

Salience driven value integration explains decision biases and preference reversal

Konstantinos Tsetso^{a,b,1}, Nick Chater^c, and Marius Usher^{d,1}

^aDepartment of Cognitive, Perceptual, and Brain Sciences, University College London, London WC1H 0AP, United Kingdom; ^bDepartment of Experimental Psychology, Oxford University, Oxford OX1 3UD, United Kingdom; ^cBehavioural Science Group, Warwick Business School, University of Warwick, Coventry CV4 7AL, England; and ^dSchool of Psychology and Sagol School of Neuroscience, Tel Aviv University, Tel Aviv 69978, Israel

Edited by Barbara Anne Doshier, University of California, Irvine, CA, and approved April 25, 2012 (received for review November 28, 2011)

Human choice behavior exhibits many paradoxical and challenging patterns. Traditional explanations focus on how values are represented, but little is known about how values are integrated. Here we outline a psychophysical task for value integration that can be used as a window on high-level, multiattribute decisions. Participants choose between alternative rapidly presented streams of numerical values. By controlling the temporal distribution of the values, we demonstrate that this process underlies many puzzling choice paradoxes, such as temporal, risk, and framing biases, as well as preference reversals. These phenomena can be explained by a simple mechanism based on the integration of values, weighted by their salience. The salience of a sampled value depends on its temporal order and momentary rank in the decision context, whereas the direction of the weighting is determined by the task framing. We show that many known choice anomalies may arise from the microstructure of the value integration process.

decision making | decoy effects | value psychophysics | expanded judgement

Recent research on the psychology and neuroscience of simple, evidence-based choices (e.g., integrating perceptual or reward information) has made impressive progress, leading to the conclusion that the brain is optimized to make the fastest decision for a specified accuracy (1–5). Accordingly, the observer is assumed to infer the most probable cause of a perceived experience by sequentially accumulating samples of noisy evidence until a response criterion is reached. The idea that simple, evidence-based decision making is optimal contrasts with findings in more complex, motivation-based decisions, focused on multiple goals with tradeoffs (e.g., choices among cars or flats). Here, a number of paradoxical and puzzling choice behaviors (6–8) have been revealed, posing a serious challenge to the development of a unified theory of choice.

Can a common theoretical framework between evidence-based and motivation-based decisions be established? A natural starting point is to propose that, in the latter, the cognitive system integrates subjective values (rather than, say, pieces of perceptual evidence), that depend on how each alternative matches the decision maker's goals (9). In particular, when alternatives are characterized by different attributes (e.g., product price and quality), preference is shaped through shifting attention across these attributes (8, 10), assessing an item's subjective value on each attribute, integrating these values across time, and finally making a choice when some threshold is reached (11–13). A detailed understanding of these computations might explain the systematic anomalies observed in motivation-based decisions.

This line of research has been difficult to pursue, however, because classical laboratory preference tasks provide little control of the moment-by-moment processes of value sampling and integration. This stands in contrast with psychophysical paradigms for studying evidence-based perceptual choice where the flow of sensory evidence can be fully controlled by the experimenter (14, 15). To obtain more precise control on the decision input, we introduce an experimental paradigm, which we call

“value psychophysics,” at the interface of psychophysics and motivation-based decisions, similar to the expanded judgment task developed in a different context (16, 17). Participants simultaneously view two or three rapidly varying sequences of numerical values, described as stock market values or slot machines' past payouts. After each presentation, they choose the sequence with either the highest overall value or the sequence they would like to “play” to obtain a reward sample. Controlling the flow of the input values (Fig. 1A) allows us directly to probe how people attend to and integrate values.

Using this task, we first demonstrated the remarkable ability of the cognitive system to rapidly integrate streams of numerical values and select the alternative with the highest mean value. However, this integration process was subject to distortions; more-salient samples are weighted more heavily. Salience was determined by (i) the temporal order of the sample, with more recent items more heavily weighted, and (ii) by the magnitude of the sample, with larger values being further amplified. Hence, we observed that the value-integration mechanism is sensitive to the variance of the sequences, favoring riskier options in the domain of gains, in direct contrast to the risk aversion predicted by expected utility theory and prospect theory (18). This risk-seeking bias was reversed in the logically equivalent task of rejecting the worst alternative. Based on these findings, we proposed a mechanism for value integration, which accounts for temporal, variance, and task-framing sensitivity by prioritizing the processing of the samples depending on their order and their momentary rank in the decision context. We showed that when this simple mechanism is extended to more than two alternatives, it provides a natural explanation of contextual preference reversal effects in multiattribute choice (6, 8). Finally, we confirmed this account by reproducing analogs of these effects in the value psychophysics paradigm, establishing a strong link between this simple task and the underlying mechanisms of complex goal-directed decisions.

Results

We report four experiments using the value psychophysics paradigm. Experiments 1 and 2 involved selecting between two numerical sequences, testing the presence of differential weighting of the values in time and in the value range, respectively. Experiments 3 and 4 involved choice among three sequences. In experiment 3, a two-stage decision process was used to test the effect of task framing on the integration process; elimination of one of the three options was followed by selection between the remaining

Author contributions: K.T., N.C., and M.U. designed research; K.T. performed research; K.T. analyzed data; and K.T., N.C., and M.U. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: konstantinos.tsetso@psy.ox.ac.uk or marius@post.tau.ac.il.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1119569109/-DCSupplemental.

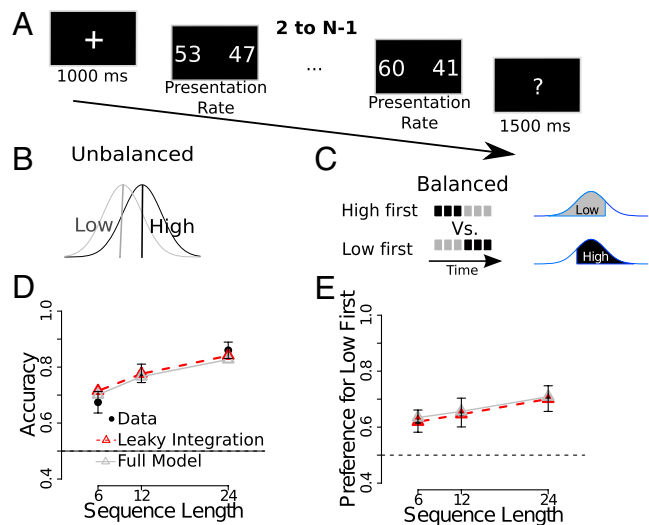


Fig. 1. Decision task and results in experiment 1. (A) The timeline of a trial. At the end of the presentation, participants decided which sequence had the highest average. The unbalanced conditions (B) consisted of two sequences generated from Gaussians with different means. In the balanced condition (C), the sequences corresponded to equal mean Gaussians, with one alternative sampled from the lower range (gray) during the first half, and from the higher range (black) during the second half of the trial (and conversely for the other alternative). (D) The decision accuracy in the unbalanced trials improves with sequence length. (E) The preference for the alternative associated with higher values at the end of the sequence shows recency, increasing with sequence length. Error bars correspond to 95% confidence interval. Red symbols (dashed line) correspond to leaky integration fits (Eq. 1), and square symbols (solid gray line) to fits of the full model (Eq. 2).

two. In experiment 4, we created analogs of preference reversal effects that typically occur in multiattribute choice problems.

Value Integration and Recency Bias. In experiment 1, we examined participants' ability to perform the task by presenting them with sequences of pairs of numbers at a rate of two or four pairs per second and asking them to select the alternative with the highest average (experiment 1; Fig. 1A). At both rates (*SI Results*), participants could select the alternative associated with the highest overall value [mean = 0.77, SD = 0.05, range: 0.68–0.90, $t(14) = 18.93$, $P < 0.001$] and accuracy increased with sequence length [$F(1, 14) = 91.09$, $P < 0.001$, mean = 0.65, SD = 0.09 for $n = 6$ items; mean = 0.85, SD = 0.06 for $n = 24$ items; Fig. 1D], indicating integration across the whole stream. This conclusion was further supported by the rejection of two candidate heuristic rules. The first heuristic we examined is to simply select the sequence containing the overall maximum value (19). However, even when the maximum value appeared in the low-average sequence (“peak-low” condition), participants still chose the high-average alternative [$t(15) = 16.43$, $P < 0.001$; Fig. S1A and see also *SI Results* for a discussion of the “peak-end” heuristic]. A second possible heuristic is to choose according to a small subset (k) of values that are maintained in a memory buffer (20); but if so, performance should have stopped improving with sequence length, unless the buffer capacity, k , is as large as the maximum sequence length (when the heuristic becomes equivalent to integration).

To further scrutinize the properties of this integration mechanism, we then examined the presence of order effects that might indicate differential weighting of values across time. We compared choice preference for alternatives with the same mean (balanced sequences; Fig. 1C), but with different temporal distribution of values, such that one option appeared better in the first half and worse in the second. Both presentation rates (*SI*

Results) revealed a clear temporal bias. The values of recent pairs were more strongly weighted [$t(15) = 7.76$, $P < 0.001$] and, moreover, recency increased with sequence length [$F(1, 14) = 15.89$, $P < 0.005$; Fig. 1E]. Both the increase of accuracy and recency as the stream length increased are consistent with a simple, leaky (decay-based) accumulation model that integrates all samples but places higher weights on the later items. Assuming two sequences of numbers, V_A and V_B , which are presented sequentially for N frames, the preference state, $P(t)$, at frame t is defined as

$$P(t) = \lambda P(t - 1) + [V_A(t) - V_B(t)] + N(0, \sigma), \quad [1]$$

with λ denoting the degree of decay and $N(0, \sigma)$ Gaussian noise. After each trial, if the preference state is positive a decision is made in favor of A and otherwise for B. The average fits of this model across all participants are presented with red symbols in Fig. 1D and E (see *SI Models* for parameters). This model roughly calculates a weighted sum of the samples, with the weights decreasing exponentially from the last to the first item (21). Thus, the longer the sequences, the less the impact of the early items (Fig. S1B).

Response to Variance and Risk Attitudes. The order effect in experiment 1 posits that the integration process assigns higher weights to items that are more recent and therefore more salient, consistent with data and theories of judgment (22). This finding opens the possibility that the decision mechanism is also susceptible to other factors that affect salience. In the value psychophysics paradigm, one such factor could be the magnitude of the value samples; larger values might be more salient and consequently overweighted. Where the alternatives have equal variances and different means, this strategy might facilitate the detection of the best sequence, because the high-mean option is naturally related to larger values. However, when the two alternatives have equal means but different variances, attention will focus on very large values at the right tail of the high variance distribution. Consequently, the decision maker would ignore the low values generated from the left tail of the high variance distribution, hence favoring the riskier option, associated with the broad Gaussian.

To test for this prurisk bias, we probed in experiment 2 the sensitivity of the choice mechanism to sequence variance. To ensure that the results apply to preference and not only to judgments of magnitude, we used two response modes: half of the respondents chose the sequence with the highest average; the other half chose from which sequence they preferred to receive a reward sample. Participants had to choose between two sequences with different variances (broad and narrow). The three conditions used are shown in Fig. 2A–C. Fig. 2C corresponds to the critical condition with two equal-mean distributions, and Fig. 2A and B correspond to two unbalanced conditions, where the broad distribution has a higher or lower mean than the narrower distribution. If choice was sensitive to the mean only, respondents should be indifferent between the two equal mean sequences of the critical condition, and performance should not differ in the two unbalanced conditions.

The results were invariant to the response mode (*SI Results*), and participants preferred the alternative with the highest mean in both unbalanced conditions [Fig. 2D, $t(15) = 8.34$, $P < 0.001$, and Fig. 2E, $t(15) = 13.36$, $P < 0.001$]. Furthermore, accuracy was higher when the broad distribution had the highest mean [$t(15) = 3.62$, $P < 0.005$], indicating a bias toward large values; this was confirmed by the choice pattern in the critical condition and the strong preference for the high-variance alternative [Fig. 2F, $t(15) = 5.39$, $P < 0.001$].

This risk-seeking pattern may seem surprising because it collides with findings from the mainstream research in risky choice—

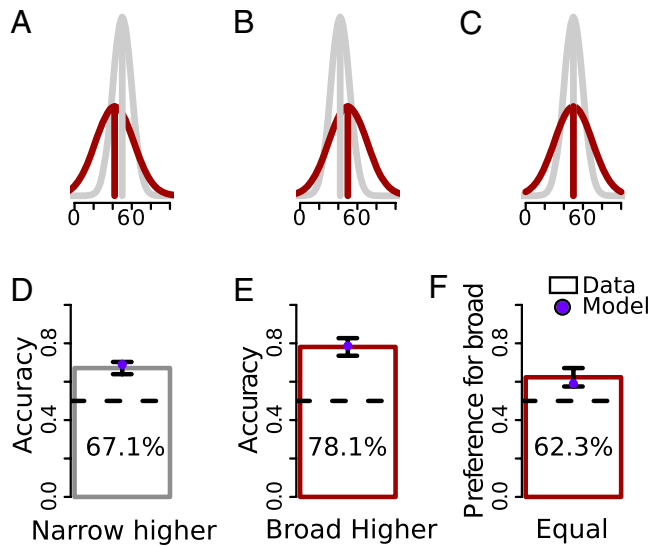


Fig. 2. Experiment 2 conditions and results. Observers decided between two alternatives, each characterized by a sequence of 12 values, presented as pairs at a rate of 2/s. In two unbalanced conditions (A and B), either the broad or the narrow distribution had the highest mean, whereas in the equal condition both distributions had equal mean and different variance (C). (D and E) Decision accuracy and (F) preference for the risky alternative, associated to the broad distribution. Purple circles indicate the fits of the rank-weighted model (Eq. 2). Individual data are given in Fig. S2A. Error bars correspond to 95% confidence interval.

decisions by descriptions (in which the probabilities and payoff information of monetary gambles is explicitly provided) where people are typically risk averse in gains (18). However, our result is in the same direction with a recent finding in experience-based decisions (23) (where people learn about probabilistic outcomes through active sampling); it is also consistent with a qualitative theory applied to scenario-based decisions: reason-based decision making (24, 25). According to this framework, decision making is highly flexible such that advantages loom larger in selection decisions (as in experiment 2) and disadvantages loom larger in rejection decisions (25).

Risk Attitudes Depend on Task Framing. If this flexibility of high-level choice applies to the mechanism that underlies the value psychophysics paradigm, a striking prediction follows: the propensity of the decision maker to choose the risky option should depend on the task framing. To test this prediction, we introduced in experiment 3 a two-stage decision task. We first presented a sequence of 12 triples, two of which corresponded to low variance/low risk and one to high variance/high risk (Fig. 3A). The samples were described as past outcomes of three slot machines, and participants were asked to first eliminate the worst option. The two remaining alternatives were shown for 12 more frames, and participants then had to select one of them. Unknown to the respondents, when the high-risk alternative was rejected during the elimination stage it was not discounted. Instead, it remained available for selection in the second stage, by covertly replacing one of the two remaining low-risk alternatives (Fig. 3A).

The results show a risk attitude reversal between the two stages. Respondents rejected the high-risk alternative in the first stage more than chance [Fig. 3B, $t(14) = 3.27$, $P < 0.001$]. However, consistent with the results in experiment 2, they subsequently showed the opposite risk-seeking preference by selecting that same alternative (now covertly replacing one of the two remaining options) above chance [Fig. 3B, $t(14) = 4.81$, $P < 0.001$]. We obtained the same reversal in a binary task (narrow vs. broad) that manipulated the task framing between participants

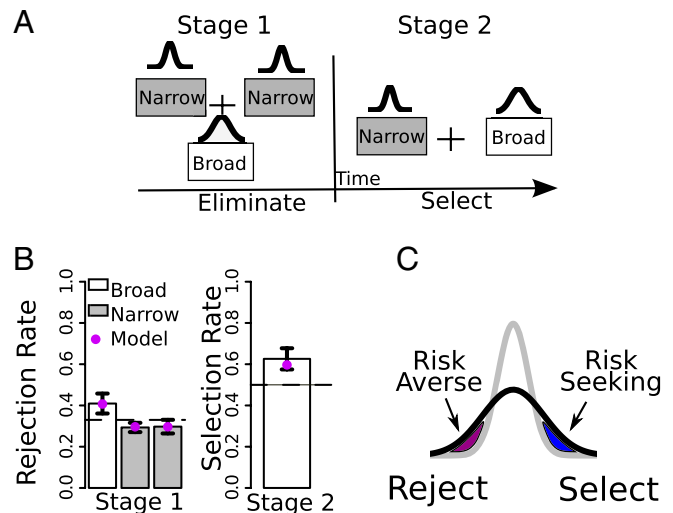


Fig. 3. Two-stage decision task and results in experiment 3. (A) Participants saw 12 triples presented at a rate of 750 ms and were first asked to eliminate one of them (stage 1), and then to select one from the remaining two (stage 2), which were presented as a second sequence of 12 pairs at a rate of 500 ms. (B) In the first stage, the rejection rate of the risky alternative was higher than chance (33%); in the second stage, the selection rate for it was also higher than chance (50%), consistent with an account that weighs different sides of the distribution depending on the task framing (C and Eq. 2). Purple circles indicate the fits of the rank-weighted model (Eq. 2). Individual data are given in Fig. S2B. Error bars correspond to 95% confidence interval.

(experiment 5 in *SI Results*). This preference reversal reveals that the task framing modulates the salience of the sampled information. Furthermore, it violates the principle of invariance (26) and is incompatible with theories of choice, which assume that risk attitudes are stable and task independent.

Rank-Dependent Weighting. Experiments 1–3 revealed the sensitivity of the decision mechanism to the prominence of the processed items. In experiment 1, salience depended on the temporal order of the information; in experiment 2, salience depended on the magnitude of the sampled value. The direction of the latter type of differential weighting was determined by the task framing as selection or rejection (experiment 3 and Fig. 3C). This flexibility in the weighting of information can be understood in terms of a top-down mechanism, which privileges the processing of sampled values as a function of their momentary ranks. To mathematically capture this pattern, we assume that the values, V_i , for each alternative, i , are weighed by their momentary ranks and integrated in separate leaky accumulators with preference states $P_i(t)$ (extending Eq. 1):

$$P_i(t) = \lambda P_i(t - 1) + [V_i(t) \cdot w(\text{rank}_i(t))] + N(0, \sigma), \quad [2]$$

with $w(\text{max}) > 1$ and $w(\text{min}) = 1$ in selection and $w(\text{max}) < 1$ and $w(\text{min}) = 1$ in rejection decisions; $\text{rank}_i(t)$ is the momentary rank of item i at time t . In selection decisions, the alternative associated to the accumulator with the highest P_i is chosen, whereas in rejection, the accumulator with the lowest P_i is eliminated. This model accounts for the data in experiments 2 and 3 (purple circles in Figs. 2D and E and 3B) as well as for the data in experiment 1 (gray lines, Fig. 1D and E; see also *SI Models* and Fig. S1A).

A direct prediction of the rank-dependent mechanism is the context sensitivity of evaluation (26); the computed value of the very same sequence will be suppressed when it is paired with a better sequence, compared with the case when it is evaluated

against an inferior sequence (27). Moreover, the preference state between two target sequences might be reversed by the addition of a third, irrelevant sequence that changes the momentary rank ordering of the values of the target alternatives.

Contextual Preference Reversal. Preference reversal due to the addition of decoy options (that are not chosen) or other irrelevant options to the choice set has been extensively studied within the multiattribute choice literature, revealing, among others, two puzzling effects: the asymmetric dominance or attraction (6) and the similarity effect (8). Assuming two options A and C that are defined on two dimensions, e.g., economy (E) and quality (Q), and people are indifferent between (E_A, Q_A) and (E_C, Q_C), with $E_A < E_C$ and $Q_A > Q_C$, the attraction effect is a preference bias in favor of A, when a decoy B, similar but overall inferior to A, is introduced in the choice set (Fig. 4A, Upper). The similarity effect corresponds to a reduction in the ratio of choices for A relative to C, when an alternative B, similar to A and of equal value, is added to the choice set (Fig. 4A, Lower). These effects are difficult to account for in many theories of choice; it is thus natural to ask whether they can be accounted within the rank-dependent framework proposed here.

One plausible way by which people may process multiattribute choice problems, suggested in Tversky's early work (8) and then extended in the decision-field theory model (11, 12), is by sequentially switching attention from one choice aspect to another. Assuming that both attributes are sampled independently and with equal probability, the rank-weighted additive value of an option will be determined by its attribute values, the corresponding (attribute-wise) ranks and how often each value/rank combination occurs. For demonstration purposes, we denote the high values of A and C as H and their low values L and assume that the highest-ranked value is multiplied by $a = 3$, the second by $b = 2$, and the third by $c = 1$. When A and C are the only available options, A ranks first in its strong dimension and

second in the other (and vice versa for C). In the attraction effect, the addition of the inferior decoy (B) leaves A's rank order intact but downgrades C to the third position 50% of the time when its weak dimension is sampled (i.e., ordering in quality $A > B > C$ and in economy $C > A > B$). Hence, although there is no preference between A and C when they are considered alone, the rank-dependent additive utility of A will be larger because it overall ranks higher: $V_A = 0.5 \cdot aH + 0.5 \cdot bL > V_C = 0.5 \cdot aH + 0.5 \cdot cL$, which is always the case for $a > b > c$.

For the similarity condition, adding a similar and equal option B reduces the preference for the target A. Assuming that A and B are identical and there is some noise fluctuations in the encoding of their values, these two options will alternate in the first/second (in quality) and second/third (in economy) positions, whereas the dissimilar alternative C will be clearly either the best (first in economy) or the worst (third in quality). In other words, A (and B) will rank 25% of the time first, 50% second, and 25% third, with its competitor C ranking 50% of the time first and 50% third. The rank-weighted additive value of A will be $V_A = 0.25 \cdot aH + 0.25 \cdot bH + 0.25 \cdot bL + 0.25 \cdot cL = 0.25[(a + b) \cdot H + (b + c) \cdot L]$. For C, $V_C = 0.5 \cdot (aH + cL)$. If V_C is better than V_A , then $0.5 \cdot (aH + cL) > 0.25[(a + b) \cdot H + (b + c) \cdot L]$ or $H(a - b) > L(b - c)$, which is always the case for $a = 3 > b = 2 > c = 1$ (very similar arguments have also been explored in the decision-by-sampling framework) (28).

To create analogs of these effects in our value psychophysics paradigm, and hence test the viability of the rank-dependent account for more complex decisions, we temporally manipulated the values of the sequences so that the instantaneous ranking of the alternatives is precisely controlled (Table 1 and Fig. 4B; ref 15). In the decoy conditions, two of the sequences (A, C) were equal overall, but temporally anticorrelated; when A received a high value, C received a low value and vice versa, as if attention switched from quality to economy in Fig. 4A, favoring each time a different option. The third alternative, B, was then temporally controlled to have similar numerical values to A. In the attraction condition, these values were, at each time frame, constrained to be slightly lower than those of A, whereas in the similarity condition A and B had overall the same mean value, without constraining their momentary ordering. In addition to these two critical conditions, the experiment had two dominance conditions in which one of the alternatives had the highest mean, to make the task engaging for participants and provide a measure of decision accuracy. All conditions were randomized within the experimental blocks.

In the dominance conditions, participants successfully chose the highest-value alternative [$t(19) = 27.77, P < 0.001$, Fig. 4D, Right]; they also showed the predicted choice patterns corresponding to preference reversals both in the attraction and in the similarity condition (see Fig. S2C for individual results). Participants preferred alternative A, which dominates the decoy (B) at every time step, instead of the anticorrelated alternative (C) [$t(19) = 5.04, P < 0.001$; Fig. 4D, Left]. In the similarity condition, where overall all three alternatives had equal net values, observers preferred the anticorrelated alternative (C) compared

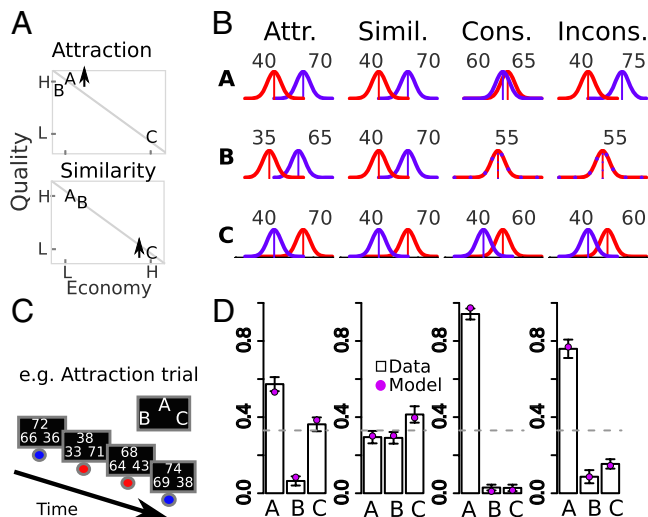


Fig. 4. Experiment 4: creating analogs of the attraction and similarity effects (A, B) using the fast value integration task with three alternatives. Each alternative is associated with two distributions, one red and one blue (colors for illustration purposes only), and at each frame the values for all three alternatives are sampled from either the red or the blue Gaussian distributions (randomly determined; see Table 1 and B for exact values). (C) Four frames from one experimental trial in the attraction condition. (D) Results for the four conditions, showing reversal effects in the decoy conditions. Individual data for the decoy conditions are given in Fig. S2C. Purple circles indicate the fits of the rank-weighted model (Eq. 3). Error bars correspond to 95% confidence interval.

Table 1. Mean values for each option in each distribution (red/blue) in the four conditions of experiment 4

Decoy conditions				Dominance conditions			
Attraction		Similarity		Consistent		Inconsistent	
B	R	B	R	B	R	B	R
70	40	70	40	60	65	75	40
65	35	70	40	55	55	55	55
40	70	40	70	40	60	40	60

with the two correlated ones (A and B) [$t(19) = 3.40, P < 0.005$; Fig. 4D, second from left). Furthermore, an analysis of the error pattern in the “dominance-inconsistent” condition shows that when failing to select the best option (A), participants chose the worst overall option (C) significantly more than the second best (B) [$t(19) = 4.37, P < 0.001$]. The increased preference for the worst option (C) is a clear signature of rank-dependent weighting, because what makes C stand out is that in half of the frames it ranks first (red distribution, Fig. 4B), whereas the overall better option B ranks always second. Extending Eq. 2 to three alternatives captures the mean choice across all participants (purple circles in Fig. 4D; *SI Models* for fitting parameters):

$$P_i(t) = \lambda P_i(t-1) + [V_i(t) \cdot w(\text{rank}_i(t))] + N(0, \sigma), \quad [3]$$

with $w(1) = a, w(2) = b, w(3) = c$ and $a > b > c$.

Discussion

Preference formation arises from the integration of multiple values that are actively sampled either from the environment (11, 12) or from memory (28). Here, using a psychophysical task where numerical values are sequentially presented, we controlled the sampling process and probed the micromechanism of value integration. We first tested our experimental method by showing that participants are sensitive to the strength of the values, accurately choosing the sequence with the highest mean. Furthermore, we showed that performance improves steadily as sequence length increases (reaching accuracy in the range: 0.75–0.96, at $n = 24$ items-long sequences), suggesting continuous integration with a large temporal span. Crucially, this averaging process was not unbiased but weighted by the temporal order of the values, with late items being more salient and thus overweighted.

Examining the sensitivity of the respondents to the variance of the sequences revealed a second source of value distortion. Participants showed increased preference for the highest variance sequence when choosing the best alternative, in contradiction to findings in description-based decisions, but consistent (18) with some findings in experience-based decisions (23). The risk-seeking bias was reversed in the logically equivalent task of rejecting the worst alternative. This reversal shows the dependence of risk attitudes on the task framing, and complements results in scenario-based choice (24). We accounted for this flexibility by assuming that the value-integration mechanism overweighted high-ranked values that are congruent with the decision-maker’s objective of selecting the best option (and vice versa for rejection decisions).

Weighting the instantaneous evidence by the momentary ranks might result in more robust decisions when the best option ranks higher more often. As we have seen, however, this strategy leads to choice anomalies when the rank order of the options does not reflect their overall goodness (i.e., options with equal means and different variances). Moreover, adding more options in the choice set can switch the time course of the rank ordering of the alternatives, even if their values remain unchanged. We confirmed this prediction by adding a third alternative in the choice set and obtaining analogs of contextual preference reversal that have been so far studied in multiattribute decisions. It is noteworthy that the attraction and similarity effects have never before been reproduced within participants and using the same experimental task. Beyond their immediate empirical significance, these results have further implications.

First, the obtained context effects validate the viability of the rank-dependent weighting mechanism, because these effects stem as a direct prediction of our proposed salience-based model; second, they link the new value psychophysics paradigm to high-order decisions, showing that our technique of controlling the sampling process is a good proxy to the study of decisions in richer domains. These two points clarify how people integrate

values across attributes and why their preference is subject to reversal in multiattribute problems.

To conclude, we introduced an experimental protocol, value psychophysics, which revealed the impressive ability of the cognitive system to integrate rapid streams of numerical values, extending the already established human ability to judge numerosity (29) and to integrate emotional affect associated with rewards (30). This protocol allowed precise moment-by-moment control of the sampling process and probed the micromechanism of value integration. Our findings indicated that choice is distorted by the differential weighting applied on the salient sampled values. Two factors were found to affect the salience of the samples: first, the temporal order (with recent pairs being more important), and second, the momentary rank in the decision context. These two aspects of value integration can account for classical decision paradoxes, such as temporal biases, risk preferences, and contextual preference reversal effects. Our findings underscore the possibility that key decision biases may derive from the nature of the basic computational operations from which the decision-making processes are built.

Methods

Participants. Overall, 67 adults (36 females) were recruited from University College London’s subject pool with ages ranging from 19 to 44 y (mean = 25.5; number of participants in each study: $N_1 = 16, N_2 = 16, N_3 = 15, N_4 = 20$). Experiment 1 was divided in two sessions with a maximum 3-d lag between them. All participants consented before the experiment and received a monetary reward of £7/h for their participation. In experiments 2 (half of the participants, $n = 8$) and 3, participants received a bonus of £2 maximum, randomly determined from their reward history during the task. Approval from the local ethics committee was provided.

Stimuli. The stimulus consisted of pairs or triples of numerical values presented simultaneously for several frames and with variable rate. The values were described as returns of stocks [experiment 1, experiment 2 (feedback group), and experiment 4] or as past outcomes of casino slot machines [experiment 2 (reward group) and experiment 3]. The presentation rate and length differed across the four experiments. The sequences were normally distributed and generated using MATLAB (MathWorks) and the COGENT toolbox (<http://www.vislab.ucl.ac.uk/cogent.php>). A detailed description of the stimuli in each experiment is given in *SI Methods*.

Task and Design. Each trial started with the presentation of a white fixation cross for 1 s. Participants saw pairs (experiments 1 and 2) or triples (experiments 3 and 4) of numbers presented sequentially. Upon the presentation of a response cue (question mark), participants had to decide, within 1.5 s, which of the sequences had the highest average (experiments 1, 4 and $n = 8$ participants in experiment 2) or which sequence they would like to draw an extra reward sample from ($n = 8$ in experiment 2 and selection stage in experiment 3). Experiment 3 consisted of two stages. In the first stage, participants, when prompted, had to eliminate one of the three options. Samples from the two remaining options were subsequently presented, and participants were asked to choose the sequence they would like to draw an extra reward sample from.

In experiment 1, the presentation rate was varied between participants (0.25, 0.45 s), and the length of the sequences was manipulated within participants (6, 12, and 24 frames). All respondents had to choose the sequence with the highest average, and error feedback was provided. In experiment 2, the presentation rate was 0.5 s, and the sequence length was fixed at 12 frames. The response mode was manipulated between subjects with half of the participants ($n = 8$, feedback group) choosing the highest average option and receiving error feedback and the other half (reward group) receiving an extra reward sample from their chosen sequence. In experiment 3, in the elimination stage, the presentation rate was 0.75 s, and the length of the sequences was set to 12 frames. In the selection stage, 12 frames were shown, but the presentation rate was faster (0.5 s). In experiment 4, the sequences length was 12 frames, and the presentation rate was 0.5 s for half of the respondents ($n = 10$) and 1 s for the other half. In all experiments, the trials were presented fully randomized across the conditions in blocks of 30 trials each. Responses were indicated by the press of the left and right (experiments 1 and 2) and top arrow buttons (experiments 3 and 4) on the keypad of a standard personal computer.

Experimental Conditions. In experiment 1 we manipulated the means of the Gaussian sequences. For each participant, the SD of the sequences was estimated using a staircase procedure before the main experiment (*SI Results*). In two of the conditions the sequences were unbalanced, differing in their means by eight units. Choice in these trials provided a measure of accuracy. In the first unbalanced “regular” condition (150 trials, 50 for each sequence length), the mean of the highest sequence was sampled from a uniform distribution (range: 45–55). The second unbalanced peak-low condition differed from the regular one in that the overall highest value was placed in the low-mean sequence. These trials were equated in the overall averaged differences between high and low, to those of the regular condition. Finally, the balanced condition involved choice between two sequences generated from the same distribution [mean sampled from $U(45, 55)$, 300 trials, and no error feedback]. In the first half of each trial, the values of the first option (labeled “high-first”) were sampled from a truncated Gaussian, clipped 1 SD below the mean, and 1 SD above the mean in the second half (and vice versa for the “low-first” option). Choice in these trials provided a direct measure of the temporal bias (primacy or recency). The overall average differences were equated for sequences of different length in all conditions.

In experiment 2, we checked the sensitivity of the participants to the variances of the sequences (150 trials). There were three conditions, and in each of them one alternative was always associated to a broad distribution [$N(\mu_b, 20)$], whereas the other was associated to a narrow distribution [$N(\mu_n, 10)$]. The first two conditions were unbalanced, and the mean of the high sequence was generated from $U(45, 55)$. In “narrow-higher,” the mean of the narrow distribution was eight units larger than the mean of the broad

($\mu_n = \mu_b + 8$; and vice versa for the “broad-higher,” i.e., $\mu_b = \mu_n + 8$). In the equal condition, the two sequences had equal means ($\mu_n = \mu_b$). The choice pattern in this condition provided a measure of risk attitude.

In experiment 3, participants performed 100 trials. All three options were generated from Gaussians with the same mean value [sampled in each trial from $U(45, 55)$]. The two options had a SD of 10 (narrow), and the third option had a SD of 20 (broad). In the trials where the broad option was eliminated at the first stage, beyond the participants’ awareness, the distribution of one of the two remaining narrow options was turned from narrow (SD = 10) to broad (SD = 20) for the 12 remaining frames of the second stage. None of the participants detected this change during the task.

In experiment 4, there were overall four conditions (55 trials each). Each option was associated with two distributions, labeled here as blue and red (colors for description purposes only), with SDs fixed at 7. At each trial, six triples were generated from the blue distributions, and the other six triples were obtained from the red ones. Before the experiment, the triples were reshuffled and thus at each frame there was a 50% probability for all three values to be sampled from either the blue or the red distributions. Table 1 shows the means of the distributions for each option in each condition. Error feedback was given only in the dominance conditions. In the attraction condition, the values were sampled such that A values were always greater or equal to B values. In the consistent condition, the values were constrained such that $V_A > V_B > V_C$ at each frame. In the inconsistent condition, the order of the blue samples was always $V_B > V_C > V_A$ (and $V_A > V_C > V_B$ for the red samples).

1. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113:700–765.
2. de Gardelle V, Summerfield C (2011) Robust averaging during perceptual judgment. *Proc Natl Acad Sci USA* 108:13341–13346.
3. Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
4. Rorie AE, Gao J, McClelland JL, Newsome WT (2010) Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS ONE* 5:e9308.
5. Sugrue LP, Corrado GS, Newsome WT (2005) Choosing the greater of two goods: Neural currencies for valuation and decision making. *Nat Rev Neurosci* 6:363–375.
6. Huber J, Payne JW, Puto C (1982) Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *J Consum Res* 9:90–98.
7. Kahneman D, Tversky A (1984) Choices, values, and frames. *Am Psychol* 39:341–350.
8. Tversky A (1972) Elimination by aspects: A theory of choice. *Psychol Rev* 79:281–299.
9. Busemeyer JR, Townsend JT (1993) Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychol Rev* 100:432–459.
10. Russo JE, Doshier BA (1983) Strategies for multiattribute binary choice. *J Exp Psychol Learn Mem Cogn* 9:676–696.
11. Diederich A, Busemeyer JR (2006) Modeling the effects of payoff on response bias in a perceptual discrimination task: Bound-change, drift-rate-change, or two-stage-processing hypothesis. *Percept Psychophys* 68:194–207.
12. Roe RM, Busemeyer JR, Townsend JT (2001) Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychol Rev* 108:370–392.
13. Usher M, McClelland JL (2004) Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol Rev* 111:757–769.
14. Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci* 28:3017–3029.
15. Tsetos K, Usher M, McClelland JL (2011) Testing multi-alternative decision models with non-stationary evidence. *Front Neurosci* 5:63.
16. Irwin FW, Smith WAS, Mayfield JF (1956) Tests of two theories of decision in an expanded judgment situation. *J Exp Psychol* 51:261–268.
17. Vickers D, Burt J, Smith P, Brown M (1985) Experimental paradigms emphasizing state or process limitations. 1. Effects on speed accuracy tradeoffs. *Acta Psychol* 59(2):129–161.
18. Kahneman D, Tversky A (1979) Prospect theory: Analysis of decision under risk. *Econometrica* 47:263–291.
19. Langer T, Sarin R, Weber M (2005) The retrospective evaluation of payment sequences: Duration neglect and peak-and-end effects. *J Econ Behav Organ* 58:157–175.
20. Hertwig R, Pleskac TJ (2010) Decisions from experience: Why small samples? *Cognition* 115:225–237.
21. Busemeyer JR, Myung IJ (1992) An adaptive approach to human decision-making: Learning theory, decision theory, and human performance. *J Exp Psychol Gen* 121(2):177–194.
22. Anderson NH (1981) *Foundations of Information Integration Theory* (Academic, New York).
23. Ludvig EA, Spetch ML (2011) Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS ONE* 6:e20262.
24. Shafir E (1993) Choosing versus rejecting: Why some options are both better and worse than others. *Mem Cognit* 21:546–556.
25. Shafir E, Simonson I, Tversky A (1993) Reason-based choice. *Cognition* 49:11–36.
26. Tversky A, Simonson I (1993) Context-dependent preferences. *Manage Sci* 39:1179–1189.
27. Louie K, Gratton LE, Glimcher PW (2011) Reward value-based gain control: Divisive normalization in parietal cortex. *J Neurosci* 31:10627–10639.
28. Stewart N, Chater N, Brown GDA (2006) Decision by sampling. *Cognit Psychol* 53:1–26.
29. Barth H, et al. (2006) Non-symbolic arithmetic in adults and young children. *Cognition* 98:199–222.
30. Bechara A, Damasio H, Tranel D, Damasio AR (2005) The Iowa Gambling Task and the somatic marker hypothesis: Some questions and answers. *Trends Cogn Sci* 9:159–162, discussion 162–164.