American Psychological Association

Decision

2325-9965/19/$12.00

2019, Vol. 6, No. 3, 223–236
http://dx.doi.org/10.1037/dec0000098

Constructing Preference From Sequential Samples: The Impact of Evaluation Format on Risk Attitudes

Yonatan Vanunu
University of New South Wales

Thorsten Pachur
Max Planck Institute for Human Development, Berlin, Germany

Marius Usher
Tel-Aviv University

The cognitive mechanism underlying decisions based on sequential samples has been found to be affected by whether multiple alternatives are evaluated together or whether each alternative is evaluated individually. In this experiment, we examined whether evaluation format can also lead to different preference orders among risky alternatives. We hypothesized that because of differences in computational demands posed by the 2 evaluation formats, there would be differences in the type of the cognitive mechanism deployed: a risk-return mechanism, that trades off the mean reward and risk of an alternative, or a selective-accumulator mechanism, that sums the rewards of each alternative, with a higher weight to more extreme payoffs. Each participant rated the same set of alternatives (sequences of payoffs from slot machines) in both a one-by-one and a grouped evaluation format. The mean and the variance of the payoff distributions of each alternative were varied orthogonally. As predicted, in the grouped (but not in the one-by-one) condition, the impact of the variance on participants' ratings interacted with the mean payoff. Specifically, participants were risk averse for alternatives with a low mean payoff and risk seeking for alternatives with a high mean payoff. Computational modeling showed that the majority of participants were best described by a risk–return model in the one-by-one condition but by a selective-accumulator model in the grouped condition. Our results underline the importance of studying the cognitive foundations of risk attitudes in order to understand how they are shaped by the structure of a given decision task.

Keywords: preference construction, risk attitudes, value integration, evaluation format, computational modeling

Supplemental materials: http://dx.doi.org/10.1037/dec0000098.supp

This article was published Online First July 30, 2018.
Yonatan Vanunu, School of Psychology, University of New South Wales; Thorsten Pachur, Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany; Marius Usher, School of Psychological Science, Tel-Aviv University.
This research was supported by a grant from the German Research Foundation (DFG) as part of the priority program on New Frameworks of Rationality (SPP 1516) to Thorsten Pachur (HE 2768/7-2) and a grant from the (NSF-BSFNIH) Computational Neuroscience (CRCNS) to Marius Usher. Thorsten Pachur and Marius Usher were joint senior authors.

We thank Michael Brusovansky for helpful discussion on the experimental design and Susannah Goss for editing the manuscript. This work was done while Yonatan Vanunu was enrolled in an master's program at Tel-Aviv University under the supervision of Marius Usher, and during an internship at Max Planck Institute for Human Development under the supervision of Thorsten Pachur.
Correspondence concerning this article should be addressed to Thorsten Pachur, Center for Adaptive Rationality, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. E-mail: Pachur@mpib-berlin.mpg.de

Judgments and decisions are often formed on the basis of sequences of observed values (or samples) from available options. For example, stocks are evaluated on the basis of sequences of returns (Betsch, Plessner, Schwieren, & Gütig, 2001), and people form impressions of others based on samples of observed behaviors or described traits (Anderson, 1981). Furthermore, a sequential sampling of values (either internally, from the mental representation of the alternatives, or externally, based on eye fixations) and an integration of these values forming a dynamic preference state has been proposed to underlie multiattribute decisions, value-based decisions, and decision making under risk (e.g., Bhatia, 2013; Busemeyer, 1985; Busemeyer & Townsend, 1993; Krajbich, Armel, & Rangel, 2010; Pleskac, Yu, Hopwood, & Liu, in press; Tsetsos, Usher, & Chater, 2010; Usher & McClelland, 2004).

The majority of studies on decision making based on sequences of values have used paradigms in which people are presented with multiple alternatives simultaneously (i.e., in groups) and are shown samples from each of them. For instance, people are shown rapid numerical streams of values (which are sampled from the alternatives' payoff distributions) and are asked to choose between them (Tsetsos, Chater, & Usher, 2012; Zeigenfuse, Pleskac, & Liu, 2014) or to rate them on a continuous scale (Betsch, Kaufmann, Lindow, Plessner, & Hoffmann, 2006; Betsch et al., 2001). However, recent studies have also started to use paradigms in which options are evaluated individually (e.g., Ashby & Rakow, 2014; Golan & Ert, 2015; Pachur & Scheibehenne, 2012; see also Brusovansky, Vanunu, & Usher, 2017; Pachur & Scheibehenne, 2017).

Given this methodological variability, one may ask whether the cognitive mechanism of preference construction from sequential samples might depend on how the alternatives are evaluated. Specifically, does it matter whether the alternatives are presented and evaluated in a *grouped* or in a *one-by-one* format? It has been shown that people show different risk preferences as a function of how many alternatives they can choose and sample from (Hills, Noguchi, & Gibbert, 2013; Noguchi & Hills, 2016; see also Tsetsos et al., 2012). It is currently unclear, however, whether an effect of evaluation format holds more generally and how its impact on the mechanism underlying risky decision making might come about.

Based on recent findings by Brusovansky et al. (2017), we propose that, because grouped and one-by-one evaluation formats pose different computational demands, decision makers might rely on different cognitive mechanisms in these two evaluation formats. Brusovansky et al. asked participants to evaluate alternatives whose payoff distribution differed in the mean payoff as well as in the sum of the sampled payoffs. It turned out that participants tended to rely on a mechanism that evaluated alternatives based on their mean payoff in one-by-one evaluation, but on a mechanism that evaluated alternatives based on the sum of the observed payoffs in grouped evaluation. Here, we extend these findings to evaluations of alternatives that differ in risk (i.e., the variance of the payoff distribution), thereby tapping into risk attitudes. Importantly, as we demonstrate in the section "Extending Summation and Averaging Models to Risk Preference", the reliance on a mechanism that relies on summary statistics of the alternatives (i.e., mean and variance) or a mechanism that relies on a (weighted) sum of the sampled payoffs can lead to systematic differences in risk attitude. We will argue that these two mechanisms differ in their robustness in the face of computational demands posed by the grouped (relative to the one-by-one) evaluation formats; as a consequence, the two evaluation formats should foster systematically different patterns of risk preference (i.e., risk aversion or risk seeking).

In what follows, we first review previous results that indicate an impact of evaluation format on the cognitive mechanism triggered; then, we present computational modeling analyses showing that decision makers in grouped evaluation are predicted to demonstrate risk-seeking behavior for options with a high average payoff and risk aversion for options with a low average payoff. By contrast, in one-by-one evaluation, the risk preference of decision makers are predicted to be independent of whether the average outcome payoff is high or low. We then present an experiment that tests these predictions.

The Influence of Evaluation Format on Decision Making

A prominent theoretical notion of how alternatives are evaluated based on sequences of samples is that values (or value differences) of the alternatives are integrated in value accumulators toward a decision boundary (Busemeyer, 1985; Busemeyer & Townsend, 1993; Krajbich et al., 2010; Ratcliff, Smith, Brown, & McKoon, 2016; Roe, Busemeyer, & Townsend, 2001).[1] An important distinction in models of value integration is whether the evaluation of an alternative is based on the summed observed payoffs or the average payoff (Betsch et al., 2001, 2006; note that in studies that do not manipulate the number of samples, this distinction is irrelevant, as both types of integration make the same prediction). Although some studies have reported support for averaging (Anderson, 1981; Kahneman, Fredrickson, Schreiber, & Redelmeier, 1993), others have supported a summation principle (Betsch et al., 2001, 2006).

Brusovansky et al. (2017) found evidence that the reliance on summation versus averaging depends on whether the alternatives are presented in a grouped or a one-by-one evaluation format. In order to illustrate the authors' approach, consider a choice between Alternative A, consisting of six high payoffs, and Alternative B, consisting of nine payoffs; six of the payoffs of Alternative B are the same as those of Alternative A, and the other three are lower payoffs but still positive and higher than the average values sampled in the experiment. It was found that whereas an averaging mechanism seemed to dominate preferences in the one-by-one format, such that participants tended to prefer Alternative A (characterized by a high average but a low number of samples), a summation mechanism seemed to dominate in the grouped format, such that participants tended to prefer Alternative B (characterized by a low average but a higher number of samples; Figure 1, right panel; Experiment 4B). Furthermore, when the average payoff and the number of samples were orthogonally manipulated, the number of samples affected only the preference in the grouped but not in the one-by-one condition (Figure 1, left panel; Experiment 4A). This suggests that a summation mechanism was more likely to operate in the former than in the latter.

Extending Summation and Averaging Models to Risk Preference

For preference construction between risky alternatives, a mechanism similar to the summary-statistics (or averaging) mechanism discussed by Brusovansky et al. (2017) is the prominent risk–return model, according to which the attractiveness of a risky alternative is an additive function of the alternative's average reward and its risk—that is, between the mean and the variance of the payoff distribution (Markowitz, 1952; see also Weber, 2010)[2]:

$$V(A) = a \times M(A) + b \times SD(A) + c, \quad (1)$$

where $V(A)$ is the subjective valuation of an alternative A, represented by a sequence of payoffs, $M(A)$ is the mean payoff, and $SD(A)$ is the standard deviation (i.e., risk) of the sequence of payoffs. The parameters a and b represent the weight given to the mean payoff and the risk, respectively, such that the ratio b/a reflects the overall risk tendency of the decision maker. A negative value of b indicates risk aversion and a positive value indicates risk seeking; c is an intercept (which is needed to map an internal preference onto an external scale). Note that the risk–return model represents what would be expected under normative considerations.

The risk–return model implicitly assumes that the decision maker is able to generate an estimate of the mean of a sequence of values and its variance. Is this a realistic assumption? Several studies have documented that people are able to extract the mean of a set of briefly presented perceptual properties in a variety of tasks, ranging from low-level properties such as circle size (Ariely, 2001) to high-level ones

[1] In binary decisions, the most known type of sequential sampling model relies on a single accumulator that integrates the total evidence difference (of log-likelihoods; Gold & Shadlen, 2001).

[2] An alternative to this risk–return model is an expected utility model, in which risk attitude is modeled via a concave value function. We focus here on the risk–return model because, although qualitatively similar in its predictions, this model provided a better quantitative account of our data.

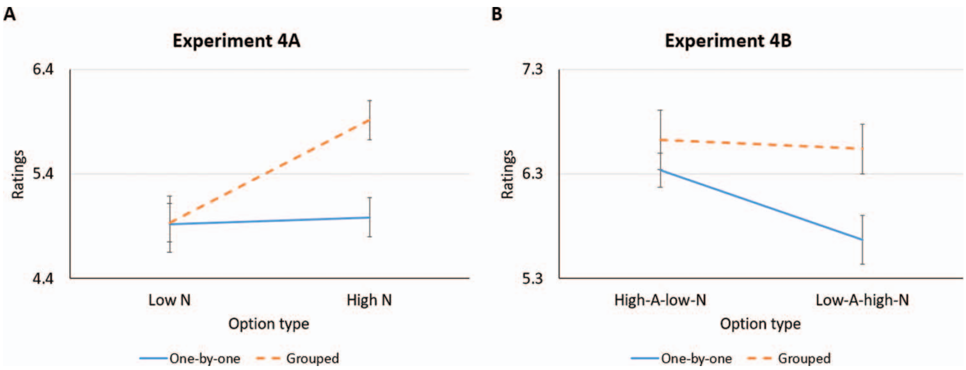

Figure 1. Results from Brusovansky, Vanunu, and Usher (2017), showing an effect of evaluation format on ratings of alternatives that consist of sequences of payoffs of slot machines (left panel) or sequences of lecturer ratings (right panel). In Experiment 4A, the average and number of samples were manipulated orthogonally (the ratings are collapsed across different levels of average the alternatives with higher averages were evaluated more positively; not shown). In Experiment 4B, the summation and averaging principles are set in opposition (among filler trials; see Brusovansky et al. for details; reproduced with permission). Figure adapted from "Why we should quit while we're ahead: When do averages matter more than sums," by M. Brusovansky, Y. Vanunu, and M. Usher, 2017, Decision. Copyright 2017 by the American Psychological Association. See the online article for the color version of this figure.

such as emotional expressions (Haberman & Whitney, 2007). In the domain of numerical cognition, people have also been shown to be able to accurately estimate the mean of rapid sequences of two- or three-digit numbers (Brezis, Bronfman, Jacoby, Lavidor, & Usher, 2016; Brezis, Bronfman, & Usher, 2015; Brezis, Bronfman, & Usher, 2018; Malmi & Samson, 1983). Brezis et al. (2016) proposed that information about an observed distribution of values might be encoded by numerosity detectors in the parietal system, and that these numerosity detectors allow one to extract the average of an observed sequence via neural population averaging (see also Malmi & Samson, 1983, for a similar suggestion).

Figure 2 (top panel) shows the predicted evaluations of the risk–return model. To derive the predictions, we first drew eight samples from each of the machines' payoff distributions, which consisted of Gaussian distributions with either low or high mean payoff (40 vs. 60) and a low or high variance (standard deviation: 5 vs. 15). The mean payoff and the payoff variance in these samples were then fed into Equation 1; the model parameters were set to $a = 1$, $b = +1$ and -1, respectively, and $c = 0$.[3] As can be seen, the model predicts independent effects of the mean (return) and the variance (risk) of an alternative on preferences without an interaction; consequently, whether the low- or high-risk alternative is judged as more attractive (i.e., risk aversion or risk seeking) depends only on the decision maker's risk attitude (i.e., whether parameter b in Equation 1 is positive or negative).

An implicit assumption of the risk–return model is that for each alternative, a representation of the sampled payoff distribution is maintained during the presentation of the sequence, from which the mean and variances of the sequence is estimated after all samples have been presented (Brezis et al., 2015, 2018; Malmi & Samson, 1983). Although this may be (relatively) feasible to implement cognitively in a one-by-one evaluation format, it poses a representational challenge in a grouped evaluation format, in which multiple alternatives are presented simultaneously. Unlike in the one-by-one condition, here, participants have to update

[3] Given that the weight given to risk, that is, risk preference, follows from to the ratio for the b and a parameters (i.e., b/a), for simplicity we set parameter $a = 1$ in this demonstration.

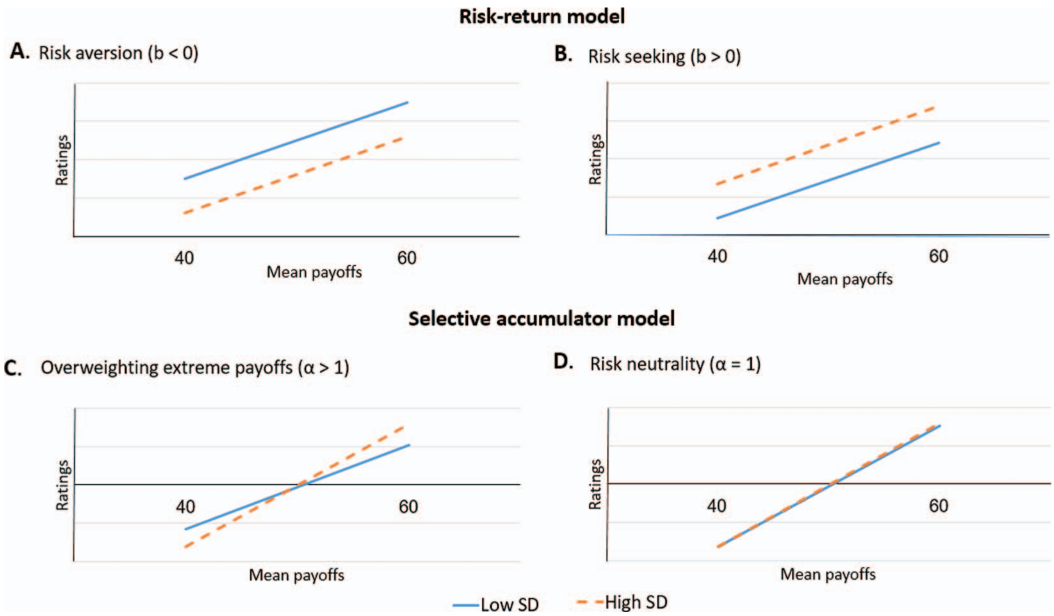

Figure 2. Top panel: Predicted evaluation of the risk–return model for a risk-averse (A, $b = 1$) and a risk-seeking (B, $b = -1$) decision makers; the other parameters in Equation 1 were set to $a = 1$ and $c = 0$. Bottom panel: Predicted evaluation of the selective-accumulator model. The predictions shown assume a value of the selectivity parameter α (see Equation 2) of 1 (Panel D) or 1.5 (Panel C); the other parameters in Equation 2 were set to $a = 1$, $\lambda = 1$, and $b = 0$. See the online article for the color version of this figure.

the representations of the different moments of the payoff distributions (i.e., mean and variance) of several alternatives without confusing them, thus requiring source monitoring (Johnson, Hashtroudi, & Lindsay, 1993). Errors in this process will make the evaluations noisier; as a consequence (all other things being equal), in a grouped format, this mechanism will yield rather similar evaluations for the different alternatives and thus decrease performance (see the online supplemental materials for a demonstration of this effect). Given that decision makers are often adaptive in their selection of decision strategies under cognitive constraints (Ariely & Zakay, 2001; Broadbent, 1971; Oh et al., 2016; Payne, Bettman, & Johnson, 1988; Rieskamp & Hoffrage, 2008; Zakay, 1993), we therefore predict that in a grouped format, people might switch to a mechanism that does not require a constant updating of both the means of the variances of multiple representations and that thus helps maintain performance.

What might such a mechanism be? Recent studies found evidence that in grouped evaluation people tend to rely on a mechanism that keeps a tally of sampled outcomes for each alternative and is subject to a weighting scheme that prioritizes extreme values (in particular high ones, but sometimes also low ones) over intermediate ones (e.g., Spitzer, Waschke, & Summerfield, 2017; Tsetsos et al., 2016; Zeigenfuse et al., 2014).[4] Here, we use such a mechanism, which we refer to as the *selective-accumulator model* (see also Spitzer et al., 2017). According to the model, the valuation of an alternative A is determined as follows:

[4] Here, we focus on a paradigm of experience-based decisions in which people are presented with a predetermined number of samples for each alternative, in contrast to the sampling paradigm (in which people also have to decide when to stop sampling) or the feedback paradigm (in which a decision is required after each value sample; e.g., Hertwig & Erev, 2009). We relate our findings to these other paradigms of experience-based decision in the Discussion section.

$$V(A) = \sum_{i=1}^{N} \lambda^{N-i} Sign(X_i - R) \times \left(|X_i - R|\right)^{\alpha} \times a + b, \quad (2)$$

where $|X_i - R|$ represents the (absolute) deviation of the magnitude of a payoff in the sequence from a reference value R, $Sign()$ is the sign function (i.e., -1 if $x < 0$, and 1 if $x > 0$) of the difference $(X_i - R)$, and N is the number of payoffs in the sequence. The summation of the sampled payoffs is subject to an attentional enhancement, implemented by the exponent α (≥ 1); with $\alpha > 1$, stronger deviations (i.e., extreme values) are weighted more heavily. The parameter λ (<1) represents temporal decay of the payoffs in the sequence, affecting earlier payoffs in the sequence more strongly than later ones (lower is have a larger exponent than higher is); λ thus implements a recency effect. The parameters a and b are scaling parameters, that map the prediction of the model on a response scale. Note that in contrast to the risk–return model, the selective-accumulator model only requires one tally per alternative and does need to maintain representations of the different moments (mean and variance) of the payoff distribution of each alternative. Moreover, it has been shown that a selective overweighting of extreme values (as implemented by the parameter α) can be adaptive in choice tasks by rendering the choice more robust to noise (Spitzer et al., 2017; Tsetsos et al., 2016). A selective-accumulator model might thus represent a mechanism that people turn to more in a grouped evaluation format.

Importantly, as shown in Figure 2 (bottom panel), in contrast to the risk–return model, the selective-accumulator model predicts an interaction between the mean and the variance of the observed payoffs. When an alternative with a high mean payoff is evaluated, the high values from the high-variance alternative receive more attention, increasing the attractiveness of that alternative and thus leading to risk seeking; when an alternative with a low mean payoff is evaluated, by contrast, the low values from the high-variance alternative receive more attention, decreasing the attractiveness of that alternative and thus leading to risk aversion.

The risk–return model and the selective-accumulator model are therefore associated with distinct risk-attitude signatures: Whereas the former predicts main effects of the mean and the variance of the sequence on preference, the latter predicts an interaction (see Figure 2C). If people are more likely to rely on the risk–return model in one-by-one evaluation, their risk attitude is predicted to be independent of whether the mean payoff is high or low. If people are more likely to rely on the selective-accumulator model in grouped evaluation, they are expected to be risk seeking for alternatives with a high mean payoff and risk averse for alternatives with a low mean payoff.[5] Such an interaction can also result from only some of the people deploying different mechanisms in the two conditions.

The aim of the following experiment was to test the predicted pattern of risk preferences across the one-by-one and grouped evaluation formats.

Method

Participants

Thirty students (16 females; age = 19–33 years, M = 25.67) participated in exchange for a participation fee of €20 and a performance-contingent bonus (see Procedure section for details). The participants were recruited from the participant pool maintained by the Max Planck Institute for Human Development, Berlin. Germany. The sample size was determined in advance.

Materials

The experiment had a 2 (evaluation format: one-by-one vs. grouped) × 2 (mean payoff of the alternative: high vs. low) × 2 (variance of the payoff sequence: high vs. low) within-subjects design. The alternatives were presented as slot machines, each presenting sequences of eight numerical values (range = 1–99) that represented possible payoffs (in €). The payoffs in each set were samples drawn from one of four Gaussian distributions that differed in terms of the mean (40 vs. 60) and the variance (standard deviation 5 vs. 15). Figure 3 provides a visual representation of

[5] The results of a preliminary study (Vanunu, 2016), showing this triple interaction, are reported in the online supplemental materials. The specific form of the selective-accumulator model (that gives higher weights to extreme values) used in our analysis was motivated by this result.

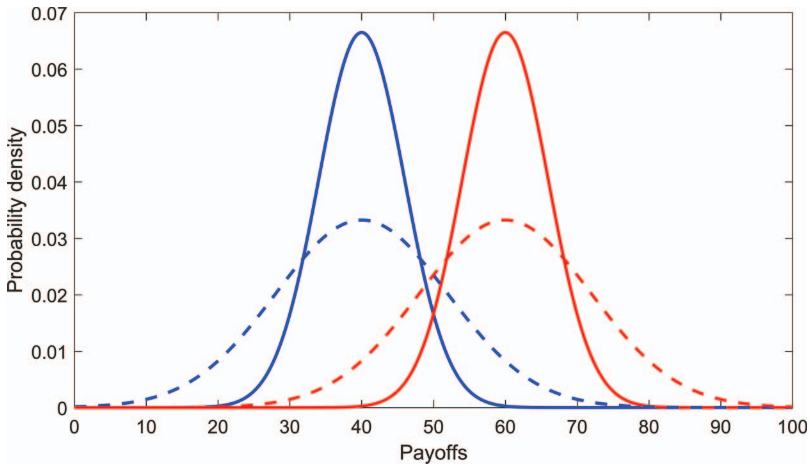

Figure 3. Payoff distributions for the four alternatives with a mean payoff of 40 (low) or 60 (high) and a standard deviation of 5 (low variance/risk) or 15 (high variance/risk), respectively. See the online article for the color version of this figure.

the four payoff distributions. No other measures and conditions than those reported here were collected.

Procedure

The one-by-one and the grouped evaluation conditions were implemented in two separate sessions, which were spaced 1 week apart and counterbalanced between participants. In both conditions, participants were told that they would see slot machines, each of which would show random samples from their payoff distribution. Having seen all samples, participants were asked to indicate how much they liked each machine on a continuous scale from 0 to 10 (0 = *I do not like the machine at all*; 10 = *I like the machine very much*), assuming that they would be able to play a single round on it. To encourage participants to express their genuine preferences, they were informed that 10 machines would be randomly picked at the end of each session, and that, of those, the machine they had given the highest evaluation would be played out to generate a single win. They received 10% of the payoff from each session (in €) as a bonus (an average of €6 per session). The order of the sessions (one-by-one vs. grouped) was counterbalanced across participants.

In the one-by-one evaluation condition, the machines were presented separately, one per trial, as shown in Figure 4A. On each trial, eight draws from the respective payoff distribution were displayed sequentially at the center of the screen, with a presentation rate of 1s per draw. After each trial, participants indicated their ratings for the machine. The session started with a short practice phase, in which eight different machines—two from each condition—were presented in random order, providing an impression of the possible outcomes of the slot machines in the test phase. In the first four practice trials, the samples of the machines were simply shown to participants without requiring a response; in the subsequent four practice trials, participants also evaluated the machines. The test phase consisted of 96 trials (24 machines for each of the four types of machines). The presentation order of the machines was randomized.

In the grouped evaluation condition, participants were told that they would see different rooms, each containing four slot machines (see Figure 4B). To help participants distinguish between the four machines, we presented each machine in one corner of the screen and gave each a unique color (red, green, blue, or yellow; randomly allocated for each room). The four machines in each room corresponded to the four alternatives resulting from crossing high and low mean payoff with high and low variance of the payoff distribution (see Figure 3). Draws were shown at a presentation rate of 1s per draw: one draw from each machine's payoff distribution in

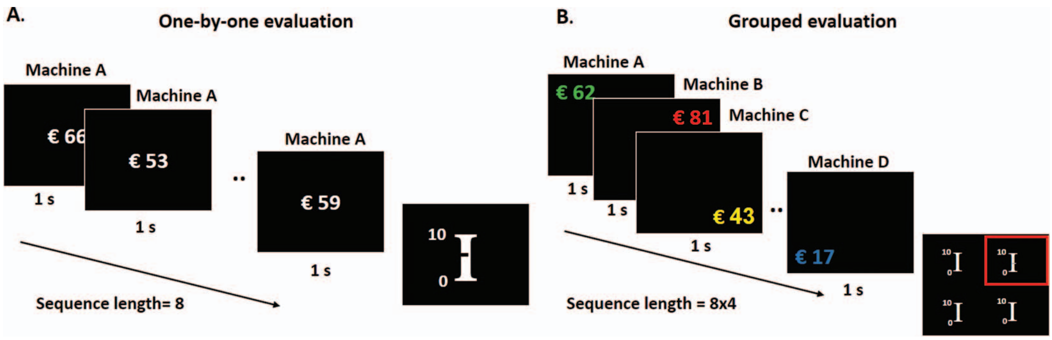

Figure 4. One-by-one (Panel A) and grouped (Panel B) evaluation conditions, consisting of a presentation of numerical sequences followed by an evaluation scale. In Panel A, the sequences were presented in the middle of the screen; in Panel B, they were presented in the four corners (each characterized by a different color) in clockwise circular fashion. See the online article for the color version of this figure.

clockwise circular fashion. This procedure was repeated for a total of eight rounds, such that eight draws were shown for each alternative (as in the one-by-one evaluation condition). Each room thus presented a total of 32 (8×4) payoffs. After the presentation of all payoffs, a scale from 0 to 10 was displayed in each of the four quadrants. Participants were asked to indicate how much they liked each machine using the mouse cursor. A rectangle around each scale (colored either red, green, blue, or yellow) indicated which scale the mouse was currently operating. The order in which the machines were evaluated was not constrained. The session started with a practice phase of two rooms (i.e., eight machines). The test phase consisted of 96 trials (24 machines for each of the four combinations of mean and variance). The presentation orders of the machines and their locations in each room were randomized.

Results

We first analyzed participants' ratings of the machines using mixed-effects models with mean payoff (low vs. high), variance (low vs. high), and evaluation format (one-by-one vs. grouped) as fixed effects and participants as random effect. We applied chi-square tests to test for differences between the log-likelihood value of the full factorial mixed model and the log-likelihood values of various null models (in which the respective factor or interaction was excluded; for the full factorial coefficient tables, see Tables S1 and S2 in the online supplemental materials). The analysis showed a main effect of mean payoff, $\chi^2(1) = 1{,}634.52$, $p < .0001$, indicating that machines with a high mean payoff were evaluated as more attractive than machines with a low mean payoff ($M_{\text{high}} = 7.11$, $SD = 0.19$ vs. $M_{\text{low}} = 3.47$, $SD = 0.21$). Critically, there was, as predicted, a triple interaction between evaluation format, mean payoff, and variance, $\chi^2(1) = 6.15$, $p < .05$ (see Figure 5).

To unpack the triple interaction, we conducted follow-up analyses for the two evaluation conditions separately. For the grouped evaluation condition, we obtained, in addition to a main effect of mean payoff, $\chi^2(1) = 1{,}176.68$, $p < .0001$ ($M_{\text{high}} = 7.47$, $SD = 0.21$ vs. $M_{\text{low}} = 3.84$, $SD = 0.20$), an interaction between mean payoff and variance, $\chi^2(1) = 35.60$, $p < .0001$ (see Figure 5, right panel): Among alternatives with a high mean payoff, participants evaluated those with high variance as more attractive, $\chi^2(1) = 32.1$, $p < .0001$ (i.e., indicating risk seeking), whereas among alternatives with a low mean payoff, participants evaluated those with low variance as more attractive, $\chi^2(1) = 14.1$, $p < .0005$ (i.e., indicating risk aversion). In the one-by-one evaluation condition, there was, in addition to a main effect of average payoff, $\chi^2(1) = 1{,}834.12$, $p < .0001$ ($M_{\text{high}} = 6.75$, $SD = 0.20$ vs. $M_{\text{low}} = 3.11$, $SD = 0.25$), a main effect of variance, $\chi^2(1) = 44.75$, $p < .0001$ ($M_{\text{lowV}} = 5.08$, $SD = 0.20$ vs. $M_{\text{highV}} = 4.78$, $SD = 0.19$): Participants evaluated alternatives with low variance as more attractive (indicating risk aversion), irrespective of whether

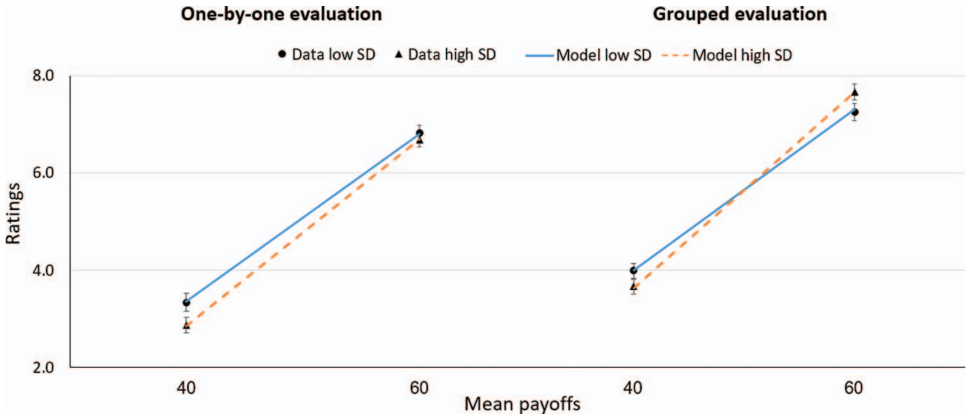

Figure 5. Average ratings for each of the four alternatives (triangles: high variance; squares: low variance) for the two evaluation conditions; error bars correspond to within-subjects standard errors. Blue/red lines show the predicted ratings, using the model that best accounted for each participant's responses and the best-fitting parameters (see Model Comparison section). See the online article for the color version of this figure.

they had a high or low mean payoff (high: $\chi^2[1] = 69.57$, $p < .0001$; low: $\chi^2[1] = 4.75$, $p < .05$; see Figure 5, left panel). Although, unlike in the grouped condition, the pattern of risk preference did not reverse between the two payoff conditions, there was an interaction between mean payoff and variance, $\chi^2(1) = 11.45$, $p < .001$ (reflecting that the differences in rated attractiveness between the low- and high-variance options were smaller in the high- than in the low-payoff condition).

Second, we examined the ratings for temporal effects, which might result from the sequential presentation of the samples. Specifically, we used linear regression to test to what extent the weight that each of the eight samples for each alternative received for the overall evaluation of the alternative depended on the sample's position in the sampling sequence. Then, we compared the resulting regression weights for each position using a repeated-measures analysis, separately for each session. As illustrated in Figure 6, there was a recency effect (i.e., higher β values for more recent samples) in the grouped condition, $F(7, 30) = 3.74$, $p < .005$, indicating that payoffs presented later in the sequence received a higher weight. In the one-by-one condition, by contrast, there were no effects of temporal position, $F(7, 30) = 0.523$, $p = .762$.

Despite the fact that participants' risk-preference patterns in the grouped and one-by-one evaluation conditions were in line with our predictions, the presence of a double interaction in both conditions (albeit of different magnitudes) indicates that the association between evaluation format and the underlying mechanism was not as clear-cut as in the idealized case depicted in Figure 2. In order to better understand this pattern and to test for individual differences between participants, we used computational modeling to compare the ability of the risk–return model and the selective-accumulator model to account for individual participants' ratings in the two evaluation format conditions.

Model Comparison

We applied both the risk–return model (Equation 1) and the selective-accumulator model (Equation 2) to each participant's ratings.[6] For simplicity, the reference parameter R for the selective-accumulator model was fixed to the middle of the range of the displayed payoffs (i.e., 50). The log-likelihoods of the models were determined using regression anal-

[6] As an alternative to the risk–return model, we also examined an EU model. This model is based on the integration of values transformed via a concave utility function and also predicts a main effect of variance on preference. It showed a worse model fit in both conditions. In addition, we examined a variant of the risk–return model, using the coefficient of variation (CV; defined as SD/M) as a measure of risk. This model also showed a worse fit than the risk–return model (Equation 1; see the online supplemental materials for details).

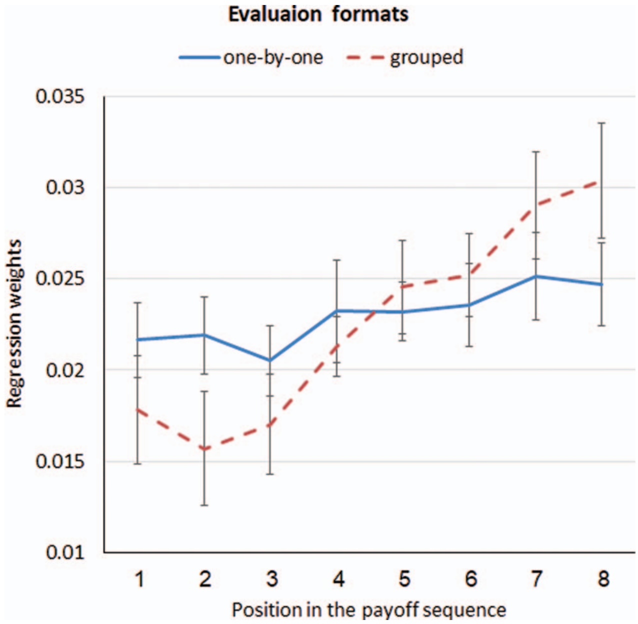

Figure 6. Temporal weights for one-by-one and grouped conditions, calculated by a linear regression analysis with the sequential payoffs as predictors for participants' evaluations. Error bars represent within-subjects standard errors. See the online article for the color version of this figure.

yses. Specifically, the payoffs for each sequence were inserted into the model equations. For each trial and each set of linear model parameters (i.e., *a*, *b*, *c* in Equation 1, and *a*, *b* in Equation 2), the model's predicted rating for each sequence was compared with each participant's actual response. Linear regression (assuming normally distributed noise) was then used to obtain the set of best-fitting (linear) parameters and the log-likelihood of the models. As the α and λ parameters of the selective-accumulator model are not part of a linear model equation, they were estimated prior to the regression estimation using a grid search (in the range [1, 2] for α and [0, 1] for λ, and with a step size of .02). Furthermore, as the risk–return model and the selective-accumulator model differ in the number of free parameters (3 vs. 4), we compared them in terms of the aggregate Akaike information criterion (AIC; Akaike, 1973; see also Pachur & Scheibehenne, 2017), with lower values indicating a better model fit. In an additional model recovery study, we showed that this model selection approach correctly identified the model from which the data were generated (see the online supplemental materials for more details).[7]

Table 1 shows the median and standard deviation (across participants) of the estimated model parameters, separately for the two evaluation conditions, as well as the aggregate AIC scores for each of the two models. For the grouped condition, the selective-accumulator model had a lower aggregate AIC score than the risk–return model; for the one-by-one condition, by contrast, the risk–return model had a lower aggregate AIC score than the selective-accumulator model. In addition, we classified the individual participants according to which of the two models provided a better fit. For the one-by-one condition, about two thirds of the participants were best described by the risk–

[7] In the model recovery study, we also applied a variant of the risk–return model that had no risk coefficient (i.e., a pure averaging model) to model the simulated data. Based on the AIC, both the actually generating model (namely, the full model, which includes a risk coefficient) and the data-generating parameters were accurately recovered.

Table 1
Results of the Computational Modeling Analysis With the Selective-Accumulator Model and the Risk–Return Model: Median (Across Participants) Best-Fitting Parameters and Standard Deviations and Summed (Across Participants) AIC Scores of Each Model, Separately for the Two Evaluation Conditions

		Evaluation condition					
		One-by-one			Grouped		
Model	Parameter	*Mdn*	*SD*	Aggregate AIC	*Mdn*	*SD*	Aggregate AIC
Selective-accumulator	α	1.02	.32	8,138	1.21	.33	10,075
	λ	.97	.05		.90	.11	
Risk–return	b/a	−.19	.41	7,958	.01	.33	10,157

Note. For the risk–return model, the ratio of the risk and return parameters (i.e., b/a) are shown, with values <1 indicating that more weight is given to return than to risk information. AIC = Akaike information criterion.

return model (67% vs. 33% for the risk–return model and the selective-accumulator model, respectively); for the grouped condition, by contrast, about two thirds were best described by the selective-accumulator model (34% vs. 66% for the risk–return and the selective-accumulator model, respectively; Table S3 in the online supplemental materials).

Finally, in order to assess absolute model fit, we determined the predicted evaluation for the different alternatives, using the model that best accounted for the respective participant's responses. As shown in Figure 5 (blue/red lines), the models captured the observed data extremely well. Overall, these results support our hypothesis that whereas people follow a risk–return mechanism in the one-by-one condition, in the grouped evaluation condition, people switch to the computationally less demanding selective-accumulator mechanism.

Discussion

We argued that grouped versus one-by-one evaluation might differ in computational demands (e.g., source monitoring) and that people would therefore rely on different mechanisms depending on evaluation format (this was also suggested by the results of a preliminary study; see the online supplemental materials). As a consequence, given the contrasting risk-preference signatures for the risk–return and the selective-accumulator mechanisms (see Figure 2), we hypothesized that in the one-by-one condition risk preferences would be unaffected by the mean sampled payoff (risk–return model); in the grouped condition (selective-accumulator model), however, participants would show risk aversion for alternatives with a low mean payoff and risk seeking for alternatives with a high mean payoff. The results of our experiment confirmed this hypothesis. In addition, computational modeling showed stronger support for the risk–return model in the one-by-one condition, whereas in the grouped condition, the selective-accumulator model received the stronger support.

Note that these effects of evaluation format are different from findings that preference construction can be affected by elicitation method, such as rating and choice (for an overview, see Lichtenstein & Slovic, 2006), as in both the grouped and one-by-conditions, participants rated the alternatives. Moreover, the pattern of risk preferences observed in the grouped condition is opposite of the *payoff-variability effect*, which refers to the phenomenon that choices become more random (i.e., closer to 50%) the higher the risk (i.e., variability) of a risky alternative when it is compared with a safe alternative (e.g., Busemeyer, 1985). In our experiment, by contrast, the responses became more extreme with higher risk (i.e., alternatives with a low average payoff were rated as less attractive and alternatives with a high average payoff were rated as more attractive the higher the variance of the alternative). One possible reason for this apparent inconsistency is that whereas the payoff-variability effect has been observed in studies in which people chose between a safe and a risky alternative (Busemeyer, 1985; Barron & Erev, 2003), in our experiment, all of the alternatives were risky. Also, note that the pattern of risk preferences we observed in the grouped condition (i.e., risk seeking for alternatives with high mean payoffs and risk aversion for alternatives with low mean payoffs) and the payoff-variability effect seem to be driven by different mechanisms. As shown by Busemeyer (1985), the payoff-variability effect can be accounted for by a sequential sampling model.[8] Conversely, in our grouped condition, the data were best accounted for by the selective-accumulator model.

Given that the paradigm used in our experiment shares features with other paradigms in experience-based decision making (such as the sampling or the feedback paradigm; for an overview, see Hertwig & Erev, 2009), one might ask to what extent our findings generalize to these other paradigms. Note that some parallel observations have been made in these paradigms. For example, Hills and colleagues (Hills et al., 2013; Noguchi & Hills, 2016) found that increasing the number of alternatives among which people have to choose in the sampling paradigm was associated with an increase in risk seeking. The authors attributed this effect to differences in the amount of sampling (note that in the sampling paradigm, people determined themselves how many samples to take from each alternative); based on our results, however, it may also be possible that with a larger choice set, participants are more likely to rely on a selective-accumulator mechanism to evaluate the alternatives. Furthermore, some recent research using the sampling paradigm has contrasted preferences in a willingness-to-pay (WTP) task, in which individually presented alternatives are evaluated, and a choice task, in which people choose between multiple alternatives. Golan and Ert (2015) found, in line with our study, that the (one-by-one) WTP evaluation format resulted in preference patterns that were closer to a normative benchmark than the (grouped) choice format did, such that there was less underweighting of rare events and the responses followed the expected values of the alternatives more closely.

To the extent that our results generalize to other paradigms of experience-based decisions, they also predict interesting boundary conditions for previous findings. For instance, Ludvig and colleagues (e.g., Ludvig, Madan, & Spetch, 2014) found that in the context of a choice task presented in the feedback paradigm, more weight seemed to be given to the extreme values of alternatives' payoff distribution. Based on our results, one may predict that such evidence for an overweighting of extreme events will disappear (or be less pronounced) in the context of a task in which the options are rated individually.

Conversely, it should also be noted that in the sampling paradigm, choices are made after every sample and are followed by feedback, which is likely to engage other or additional psychological processes—such as reinforcement learning (e.g., Sutton & Barto, 1998) and prediction (e.g., Plonsky, Teodorescu, & Erev, 2015)—than in our sequential sampling paradigm. Whether and when the differences between our paradigm and the sampling and feedback paradigms limit the generality of the findings of our study to other types of decisions from experience is a fascinating issue for future research.

Although our results generally followed the pattern predicted in our model analysis (Figure 2), there were also considerable individual differences: About one third of the participants seemed to rely on the risk–return mechanism even in the grouped condition, and about one third seemed to rely on the selective-accumulator mechanism in the one-by-one condition. Further studies are required to better understand the observed individual differences. For example, such studies could correlate strategy use in the grouped conditions with the decision maker's cognitive capacity (e.g., working memory) or with the ability to maintain differentiated distributions of values in long-term memory (e.g., source memory). From an adaptivity point of view, one would predict that participants with low memory ability or differentiation capacity are more likely to rely on an accumulator mechanism than participants with high memory ability.

Our finding that the evaluation format affects preference construction has important implications for improving decision quality. To date, most decision-making studies that reported deviations from normative benchmarks have used choice tasks (in which multiple alternatives are presented in groups); few studies, by contrast, have involved one-by-one evaluations. If a one-by-one evaluation format is more likely to trigger population coding mechanisms that extract the normatively relevant mean and variance of a sequence of payoffs, decisions and evaluations

[8] In this model, the variability in the evaluation of the risky alternative therefore negatively impacts the probability of choosing it when its mean is higher than that of the certain one (and vice versa; Busemeyer, 1985).

should be closer to standards of rationality when alternatives are evaluated one at a time.

References

Akaike, H. (1973). Information theory as an extension of the maximum likelihood principle. In B. N. Petrov & F. Csaki (Eds.), *Second international symposium on information theory* (pp. 267–281). Budapest, Hungary: Akademiai Kiado.

Anderson, N. H. (1981). Integration theory applied to cognitive responses and attitudes. In R. E. Petty, T. M. Ostrom, & T. C. Brock (Eds.), *Cognitive responses in persuasion* (pp. 361–397). Hillsdale, NJ: Erlbaum.

Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science, 12,* 157–162. http://dx.doi.org/10.1111/1467-9280.00327

Ariely, D., & Zakay, D. (2001). A timely account of the role of duration in decision making. *Acta Psychologica, 108,* 187–207. http://dx.doi.org/10.1016/S0001-6918(01)00034-8

Ashby, N. J., & Rakow, T. (2014). Forgetting the past: Individual differences in recency in subjective valuations from experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 40,* 1153–1162. http://dx.doi.org/10.1037/a0036352

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16,* 215–233. http://dx.doi.org/10.1002/bdm.443

Betsch, T., Kaufmann, M., Lindow, F., Plessner, H., & Hoffmann, K. (2006). Different principles of information integration in implicit and explicit attitude formation. *European Journal of Social Psychology, 36,* 887–905. http://dx.doi.org/10.1002/ejsp.328

Betsch, T., Plessner, H., Schwieren, C., & Gütig, R. (2001). I like it but I don't know why: A value-account approach to implicit attitude formation. *Personality and Social Psychology Bulletin, 27,* 242–253. http://dx.doi.org/10.1177/0146167201272009

Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review, 120,* 522–543. http://dx.doi.org/10.1037/a0032457

Brezis, N., Bronfman, Z. Z., Jacoby, N., Lavidor, M., & Usher, M. (2016). Transcranial direct current stimulation over the parietal cortex improves approximate numerical averaging. *Journal of Cognitive Neuroscience, 28,* 1700–1713. http://dx.doi.org/10.1162/jocn_a_00991

Brezis, N., Bronfman, Z. Z., & Usher, M. (2015). Adaptive spontaneous transitions between two mechanisms of numerical averaging. *Scientific Reports, 5,* 10415. http://dx.doi.org/10.1038/srep10415

Brezis, N., Bronfman, Z. Z., & Usher, M. (2018). A perceptual-like population-coding mechanism of approximate numerical averaging. *Neural Computation, 30,* 428–446. http://dx.doi.org/10.1162/neco_a_01037

Broadbent, D. E. (1971). *Decision and stress*. Oxford, UK: Academic Press.

Brusovansky, M., Vanunu, Y., & Usher, M. (2017). Why we should quit while we're ahead: When do averages matter more than sums? *Decision*. Advance online publication. http://dx.doi.org/10.1037/dec0000087

Busemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 538–564. http://dx.doi.org/10.1037/0278-7393.11.3.538

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. http://dx.doi.org/10.1037/0033-295X.100.3.432

Golan, H., & Ert, E. (2015). Pricing decisions from experience: The roles of information-acquisition and response modes. *Cognition, 136,* 9–13. http://dx.doi.org/10.1016/j.cognition.2014.11.008

Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences, 5,* 10–16. http://dx.doi.org/10.1016/S1364-6613(00)01567-9

Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current Biology, 17,* R751–R753. http://dx.doi.org/10.1016/j.cub.2007.06.039

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences, 13,* 517–523. http://dx.doi.org/10.1016/j.tics.2009.09.004

Hills, T. T., Noguchi, T., & Gibbert, M. (2013). Information overload or search-amplified risk? Set size and order effects on decisions from experience. *Psychonomic Bulletin & Review, 20,* 1023–1031. http://dx.doi.org/10.3758/s13423-013-0422-3

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114,* 3–28. http://dx.doi.org/10.1037/0033-2909.114.1.3

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science, 4,* 401–405. http://dx.doi.org/10.1111/j.1467-9280.1993.tb00589.x

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience, 13,* 1292–1298. http://dx.doi.org/10.1038/nn.2635

Lichtenstein, S., & Slovic, P. (Eds.). (2006). *The construction of preference*. New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9780511618031

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making, 27,* 146–156. http://dx.doi.org/10.1002/bdm.1792

Malmi, R. A., & Samson, D. J. (1983). Intuitive averaging of categorized numerical stimuli. *Journal of Verbal Learning & Verbal Behavior, 22,* 547–559. http://dx.doi.org/10.1016/S0022-5371(83)90337-7

Markowitz, H. M. (1952). Portfolio selection. *The Journal of Finance, 7,* 77–91.

Noguchi, T., & Hills, T. T. (2016). Experience-based decisions favor riskier alternatives in large sets. *Journal of Behavioral Decision Making, 29,* 489–498. http://dx.doi.org/10.1002/bdm.1893

Oh, H., Beck, J. M., Zhu, P., Sommer, M. A., Ferrari, S., & Egner, T. (2016). Satisficing in split-second decision making is characterized by strategic cue discounting. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 42,* 1937–1956. http://dx.doi.org/10.1037/xlm0000284

Pachur, T., & Scheibehenne, B. (2012). Constructing preference from experience: The endowment effect reflected in external information search. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 1108–1116. http://dx.doi.org/10.1037/a0027637

Pachur, T., & Scheibehenne, B. (2017). Unpacking buyer-seller differences in valuation from experience: A cognitive modeling approach. *Psychonomic Bulletin & Review, 24,* 1742–1773. http://dx.doi.org/10.3758/s13423-017-1237-4

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. http://dx.doi.org/10.1037/0278-7393.14.3.534

Pleskac, T. J., Yu, S., Hopwood, C., & Liu, T. (in press). Mechanisms of deliberation during preferential choice: Perspectives from computational modeling and individual differences. *Decision*.

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review, 122,* 621–647. http://dx.doi.org/10.1037/a0039413

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences, 20,* 260–281. http://dx.doi.org/10.1016/j.tics.2016.01.007

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica, 127,* 258–276. http://dx.doi.org/10.1016/j.actpsy.2007.05.004

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. http://dx.doi.org/10.1037/0033-295X.108.2.370

Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour, 1*(8), 0145. http://dx.doi.org/10.1038/s41562-017-0145

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Boston, MA: MIT Press.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences of the United States of America, 109,* 9659–9664. http://dx.doi.org/10.1073/pnas.1119569109

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences of the United States of America, 113,* 3102–3107. http://dx.doi.org/10.1073/pnas.1519157113

Tsetsos, K., Usher, M., & Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological Review, 117,* 1275–1291. http://dx.doi.org/10.1037/a0020580

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review, 111,* 757–769.

Vanunu, Y. (2016). *Constructing preference from rapid numerical sequences: Task demands and attentional risk bias* (Unpublished master's thesis). Tel Aviv University, Tel Aviv, Israel.

Weber, E. U. (2010). Risk attitude and preference. *WIREs Cognitive Science, 1,* 79–88. http://dx.doi.org/10.1002/wcs.5

Zakay, D. (1993). The impact of time perception processes on decision making under time stress. In O. Svenson & A. J. Maule (Eds.), *Time pressure and stress in human judgment and decision making* (pp. 59–72). New York, NY: Plenum Press. http://dx.doi.org/10.1007/978-1-4757-6846-6_4

Zeigenfuse, M. D., Pleskac, T. J., & Liu, T. (2014). Rapid decisions from experience. *Cognition, 131,* 181–194. http://dx.doi.org/10.1016/j.cognition.2013.12.012

Received June 21, 2017
Revision received February 9, 2018
Accepted May 23, 2018 ■