values differed substantially from the assumed parameter values of the priority heuristic (PH). A: The distribution of the standard deviation of the error component. Two outliers with a sigma of 8.4 and 10.0 are not shown. B: The distribution of the estimated probability threshold. C: The distribution of the estimated outcome threshold. D: The distribution of the estimated best order of aspects (implemented with two free parameters). Min = minimum outcome; Max = maximum outcome; PrM = probability of minimum outcomes.

Thus, with only one free parameter DFT reached a similar fit as the full priority model did with five parameters. When taking the models' complexity into account by comparing their AIC values, for 22 of 30 participants DFT had a better AIC value in comparison with the priority model ($p = .02$, according to a sign test). DFT had a better AIC value in comparison with that of cumulative prospect theory for 12 of the 30 participants ($p = .36$, according to a sign test). Thus, although DFT did significantly better than the priority model in describing participants' choices, whether DFT or prospect theory provides a better description of the choice behavior could not clearly be determined because both theories predicted the choices of approximately half of the participants best.

Discussion of Study 2

Study 2 tested the different theories for a wide range of choice situations and provided clear evidence that the probabilistic models did better than their deterministic counterparts. Comparing the priority model with the priority heuristic made clear that assuming an error component when aspects are compared and using larger probability and outcome thresholds than assumed by the priority heuristic led to a better description of the observed choices. Comparing the priority model with prospect theory, the latter provided a better account of the data, according to their AIC values. DFT also described the choices better than did the priority model. Yet it could not be decided whether prospect theory or DFT described the choices best. The comparisons of the three theories have to be interpreted carefully because the AIC criterion does not take the parameters' functional relationship into account, as described above. Therefore Study 3 was conducted in which new predictions of the models were determined on the basis of the estimated parameters from Study 2. In sum, for the purpose of comparing the probabilistic and the deterministic versions of the different models, Study 2 provided clear evidence in favor of the probabilistic models. For the purpose of comparing the three different theories, Study 2's main function was to provide estimates for the models' free parameters.

Study 3

The probabilistic theories tested in this article have free parameters, making them flexible in describing various kinds of behavior. The models' flexibility—that is, their complexity—needs to be taken into account in any model comparison test. In Study 2, this was done either by performing statistical tests that took the models' degrees of freedom into account or by using AIC as a model selection criterion. However, as has been discussed, AIC is less appropriate for non-nested models (Pitt et al., 2002; for an opposing view, see also Burnham & Anderson, 2002). Therefore, the main purpose of Study 3 was to provide an unambiguous model comparison test.

In Study 3, the three models' predictions were tested completely independently of the observed data—that is, without fitting the models' parameters. Instead, the models' predictions were based on Study 2's estimated parameters, so that Study 3 represents a generalization test of the theories (Busemeyer & Wang, 2000). In Study 2, a random, "representative" set of gambles was selected for estimating the models' parameters. To perform a strong test of the theories in Study 3, choice situations were selected for which the models made diametrically opposed predictions; thus, here the models also differ in their qualitative predictions. In this way, Study 3 presents the crucial model comparison test, capable of providing a clear-cut answer as to which theory is best in predicting decisions under risk.

Method

Study 3 had three between-subjects conditions, with 30 participants each. The first condition tested the priority heuristic against the priority model, the second condition tested the priority model against cumulative prospect theory, and the third condition tested cumulative prospect theory against DFT. For each condition, a specific item set was generated.

Materials. For the first condition, an item set was generated by first determining the priority heuristic's predictions for all 9,598 pairs of stochastically non-dominated gambles from Study 2. To determine the priority model's predictions, a bootstrapping technique was employed (Efron, 1979). For each of the 9,598 pairs, an experiment with 30 participants was simulated, assuming that each participant would make choices according to the priority model. For each simulated participant, the priority model's parameter values were selected with replacement from the joint distribution of parameter values estimated for the 30 participants in Study 2. Then, for the simulated 30 participants, the average predicted choice probability of choosing Gamble A over Gamble B was determined. Because the predictions for these 30 simulated participants differed depending on the sampled parameter values, the simulated experiment was repeated 100 times, and the average choice probabilities were determined. Finally, 60 pairs for the gain, loss, and mixed payoff domains were selected from the 9,598 pairs of gambles for which the probabilistic predictions of the priority model differed most from the priority heuristic's prediction. For the selected 180 pairs of gambles, the priority heuristic predicted the choice of one gamble, whereas the priority model always predicted the choice of the other gamble with a probability above .50; this predicted choice probability was on average .78 (ranging from .70 to .87).

For the second condition of Study 3, an item set was generated that allowed an indicative test of the predictions of the priority model and prospect theory. Again, for each of the 9,598 pairs of gambles, 100 experiments with 30 participants were simulated, and the average choice probabilities were determined. For each simulated participant, the parameter values for prospect theory were selected with replacement from the joint distribution of parameter values estimated in Study 2, and each simulated participant made choices according to prospect theory. Finally, from the 9,598 pairs of gambles, all pairs for which the priority model and cumulative prospect theory predicted opposite gambles with an average probability larger than .50 were selected. From these pairs, 60 pairs of gambles for each domain were selected for which the predicted choice probabilities differed most. This led to a set of 180 pairs for which the priority model predicted the choice of one gamble with an average probability of .69 (SD = .08) and for which cumulative prospect theory predicted the choice of the other gamble with an average probability of .77 (SD = .09).

For the third condition in Study 3, an item set was generated applying the same method used for the second condition to allow an indicative test of the predictions of prospect theory and DFT. This led to a set of 180 pairs for which prospect theory predicted the choice of one gamble with an average probability of .73 (SD = .12) and for which DFT predicted the choice of the other gamble with an average probability of .71 (SD = .13).

Participants and procedure. Ninety people (54 women and 36 men) with an average age of 23 years participated in the experiment. Participants were students from various departments at the Free University of Berlin. The experimental task lasted approximately 1 hr and was the same as in Study 2, with the only difference that different item sets were used. The participants received a payment of €15 (the equivalent of $18 at the time of the study). In addition, after completing the task, one pair of the 180 gambles was selected randomly, the participant's chosen gamble was played out in front of him or her, and 5% of the gamble's payoff was added to or subtracted from the €15 payment, resulting in possible payments between €10 and €20. The instructions, procedure, and the experimental software were the same as in Study 2 (see Figure 3).

Results

First, I analyzed whether participants' choices were systematic. Thereafter, I examined which of the three theories predicted the choices best.

Choice proportions. The correlation between the differences of the first and second gamble's expected values and the proportion of participants choosing the first gamble was determined for all three conditions, resulting in $r = .88$ ($p < .01$) when testing the priority model against the priority heuristic; $r = .61$ ($p < .01$) when testing the priority model against prospect theory; and $r = .71$ ($p < .01$) when testing prospect theory against DFT. Thus, participants showed systematic behavior. The gamble with the larger expected value was selected with a maximum choice proportion of 100% for five, two, and two pairs of gambles for the first, second, and third experimental condition, respectively.

Model comparison. The priority model was tested against the priority heuristic by comparing which model predicted the choices best for each participant. For all 30 participants, the priority model achieved a better G^2 score than did the priority heuristic, with *Mdn* $G^2 = 186$ for the former and *Mdn* $G^2 = 1{,}276$ for the latter. Note, when predicting the choice between two gambles with an equal probability of .50, a G^2 of 249.5 results. When comparing the priority model with the priority heuristic with a naïve error theory (with a *Mdn* $G^2 = 281$), the priority model was preferable for 24 of the 30 participants.

In addition, I determined how often each participant chose the gamble predicted as most likely by the priority model versus the priority heuristic. For the majority of pairs of gambles, 29 participants chose the gamble predicted as most likely according to the priority model. Only 1 participant made the choices predicted as most likely according to the priority model in as few as 50% of all cases. In 74% of all pairs of gambles, participants made the choice most likely predicted by the priority model. Consequently, the gambles predicted by the priority heuristic were chosen in only 26% of all cases. For the gain, loss, and mixed gamble domains, the gambles predicted by the priority heuristic were chosen in 21%, 20%, and 35% of all cases, respectively. In sum, the priority model clearly outperformed the priority heuristic in predicting participants' preferential choices.

In the second condition, the priority model was tested against cumulative prospect theory. The priority model's fit of *Mdn* $G^2 = 314$ was inferior to prospect theory's fit of *Mdn* $G^2 = 261$. For the majority of the participants—22 participants of 30—prospect theory achieved a better G^2 score. In addition, I determined how often each participant had chosen the gamble predicted as most likely according to the two models: Here, the majority—27 participants—chose mainly the gambles predicted by prospect theory, whereas only 3 chose mainly those predicted by the priority model. The gambles predicted as most likely by prospect theory were chosen in 63% of all cases, and consequently, those of the priority model in only the remaining 37%. In sum, prospect theory predicted the choice better than did the priority model. However, note that prospect theory did not do better than a pure chance prediction, which would yield a fit of $G^2 = 249.5$.

Finally, in the third condition, cumulative prospect theory was tested against DFT. Cumulative prospect theory's fit of *Mdn* $G^2 = 412$ was inferior to DFT's fit of *Mdn* $G^2 = 206$. For all 30 participants, DFT achieved a better G^2 score compared with that of prospect theory. In addition, I determined how often each participant had chosen the gamble predicted as most likely according to the two models: Here, the majority—27 participants—chose mainly the gambles predicted by DFT, whereas only 3 chose mainly those predicted by prospect theory. The gambles predicted as most likely by DFT were chosen in 66% of all cases, and consequently, those of prospect theory in only the remaining 34%.

To illustrate the results, Figure 6 shows for each condition and all 180 pairs of gambles a scatter plot of the models' predicted choice probabilities of choosing one of the two gambles as well as the proportion of participants actually choosing the gamble. The results are clear cut: In the first condition, when testing the priority model against the priority heuristic, the former comes out the clear winner in predicting participants' choices (see Panels A1 and A2 of Figure 6). The observed choice proportions correlate strongly with the priority models' predictions ($r = .79$) and consequently were negatively correlated with the priority heuristic's predictions ($r = -.76$). In the second condition, when testing the priority model against prospect theory, prospect theory proves superior (see Panels B1 and B2 of Figure 6). Here the observed choice proportions correlate strongly with prospect theory's predictions ($r = .50$) and consequently were negatively correlated with the priority model's predictions ($r = -.41$). Finally, in the third condition, when comparing prospect theory with DFT, the best model to predict the choices was DFT (see Panels C1 and C2 of Figure 6). The observed choice proportions correlate strongly with DFT's predictions ($r = .77$) and consequently were negatively correlated with the prospect theory's predictions ($r = -.69$).

Only the third condition of Study 3 was designed to test DFT's predictions. Because DFT performed so well, I additionally analyzed how well DFT predicted the choices in the first two conditions. DFT's predictions were determined by using the bootstrapping technique described above. In the first condition, DFT reached a fit of *Mdn* $G^2 = 157$ compared with the priority model with *Mdn* $G^2 = 186$ and cumulative prospect theory with *Mdn* $G^2 = 168$. For 24 of the 30 participants, DFT predicted the choices better than did the priority model (i.e., DFT had a lower G^2; $p < .01$ according to a sign test). In comparison with cumulative prospect theory, DFT predicted the choices better for 18 of the 30 participants ($p = .36$). In the second condition, in which prospect theory performed better than the priority model, DFT reached a fit of *Mdn* $G^2 = 223$ compared with cumulative prospect theory with *Mdn* $G^2 = 261$ and the priority model with *Mdn* $G^2 = 314$. DFT predicted the choices better than did the priority model for 29 of the 30 participants and better than did prospect theory for all participants. In sum, DFT overall provided the best account of the choice behavior for all conditions in Study 3. It always performed better than the priority model. When compared with prospect theory, DFT in particular performed better in the second and third conditions. Condition 3 is most instructive, because here the two models made qualitatively different predictions; that is, one model predicted the choice of one gamble, whereas the other model predicted the opposite choice.

General Discussion

The general goal of this article has been to illustrate the superiority of probabilistic choice models in comparison with their deterministic counterparts in predicting decisions under risk. To accomplish this goal, I constructed a probabilistic choice model that generalizes the deterministic priority heuristic proposed by Brandstätter et al. (2006). Likewise, I constructed a probabilistic version of the deterministic cumulative prospect theory of Tversky and Kahneman (1992). Moreover, I considered decision field theory (Busemeyer & Townsend, 1993), a model developed specifically to explain the probabilistic nature of preferential choice, and compared it with its deterministic counterpart. The probabilistic versions for all three theories were always more suitable for describing decisions under risk than were the deterministic versions, even when taking the larger complexity of the probabilistic models into account. When testing the three probabilistic models against each other, DFT did best, followed by cumulative prospect theory.

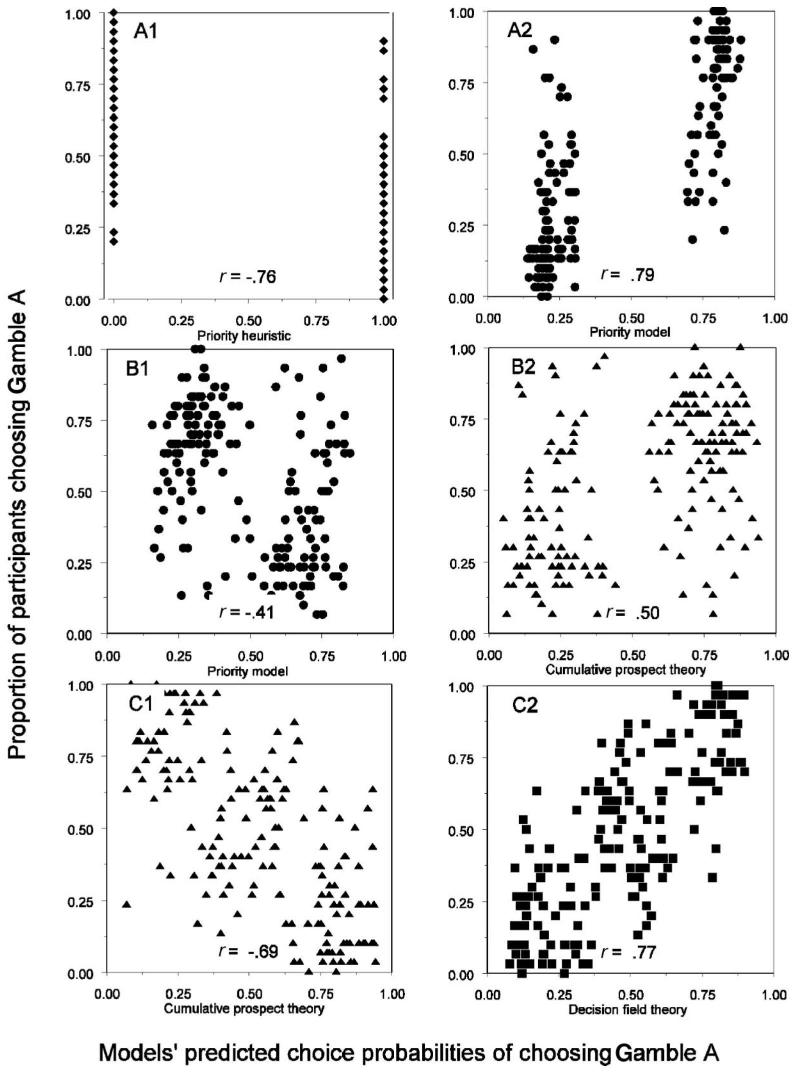

Figure 6. The theories' predicted probabilities of choosing Gamble A and the proportions of participants actually choosing Gamble A for all three experimental conditions. Each panel shows the correlation between the predicted choice probabilities and the observed choice proportions. A1 and A2: Results for the first experimental condition of comparing the priority heuristic (A1) with the priority model (A2). B1 and B2: Results for the second condition of comparing the priority model (B1) with cumulative prospect theory (B2). C1 and C2: Results for the third condition of comparing cumulative prospect theory (C1) with decision field theory (C2).

Probabilistic Models of Preferences

Deterministic theories of human behavior are popular in the social sciences, for several reasons. First, they are often more transparent than probabilistic models. Therefore, from a pragmatic scientific perspective, a deterministic theory often has the advantage of allowing an unambiguous interpretation of observations. Second, by not specifying an error theory, a deterministic model is often simpler than a probabilistic model. When defining the complexity of a model by the ratio of the data the theory can predict to the theoretically possible data (Myung, 2000), probabilistic theories are often more complex because they are able to predict a larger data space. Naturally, when applying Occam's razor to two theories that are equally good at predicting behavior, the simpler one should be preferred (Forster, 2000; Myung & Pitt, 1997). Finally, due to their usually lower complexity, deterministic models can be easily falsified. This is presumably their most important advantage because when the data do not succeed in falsifying a theory that can in principle be easily falsified, this provides strong empirical support for that theory. In contrast, when the data do not falsify a theory—because the theory can produce any possible pattern of data—the empirical support of the data is of no value (Roberts & Pashler, 2000).

Because of these advantages, deterministic theories often provide a good starting point for explaining behavioral phenomena. However, I argue that a probabilistic model provides a more adequate description of human behavior. The research approach taken in the present article, of generalizing a deterministic model to a probabilistic choice model, could in principle also provide empirical support for the idea of the deterministic model. For instance, behavioral observations could falsify a deterministic model but could still be consistent with the probabilistic model. For an illustration of this point, consider two recent studies by Lee and Cummins (2004) and Bergert and Nosofsky (2007), below.

Lee and Cummins (2004) tested two deterministic models of probabilistic inferences. The inference task was to judge which of two objects had a higher criterion value based on various cues. The first model examined was a lexicographic heuristic (Gigerenzer & Goldstein, 1996) that assumes cues are considered sequentially in a specific order to make an inference. The second deterministic model was a Bayesian linear model that assumes the information from all cues is integrated for an inference. Lee and Cummins found that the Bayesian model was superior to the heuristic. Bergert and Nosofsky (2007) replicated Lee and Cummins's study but tested two probabilistic generalizations of the two deterministic models. In particular, the probabilistic generalizations of the lexicographic heuristic allowed for varying orders in which cues are considered. Surprisingly—and in contrast to Lee and Cummins's results—the generalized version of the lexicographic heuristic provided a better account of the data than did the Bayesian model. Apparently the participants did not integrate the available information as assumed by the Bayesian model. Instead, they made their inferences on the basis of single cues they searched for sequentially, although the order in which the cues were considered varied across participants and differed from the order assumed by the deterministic heuristic. This illustrates that an experimental test could reject a deterministic model, although the model's basic conceptualization is accurate.

The Probabilistic Nature of Preferential Choice

The probabilistic priority model does better than the priority heuristic in predicting people's preferential choices. First, it explains varying choice proportions across different pairs of gambles, as Figure 2 illustrates. This shows, first, that a probabilistic model of behavior accounts for constant inconsistencies, that is, for why a person makes different choices when presented with an identical choice problem, as illustrated by Hey (2001). Second, and more importantly, the model explains why the choice proportions vary for different problems. Implementing a naïve error theory within the priority heuristic cannot explain these varying choice proportions. Brandstätter et al. (2006) argued that the priority heuristic predicts a rank order of choice proportions by the "assumption that the earlier the examination stops, the more extreme the choice proportions will be" (p. 428). The authors did not explain the rationale underlying this assumption, but presumably, the later the examinations stop, the greater the number of potential errors, leading to lower choice proportions. The priority model is in line with this assumption, but it makes fine-grained, quantitative predictions that may differ from the priority heuristic. For instance, it predicts a moderate choice probability even if the first aspect considered is decisive, namely, when the comparison difference is barely larger than the decision threshold.

Study 2 illustrated that the assumption of different parameter values for the building blocks of the priority heuristic provides a more accurate description of individuals' choices. With larger decision thresholds for the probabilities and the payoffs, the priority model made superior predictions. Likewise, when allowing for different orders in which aspects were considered, the choice process could be described better. Finally, by fitting the priority model to each individual separately, it became clear that individuals differ substantially in the way they make choices. For instance, three different orders in which the gambles are compared were obtained for large groups of participants (cf. Figure 5D). In sum, the priority model has several advantages in comparison with the priority heuristic for explaining preferential choice.

Cumulative Prospect Theory and Decision Field Theory

The present article illustrated the advantages of probabilistic theories over their deterministic counterparts. I have shown that a probabilistic version of cumulative prospect theory did better than the deterministic version. Likewise, the probabilistic DFT did better than a deterministic DFT. In comparison with the priority model and prospect theory, DFT performed best in predicting choices, in particular when considering Study 3's crucial model comparison tests. These findings are in contrast to the conclusions of Brandstätter et al. (2006), who reported that the priority heuristic did better than prospect theory. How can this difference in findings be explained? To understand the theories' qualitative predictions, I analyzed the predictions of the priority heuristic, cumulative prospect theory, and DFT for the set of 9,598 pairs of nondominated gambles described in Study 2. First, for each of these pairs, DFT always predicts that the gamble with the larger expected value will be chosen with a probability larger than 50%. Second, the gamble predicted by the priority heuristic and the most likely gamble predicted by cumulative prospect theory were determined (using the bootstrapping technique of Study 3 and prospect theory's estimated parameters of Study 2). Table 3 illustrates the qualitatively different relationships among the theories' predictions. Cumulative prospect theory predicts that the gamble with the larger expected value will be most likely chosen for 91% of the pairs. In stark contrast, the priority heuristic chooses the gamble with the larger expected value in only 60% of all pairs. This analysis provides an explanation of why Brandstätter et al. (2006) found support for the priority heuristic for the experimental results of Tversky and Kahneman (1992). In this experiment, gambles were specifically selected to illustrate violations of the expected utility approach, which the priority heuristic can explain.

In general, the priority heuristic is able to predict various violations of the expected utility approach, including intransitive choices. A theory that can predict intransitive choices will often predict very different behavior than will the expected utility theory. Therefore, the ability to predict intransitive choices is a very useful characteristic to compare theories. In principle, there is no question that people make intransitive choices (e.g., Birnbaum, Patton, & Lott, 1999), but whether intransitive choices represent a general and robust behavioral phenomenon is open to debate (see Rieskamp et al., 2006). I argue that a theory of decisions under risk should predict violations of transitivity as the exception rather than the rule. Therefore the priority heuristic appears to be positioned in too stark of an opposition to the expected utility perspective. One could speculate that the priority heuristic accurately represents difficult choices. It may be that people first try to choose the gamble with the larger expected value (utility) and select the priority heuristic when this becomes too complicated. It has been argued that one needs to give up the idea of a single heuristic underlying people's choices (Johnson, Schulte-Mecklenbeck, &

Table 3
Cross-Table of the Predictions of the Priority Heuristic, Cumulative Prospect Theory, and Decision Field Theory in Relation to the Gambles' Expected Values

	Priority heuristic		
CPT	PH = EV/DFT	PH ≠ EV/DFT	Total
CPT = EV/DFT	57%	34%	91%
CPT ≠ EV/DFT	3%	6%	9%
Total	60%	40%	

Note. PH = priority heuristic; CPT = cumulative prospect theory; DFT = decision field theory; EV = expected value. Predictions = the most likely gamble predicted by DFT always corresponds with the gamble that has the larger expected value; CPT = EV/DFT (CPT ≠ EV/DFT) = pairs of gambles where the most likely gamble predicted by CPT corresponds with the gamble that has the larger (lower) expected value; PH = EV/DFT (PH ≠ EV/DFT) = pairs of gambles where the gamble predicted by the PH corresponds with the gamble that has the larger (lower) expected value.

Willemsen, 2008). The idea that people select different strategies depending on the decision context has been successfully examined (see Dieckmann & Rieskamp, 2007; Erev & Barron, 2005; Gaissmaier, Schooler, & Rieskamp, 2006; Mata, Schooler, & Rieskamp, 2007; Payne et al., 1988; Rieskamp, 2006, 2008; Rieskamp & Hoffrage, 2008; Rieskamp & Otto, 2006; von Helversen & Rieskamp, 2008).

In general, it has been argued that for model comparisons tests, such as testing the priority heuristic against alternative models, one should focus on decision situations in which the models make different predictions (e.g., Birnbaum, 2008; Rieger & Wang, 2008). Furthermore, for estimating a model's free parameters on which its predictions are based, and for testing the model's predictions, the same goodness-of-fit criterion should be used. This is precisely what has been accomplished in the present studies. Studies 2 and 3 jointly present a transparent, fair, and strong model comparison approach, which can also be generalized to future research. The results of the model comparison tests reveal that when the priority heuristic is applied to describe behavior for a representative set of gambles, as in Study 2, it performs worse than does prospect theory. In general, depending on how the set of gambles is selected, different conclusions about the theories can result. For this reason, Study 3 was conducted with a set of gambles for which the pairs of models under consideration made diametrically opposed predictions, and in this crucial model comparison study, DFT and cumulative prospect theory did better than the priority heuristic and the priority model. DFT also outperformed cumulative prospect theory, in particular when focusing on the crucial model comparison test of Study 3. Because DFT predicts that the gamble with the larger expected value will be chosen with a probability larger than 50%, this implies that cumulative prospect theory always predicts the choice of the gamble with the lower expected value. In such a situation, the strong empirical support for DFT shows that people do not generally choose the gamble with the lower expected value. Thus, although past research has shown that people violate the expected utility approach, the present studies suggest that these violations are more the exception than the rule.

What are the underlying principles of DFT that make the model so successful? First, the model in general predicts that the larger the difference in the gambles' expected values, the larger the probability that the gamble with the larger expected value will be chosen. This seems to be a trivial assumption, but it is not incorporated in many choice theories, including, for instance, the priority heuristic. Second, the model assumes that the gambles are compared with each other by computing the differences of the gambles' outcomes. These differences are accumulated in an overall difference score until the score passes an individual decision threshold. This accumulation process implies that the lower the variance of the differences accumulated, the more quickly the decision threshold is reached, thus implying a larger choice probability for the gamble with the higher expected value. For instance, suppose a decision maker has to choose between Gamble A that leads with an equal chance to a payoff of either \$1 or \$5, Gamble B that leads with an equal chance to either \$4 or \$6, or Gamble B′ that leads with an equal chance to either \$2 or \$8. Although Gamble B and Gamble B′ have the same expected value, Gamble B will be chosen with a greater probability than Gamble B′ when compared with Gamble A, due to the smaller variance of the outcomes of Gamble B. Moreover, although Gamble B or Gamble B′ will be chosen with a probability larger than 50%, Gamble A will nevertheless be chosen with a substantial probability, representing a violation of stochastic dominance. The accumulation of differences allows DFT to explain other expected utility violations, such as violations of the independence from irrelevant alternatives principle (see Busemeyer & Townsend, 1993).

When examining DFT's specification in Equation 2 above, one might have noticed the similarities between DFT and Thurstone's "law of comparative judgments" (Thurstone, 1927). Thurstone's model (Case 1) differs from DFT by assuming that the variance and the covariance of the differences depend on an individual's idiosyncratic ability to discriminate between options; in DFT the variance is the result of a cognitive process of accumulated differences, so that the variance and the covariance are defined by the options' attribute values (i.e., the gambles' outcomes and their probabilities). In Thurstone's model, the variance and the covariance have to be estimated, but for an unequivocal estimation the model has too many free parameters, so that it is additionally assumed that the covariance stays constant across decision trials, whereas for DFT the covariance varies across trials. Moreover, DFT assumes that a decision is made only when the accumulated difference passes a threshold, whereas Thurstone assumed that people can discriminate even very small differences.

However, it should be emphasized that I used only a simplified version of DFT. For this simplification, I assumed that the gambles' outcomes were represented by identical subjective values and likewise that the objective probabilities were represented with identical subjective weights. By using psychologically more plausible subjective representations of outcomes and probabilities, DFT might provide an even better account of people's decisions under risk. In sum, the most striking result of the present research is that it shows that DFT outperforms prospect theory, the most influential alternative view to the expected utility approach.

Final Conclusions

The present studies demonstrated that probabilistic models of decisions under risk provide a better account of human decision making than do their deterministic counterparts. Of these probabilistic models, DFT and cumulative prospect theory outperformed the priority model, and when testing DFT against cumulative prospect theory, the former predicted people's decisions more accurately. Most importantly, the present studies have shown that any accurate theory of decisions under risk needs to take the probabilistic nature of preferential choice into account.

References

Albers, W. (2001). Prominence theory as a tool to model boundedly rational decisions. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 297–317). Cambridge, MA: MIT Press.

Bell, D. E. (1982). Regret in decision-making under uncertainty. *Operations Research, 30,* 961–981.

Bell, D. E. (1985). Disappointment in decision-making under uncertainty. *Operations Research, 33,* 1–27.

Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33,* 107–129.

Birnbaum, M. H. (2005). A comparison of five models that predict violations of first-order stochastic dominance in risky decision making. *Journal of Risk & Uncertainty, 31,* 263–287.

Birnbaum, M. H. (2008). Evaluation of the priority heuristic as a descriptive model of risky decision making: Comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review, 115,* 235–260.

Birnbaum, M. H., & Chavez, A. (1997). Tests of theories of decision making: Violations of branch independence and distribution independence. *Organizational Behavior and Human Decision Processes, 71,* 161–194.

Birnbaum, M. H., Patton, J. N., & Lott, M. K. (1999). Evidence against rank-dependent utility theories: Tests of cumulative independence, interval independence, stochastic dominance, and transitivity. *Organizational Behavior and Human Decision Processes, 77,* 44–83.

Block, H. D., & Marschak, J. (1960). Random orderings and stochastic theories of responses. In I. Olkin, S. Ghurye, W. Hoeffding, W. Madow, & H. Mann (Eds.), *Contributions to probability and statistics* (pp. 90–132). Stanford, CA: Stanford University Press.

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review, 113,* 409–432.

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multimodel inference.* New York: Springer.

Busemeyer, J. R. (1980). Importance of measurement theory, error theory, and experimental design for testing the significance of interactions. *Psychological Bulletin, 88,* 237–244.

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.

Busemeyer, J. R., & Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology, 44,* 171–189.

Camerer, C. F. (1989). An experimental test of several generalized utility theories. *Journal of Risk & Uncertainty, 2,* 61–104.

Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition, 35,* 1801–1813.

Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics, 7,* 1–26.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112,* 912–931.

Erev, I., Roth, A., Slonim, R. L., & Barron, G. (2002). *Combining a theoretical prediction with experimental evidence to yield a new prediction: An experimental design with a random sample of tasks.* Unpublished manuscript.

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review, 101,* 519–527.

Estes, W. K., & Maddox, W. T. (2005). Risks of drawing inferences about cognitive processes from model fits to individual versus average performance. *Psychonomic Bulletin & Review, 12,* 403–408.

Falmagne, J. C. (1976). Random conjoint measurement and loudness summation. *Psychological Review, 83,* 65–79.

Fishburn, P. C. (1983). Transitive measurable utility. *Journal of Economic Theory, 31,* 293–317.

Forster, M. R. (2000). Key concepts in model selection: Performance and generalizability. *Journal of Mathematical Psychology, 44,* 205–231.

Gaissmaier, W., Schooler, L. J., & Rieskamp, J. (2006). Simple predictions fueled by capacity limitations: When are they successful? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 966–982.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103,* 650–669.

González-Vallejo, C. (2002). Making trade-offs: A probabilistic and context-sensitive model of choice behavior. *Psychological Review, 109,* 137–155.

Haider, H., & Frensch, P. A. (2002). Why aggregated learning follows the power law of practice when individual learning does not: Comment on Rickard (1997, 1999), Delaney et al. (1998), and Palmeri (1999). *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 392–406.

Hey, J. D. (2001). Does repetition improve consistency? *Experimental Economics, 4,* 5–54.

Johnson, E. J., Schulte-Mecklenbeck, M., & Willemsen, M. C. (2008). Process models deserve process data: Comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review, 115,* 263–273.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291.

Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the "take the best" and the "rational" models. *Psychonomic Bulletin & Review, 11,* 343–352.

Leland, J. W. (1994). Generalized similarity judgments: An alternative explanation for choice anomalies. *Journal of Risk & Uncertainty, 9,* 151–172.

Loomes, G., & Sugden, R. (1982). Regret theory—An alternative theory of rational choice under uncertainty. *Economic Journal, 92,* 805–824.

Lopes, L. L., & Oden, G. C. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and SP/A theory. *Journal of Mathematical Psychology, 43,* 286–313.

Luce, R. D. (1959). *Individual choice behavior.* Oxford: Wiley.

Luce, R. D. (1977). The choice axiom after twenty years. *Journal of Mathematical Psychology, 15,* 215–233.

Luce, R. D. (2000). *Utility of gains and losses: Measurement-theoretical and experimental approaches.* Mahwah, NJ: Erlbaum.

Luce, R. D., & Fishburn, P. C. (1991). Rank- and sign-dependent linear utility models for finite first order gambles. *Journal of Risk & Uncertainty, 4,* 29–59.

Luce, R. D., & Fishburn, P. C. (1995). A note on deriving rank-dependent utility using additive joint receipts. *Journal of Risk & Uncertainty, 11,* 5–16.

Luce, R. D., & Suppes, P. (1965). Preference, utility, and subjective probability. In R. D. Luce, R. B. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 3, pp. 249–410). New York: Wiley.

Machina, M. J. (1989). Dynamic consistency and non-expected utility-models of choice under uncertainty. *Journal of Economic Literature, 27,* 1622–1668.

Marley, A. A. J. (2002). Random utility models and their applications: Recent developments. *Mathematical Social Sciences, 43,* 289–302.

Mata, R., Schooler, L. J., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of decision strategies. *Psychology & Aging, 22,* 796–810.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics* (pp. 105–142). New York: Academic Press.

Mellers, B. A. (2000). Choice and the relative pleasure of consequences. *Psychological Bulletin, 126,* 910–924.

Mellers, B. A., & Biagini, K. (1994). Similarity and choice. *Psychological Review, 101,* 505–518.

Mosteller, F., & Nogee, P. (1951). An experimental measurement of utility. *Journal of Political Economy, 59,* 371–404.

Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology, 44,* 190–204.

Myung, I. J., & Pitt, M. A. (1997). Applying Occam's razor in modeling cognition: A Bayesian approach. *Psychonomic Bulletin & Review, 4,* 79–95.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal, 7,* 308–313.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 534–552.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109,* 472–491.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization, 3,* 323–343.

Quiggin, J. (1993). *Generalized expected utility theory: The rank-dependent model.* Boston: Kluwer.

Rieger, M. O., & Wang, M. (2008). What is behind the priority heuristic? A mathematical analysis and comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review, 115,* 274–280.

Rieskamp, J. (2006). Perspectives of probabilistic inferences: A reinforcement and an adaptive network model compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 1355–1370.

Rieskamp, J. (2008). The importance of learning when making inferences. *Judgment and Decision Making, 3,* 261–277.

Rieskamp, J., Busemeyer, J. R., & Mellers, B. (2006). Extending the bounds of rationality: A review of research on preferential choice. *Journal of Economic Literature, 44,* 631–661.

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica, 127,* 258–276.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General, 135,* 207–236.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review, 107,* 358–367.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review, 108,* 370–392.

Rubinstein, A. (1988). Similarity and decision-making under risk (is there a utility-theory resolution to the Allais paradox?). *Journal of Economic Theory, 46,* 145–153.

Savage, L. J. (1954). *The foundations of statistics.* New York: Wiley.

Sokal, R. R., & Rohlf, F. J. (1994). *Biometry: The principles and practice of statistics in biological research.* New York: Freeman.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature, 38,* 332–382.

Thorngate, W. (1980). Efficient decision heuristics. *Behavioral Science, 25,* 219–225.

Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review, 34,* 273–286.

Thurstone, L. L. (1959). *The measurement of values.* Chicago: University of Chicago Press.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review, 79,* 281–299.

Tversky, A., & Kahneman, D. (1992). Advances in prospect-theory—Cumulative representation of uncertainty. *Journal of Risk & Uncertainty, 5,* 297–323.

von Helversen, B., & Rieskamp, J. (2008). The mapping model: A cognitive theory of quantitative estimation. *Journal of Experimental Psychology: General, 137,* 73–96.

von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton, NJ: Princeton University Press.

Wong, R. S. K. (1994). Model selection-strategies and the use of association models to detect group-differences. *Sociological Methods & Research, 22,* 460–491.

(*Appendixes follow*)

Appendix A

Specification of the Priority Model

The priority model assumes that when comparing two gambles with two possible outcomes each, three aspects are considered sequentially: the minimum outcome of each gamble; the probability of each minimum outcome occurring, called the *minimum probability*; and the maximum outcome of each gamble. For gambles that lead only to positive payoffs, or to both positive and negative payoffs, the minimum outcome refers to the minimum gain (i.e., the lowest payoff), and the maximum outcome refers to the maximum gain (i.e., the highest payoff). For gambles with only negative payoffs, the minimum outcome refers to the minimum loss (i.e., the highest payoff), and the maximum outcome refers to the maximum loss (i.e., the lowest payoff). When comparing the minimum outcomes of two Gambles A and B, the subjective difference

$$d_{s\_min}(\min_A,\min_B) = d_{min}(\min_A,\min_B) + \varepsilon_{min} \quad \text{(A1)}$$

is determined, where $\min_A$ and $\min_B$ refer to the minimum outcomes of gambles A and B, respectively, $d_{min} = \min_A - \min_B$, and where ε_{min} is an error component. Following a Thurstonian approach, it is assumed that the error component is identically and independently normally distributed with mean zero and the standard deviation of the error component σ_{min}, which is a free parameter. The subjective difference is then compared with the threshold $\delta_{out} = \delta^*_{out} \times \max(\max_A,\max_B)$, where $0 < \delta^*_{out} < 1$ is a free outcome threshold parameter, and $\max(\max_A,\max_B)$ is the maximum of the two maximum outcomes of the two gambles being compared. When d_{s_min} is larger than the threshold δ_{out}, Gamble A is chosen; when d_{s_min} is smaller than $-\delta_{out}$, Gamble B is chosen, and when d_{s_min} lies in between, the next aspect is considered. The probability of choosing Gamble A over Gamble B based on the minimum outcomes is determined by the probability with which d_{min} falls above the threshold; that is, $p_{min}(A,B) = p[d_{s\_min} > \delta_{out}) = F[(d_{min} - \delta_{out})/\sigma_{min}]$ where F is the standard normal cumulative distribution function.

When comparing the probability with which the minimum outcomes occur, the subjective difference

$$d_{s\_pr}(pr_{minA},pr_{minB}) = d_{pr}(pr_{minA},pr_{minB}) + \varepsilon_{pr} \quad \text{(A2)}$$

is determined, where pr_{minA} and pr_{minB} refer to the probabilities with which the minimum outcomes of Gambles A and B, respectively, occur, $d_{pr} = pr_{minA} - pr_{minB}$, and where the error ε_{pr} is normally distributed with mean zero and σ_{pr} as a free parameter. Because probabilities vary between 0 and 1, the distribution of ε_{pr} is truncated, such that d_{s_pr} only varies between -1 and $+1$. The difference d_{s_pr} is compared with the probability threshold δ_{pr}, which is a free parameter varying between 0 and 1. When d_{s_pr} is larger than δ_{pr}, Gamble A is chosen; when d_{s_pr} is smaller than $-\delta_{pr}$, Gamble B is chosen; and when d_{s_pr} lies in between, the next aspect is considered. The probability of choosing Gamble A over Gamble B based on the probabilities of the minimum outcomes is determined by the probability with which d_{s_pr} falls above the threshold; that is, $p_{pr}(A,B) = p[d_{s\_pr} > \delta_{pr}) = F[(d_{pr} - \delta_{pr})/\sigma_{pr}]$ where F is the truncated standard normal cumulative distribution function (normalized to a cumulative probability of one).

When comparing the maximum outcomes, the difference

$$d_{s\_max}(\max_A,\max_B) = d_{max}(\max_A,\max_B) + \varepsilon_{max} \quad \text{(A3)}$$

is determined, where $\max_A$ and $\max_B$ refer to the maximum outcomes of Gambles A and B, respectively, $d_{max} = \max_A - \max_B$, and where ε_{max} is normally distributed with mean zero and σ_{max} as a free parameter. The difference is then compared with the threshold δ_{out} described above. When d_{s_max} is larger than the threshold δ_{out}, Gamble A is chosen, when d_{s_max} is smaller than $-\delta_{out}$, Gamble B is chosen, and when d_{s_max} lies in between, the next aspect is considered. The probability of choosing Gamble A over Gamble B based on the maximum outcomes is determined by the probability with which d_{s_max} falls above the threshold; that is, $p_{max}(A,B) = p[d_{s\_max} > \delta_{out}) = F[(d_{max} - \delta_{out})/\sigma_{max}]$ where F is the standard normal cumulative distribution function.

The priority model assumes that an individual may choose any one of the possible orders in which the three aspects can be considered, in contrast to the priority heuristic, which assumes a specific order. To parameterize the order in which the aspects are considered, two free parameters are used. Following Bergert and Nosofsky's (2007) method, it is assumed that each of the three aspects has a specific importance weight w varying between a value of 0 and 1, with w_{min} for the minimum outcomes, w_{pr} for the probability of the minimum outcomes, and w_{max} for the maximum outcomes. It is further assumed that the three weights sum up to a value of 1, so that when setting the value for the first two weights, the value of the third weight results (i.e., $w_{max} = 1 - w_{pr} - w_{min}$) providing two free parameters. The probability of selecting a specific aspect is determined by the ratio of the specific weight to the sum of the importance weights of all aspects not considered so far. For instance, the probability of considering the minimum outcome first is determined by $w_{min}/(w_{min} + w_{pr} + w_{max})$. To implement a deterministic order in which the decision maker considers the different aspects, a final restriction is made (contrary to Bergert & Nosofsky, 2007). It is assumed that all weights have to be different and can take only the three values 0.999, 0.999×10^{-3}, and 10^{-6}. The parameter space for the three weights w_{min}, w_{pr}, and w_{min} is therefore restricted to the set $\{(0.999, 0.999 \times 10^{-3}, 10^{-6}), (0.999, 10^{-6}, 0.999 \times 10^{-3}), (0.999 \times 10^{-3}, 0.999, 10^{-6}), (0.999 \times 10^{-3}, 10^{-6}, 0.999), (10^{-6}, 0.999, 0.999 \times 10^{-3})$, and $(10^{-6}, 0.999 \times 10^{-3}, 0.999)\}$. With this discrete distribution of possible weights, the priority model essentially considers the different aspects of gambles deterministically. Whenever an aspect is considered last, the difference is not compared with the specific threshold for that aspect but with a threshold value of zero. For instance, when the maximum outcomes are considered last, the difference d_{max} is compared with a value of zero, so that when d_{max} is larger than zero, Gamble A is chosen. The probability of choosing Gamble A over Gamble B is then defined as $p_{max}(A,B) = p[d_{s\_max} > 0) = F[d_{max}/\sigma_{max}]$ where F is the standard normal cumulative distribution function.

Finally, to determine the total probability of choosing Gamble A over Gamble B, the choice probabilities resulting from the sequential comparison processes need to be combined. When the gambles are first compared by their minimum outcomes, followed by their minimum probabilities, and finally by their maximum outcomes, then the probability of choosing Gamble A over Gamble B is defined as

$$p(\text{A,B}) = p_{\min}(\text{A,B}) + [1 - p_{\min}(\text{A,B}) - p_{\min}(\text{B,A})] \times p_{\text{pr}}(\text{A,B}) + [1 - p_{\min}(\text{A,B}) - p_{\min}(\text{B,A})] \times [1 - p_{\text{pr}}(\text{A,B}) - p_{\text{pr}}(\text{B,A})] \times p_{\max}(\text{A,B}). \quad \text{(A4)}$$

When aspects of the gambles are compared in a different order, Equation A4 needs to be modified accordingly. For the sake of parsimony, it is assumed in the following that the standard deviation of the error components is determined by one single error parameter sigma (i.e., σ), so that $\sigma_{\text{pr}} = \sigma$, and $\sigma_{\min} = \sigma_{\max} = \sigma \times \max(\text{payoff}(i))$, where payoff($i$) refers to the possible payoffs of all gambles presented in each individual study. In sum, the priority model has five free parameters: sigma, specifying the standard deviation of the error component (with $\sigma > 0$); the outcome threshold δ^*_{out} ($0 < \delta^*_{\text{out}} < 1$), the probability threshold δ_{pr} ($0 < \delta_{\text{pr}} < 1$), and the two importance weights $w_{\min}$ and w_{pr} that determine the order in which the aspects are considered.

Appendix B

Specification of Cumulative Prospect Theory

According to cumulative prospect theory (Tversky & Kahneman, 1992), the subjective value of a payoff x is defined as

$$v(x) = \begin{cases} x^{\alpha}, \text{ if } x \geq 0 \\ -\lambda(-x)^{\beta}, \text{ if } x < 0 \end{cases}, \quad \text{(B1)}$$

where α and β are free parameters that modulate the curvature of the subjective value function separately for gains and losses, with $\lambda \geq 1$ specifying loss aversion. The probabilities p_x of the payoff x are transformed by the following weighting functions:

$$w(p_x, c) = p_x^{\,c}/[p_x^{\,c} + (1 - p_x)^c]^{1/c}, \quad \text{(B2)}$$

with $c = \gamma$ for positive payoffs and $c = \delta$ for negative payoffs, which are two free parameters that specify the inverse S-shaped transformation. This transformation is used to assign the probability of every payoff with a specific weight π:

$$\pi(p_x) = \begin{cases} \sum_{x_i \geq x} w(p_{x_i}, \gamma) - \sum_{x_i > x} w(p_{x_i}, \gamma), \text{ if } x \geq 0 \\ \sum_{x_i \leq x} w(p_{x_i}, \delta) - \sum_{x_i < x} w(p_{x_i}, \delta), \text{ if } x < 0 \end{cases}, \quad \text{(B3)}$$

where x_i refers to all possible payoffs that are at least as good or better than the considered payoff x or to all possible payoffs that are at least as bad or worse than the considered outcome x, respectively. Finally, an overall subjective value for a Gamble A can be determined by

$$V(\text{A}) = \sum_{x_i} \pi(p_{x_i}) v(x_i), \quad \text{(B4)}$$

and the gamble with the larger value is preferred by the decision maker.

Appendix C

Specification of Decision Field Theory

Decision field theory (DFT) assumes that during the decision process a difference score is determined by integrating comparisons of the alternatives' aspects. Mathematically, the difference score can be determined by

$$d_{\text{DFT}} = v(\text{A}) - v(\text{B}), \quad \text{(C1)}$$

where v represents the subjective value (i.e., the valence) of Gambles A and B. The valence v of each gamble is defined as

$$v(\text{A}) = \sum_i W(p_{i_\text{A}}) u(x_{i_\text{A}}), \quad \text{(C2)}$$

with W being continuous random variables that represent the subjective weight assigned to the different outcomes and that can vary from time to time. The subjective value of the various outcomes x_i are determined by the utility function $u(x_i)$. For parsimony, I assume that the expected value of each attention weight equals the probability with which the associated outcome occurs, that is, $\text{EV}[W(p_i)] = p_i$. Furthermore, I assume that the subjective value of an outcome is identical with the monetary outcome, that is $u(x_i) = x_i$.

DFT assumes that the difference score is determined by a dynamic process in which the differences of the gambles' outcomes are integrated over time. Whenever this difference score passes the decision threshold θ_{DFT} a decision is made. The probability of passing the threshold is also a function of the standard deviation of the differences (i.e., σ_{DFT}). The variance of the difference is conventionally defined as

$$\sigma^2_{\text{DFT}} = \sigma^2_{\text{A}} + \sigma^2_{\text{B}} - 2 \times \sigma_{\text{AB}}, \quad \text{(C3)}$$

where σ^2_{A} and σ^2_{B} are the variances of the valences of each gamble, and σ_{AB} denotes the covariance between the two valences (which is zero for independent gambles).

Received February 23, 2007
Revision received May 27, 2008
Accepted June 11, 2008 ■