Economic irrationality is optimal during noisy decision making

Konstantinos Tsetsos[a,b,1], Rani Moran[c,d], James Moreland[e], Nick Chater[f,2], Marius Usher[c,d,2], and Christopher Summerfield[a,2]

[a]Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, United Kingdom; [b]Department of Psychological Sciences, Birkbeck, University of London, London WC1E 7HX, United Kingdom; [c]School of Psychology, University of Tel Aviv, Tel Aviv 69978, Israel; [d]Sagol School of Neuroscience, University of Tel Aviv, Tel Aviv 69978, Israel; [e]Department of Psychology, University of Washington, Seattle, WA 98195; and [f]Behavioural Science Group, Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom

Edited by James L. McClelland, Stanford University, Stanford, CA, and approved January 29, 2016 (received for review September 25, 2015)

According to normative theories, reward-maximizing agents should have consistent preferences. Thus, when faced with alternatives *A*, *B*, and *C*, an individual preferring *A* to *B* and *B* to *C* should prefer *A* to *C*. However, it has been widely argued that humans can incur losses by violating this axiom of transitivity, despite strong evolutionary pressure for reward-maximizing choices. Here, adopting a biologically plausible computational framework, we show that intransitive (and thus economically irrational) choices paradoxically improve accuracy (and subsequent economic rewards) when decision formation is corrupted by internal neural noise. Over three experiments, we show that humans accumulate evidence over time using a "selective integration" policy that discards information about alternatives with momentarily lower value. This policy predicts violations of the axiom of transitivity when three equally valued alternatives differ circularly in their number of winning samples. We confirm this prediction in a fourth experiment reporting significant violations of weak stochastic transitivity in human observers. Crucially, we show that relying on selective integration protects choices against "late" noise that otherwise corrupts decision formation beyond the sensory stage. Indeed, we report that individuals with higher late noise relied more strongly on selective integration. These findings suggest that violations of rational choice theory reflect adaptive computations that have evolved in response to irreducible noise during neural information processing.

decision making | irrationality | choice optimality | selective integration | evidence accumulation

Daily decisions, such as choosing a holiday destination or accepting a job offer, involve comparing alternatives that are characterized by different attributes (1, 2). Understanding how the brain combines information from different attributes into unitary decision values is a key challenge in psychology and the neurosciences (3, 4). From a normative perspective, the value of an alternative should be independent of factors, such as the attractiveness of competing alternatives or the context in which preferences are elicited (5). Thus, the preference relationship between two alternatives ought to remain stable, regardless of changes to the choice set, incurred for example by the addition or removal of other choice alternatives (6).

However, human preferences are often driven by irrelevant factors (7, 8). For instance, an initial preference for one holiday destination (e.g., Bali) over another (e.g., Berlin) can reverse when an inferior alternative (e.g., Dresden) is added to the choice set, even if this "decoy" alternative is never chosen (9, 10). Similarly, an individual preferring a holiday in Bali to Berlin, and Berlin to Boston, will sometimes show a systematic "intransitive" (or inconsistent) preference for Boston over Bali (11). A canonical argument states that such violations of decision theory (hereafter "economic" or "choice irrationality") disclose fundamental limitations in human processing capacity and of the executive system (12, 13). However, this argument does not have an obvious normative justification. Why did the computations that underlie irrational choices survive millions of years of evolutionary pressure for optimal behaviors that maximize reward? [The term "optimal" is most often used to refer to a policy that yields the highest reward rate, given the likely sources of uncertainty in the environment (e.g., the performance-limiting visibility of a stimulus), and the structure of the task (e.g. the monetary payoff for one action over another). Where independent variables are often economic goods of unknown subjective worth, an optimal policy may be hard to specify, but a rational policy can be defined as one that discloses stable and consistent preferences, as if agents made choices by maximizing a latent subjective utility function.]

Here, we describe an alternative theory, known as "selective integration" (14), that overcomes this challenge by offering a normative justification of choice irrationality. Building on psychophysical research into the neural and computational mechanisms by which decisions are formed via sequential sampling (15), we assume that choice attributes (e.g., the expense, weather, or culture encountered on holiday) are sampled in turn (2), and integrated toward a cumulative decision variable (16). Under selective integration, the gain of processing on each attribute i of an alternative depends on its rank within that attribute, with a selective gating parameter w ($0 < w < 1$) controlling the reduction in gain for the weakest attribute value (e.g., B_i when A, B are offered and $A_i > B_i$). Selective integration thus makes decisions sensitive to the relative ranks of the alternatives within each attribute, over and above their cumulative average value.

Significance

Healthy individuals appear to display inconsistent preferences, preferring *A* over *B*, *B* over *C*, and *C* over *A*. Inconsistent, intransitive preferences of this form are hallmark manifestations of irrational choice behavior and breach the very assumptions of economic theory. Nevertheless, the neurocognitive mechanisms that mediate the formation of intransitive preferences remain elusive. We show that intransitivity arises from a bottleneck mechanism that blocks the processing of momentarily less valuable information. Although this algorithm is by classical definitions suboptimal (permitting the loss of information), we theoretically and empirically demonstrate that it leads to better decisions when accuracy can be compromised by neural noise beyond the sensory stage. Thus, contrary to common belief, choice irrationality is a by-product of purposeful neural computations.

Author contributions: K.T., N.C., M.U., and C.S. designed research; K.T. and J.M. performed research; K.T. and R.M. contributed new reagents/analytic tools; K.T., R.M., J.M., and C.S. analyzed data; and K.T., R.M., N.C., M.U., and C.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

Data deposition: Raw data and simulation codes have been deposited in Zendo (dx.doi.org/10.5281/zenodo.46426).

[1]To whom correspondence should be addressed. Email: k.tsetsos62@gmail.com.

[2]N.C., M.U., and C.S. contributed equally to this work.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1519157113/-/DCSupplemental.

Previously, we showed that, when a third (decoy) alternative (D) alters the ranking of two existing alternatives, A and B, the model successfully predicts a preference reversal, despite the fact that the attribute values of A and B remain intact (Table S1) (14). In the current report, we show that selective integration also explains intransitive choice behavior in humans. Using a psychophysical experimental approach, we show that intransitivity can be provoked in most individuals by simple changes to the relative ordering of the decision information, as predicted by selective integration. Critically, however, using a computational simulation based on the sequential sampling framework, we demonstrate that selective integration paradoxically maximizes the accuracy (and subsequent economic outcomes) in the presence of "late" internal noise arising during decision formation, beyond the sensory stage. Importantly, we show that humans with higher estimated late noise are more prone to integrate selectively. Together, these findings offer a biologically viable, descriptively extended, and normatively motivated explanation of economic irrationality.

Results

Selective Integration and Intransitivity. Under a popular computational framework, decisions between competing alternatives are optimized via sequential sampling and integration (17, 18). In choices between multiattribute economic alternatives, such as holidays, this involves sampling attributes in turn (e.g., expense, weather, culture), accumulating their respective values for each alternative (e.g., Bali, Berlin), and comparing the resulting cumulative decision values to select an action (19). We implement selective integration by adding to this framework a "selective gating" parameter, w, which reduces the gain of accumulation for the weaker attribute value (e.g., weather in Berlin) on each sample from 1 (lossless processing, which is optimal according to decision theory) to $1 - w$ (Fig. 1*A* and *Methods*). [An equivalent implementation of selective integration would overweight the stronger attribute value, leaving intact the weaker value. Although this implementation is functionally analogous to the one depicted in Fig. 1*A*, the two differ in terms of metabolic costs: discounting the value of the local loser engenders a reduction in neural firing rate in contrast to the more costly strategy of amplifying the value (and associated firing rate) of the local winner. Because the experiments presented here were not designed to dissociate these two implementations, we chose to adhere to the less costly one.] After gating (where $w > 0$), the cumulative value of each alternative is not only a function of its attribute values, as normative theory prescribes, but also a function of the ordinal positions of these values within the different attributes (e.g., Table S1).

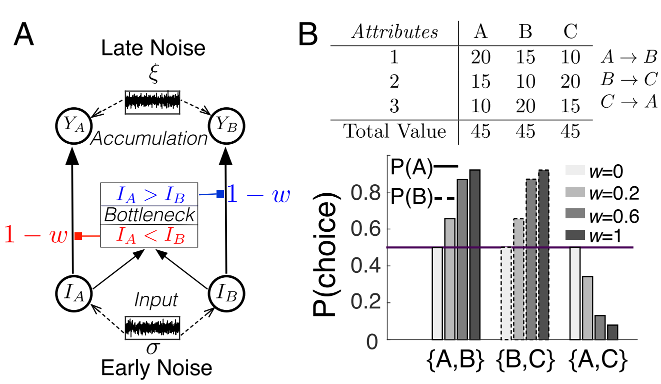

Fig. 1. Selective integration and intransitivity. (*A*) Schematic of the selective integration model. On each time step, the values of two different alternatives on a single attribute are considered. Input samples (I_A, I_B), corresponding to attribute values, feed to a bottleneck that discounts the gain of the weakest sample (via selective gating, w) before relaying the inputs to the accumulators (Y_A, Y_B). Noise can arise both at the input (σ) and accumulation levels (ξ). (*B*) Choice probability for different values of w and for $\sigma = \xi = 2$, for pairwise comparisons between three equally valued multiattribute alternatives (table). $A \rightarrow B$: A wins in more samples than B. WST is violated for $w > 0$ [i.e., $P(A|\{A,B\}) > 0.5$, $P(B|\{B,C\}) > 0.5$, $P(A|\{A,C\}) < 0.5$, with $P(X|\{X,Y\})$ denoting the probability of choosing X over Y].

To illustrate how violations of choice rationality in decisions between two alternatives can arise from selective integration, consider two equally valued alternatives (e.g., A and B in Fig. 1*B*) that differ along three equally important attributes, which are sampled in turn. For $w > 0$, the alternative with two (out of three) winning attributes (A) will (on average) be chosen over an alternative that wins by a larger margin on a single attribute (B), because the input to the latter is more often dampened yielding a lower cumulative value. Thus, when the same three values are permuted circularly in three alternatives, the model predicts a violation of "weak stochastic transitivity" (WST) (11, 20): A is chosen more often over B, B over C, and C over A (Fig. 1*B*; Table S2 for an illustration).

Violations of WST are not only incompatible with normative theories but also with a large class of descriptive theories of choice in which preference tendencies are perturbed by normally distributed noise (16, 21). Thus, when empirically obtained, such violations offer important theoretical constraints. Although WST violations have been reported in humans (11), recent research has shown that the vast majority of these putative violations were not statistically significant when a more appropriate statistical test is applied (20, 22). It is thus an empirical question whether intransitivity, as predicted by our framework, will occur in human observers. Before examining whether humans violate WST in the direction predicted by our model, we first set out to examine how well selective integration characterizes the way humans accumulate evidence over time while forming preferences for different alternatives.

Selective Integration in a Psychophysical Task. We gathered data from human participants performing a psychophysical choice task with real economic incentives for accurate choices (Fig. 2*A*; *Methods*). In experiment 1, participants ($n = 28$) chose between two alternatives each characterized by nine sequentially occurring bars of different heights, presented in two simultaneous streams (at a rate of 400 ms per frame) on the left and right of the screen (Fig. 2*A*). Participants were instructed that the two bars in each presentation frame correspond to two "attribute" values as in the example of Fig. 1*B*. At the end of each trial, they were asked to choose the stream with the larger average height, receiving monetary reward proportional to their choice accuracy. Using the notation $A \rightarrow B$ to indicate that "alternative A has more winning attributes than B" [here, six vs. three winning attributes], we constructed sequences of equal average value, such that $A \rightarrow B$, $B \rightarrow C$, and $C \rightarrow A$, as in the alternatives in Fig. 1*B* ("cyclic" trials). Although participants performed the task accurately (range of 62–92% correct on intermixed "standard" trials where the attribute values were randomly generated from Gaussian distributions), they exhibited a higher preference for the frequently winning stream in the cyclic trials, as predicted by the model [Fig. 2*B*; $P(A|\{A,B\}) = 0.61$, $P(B|\{B, C\}) = 0.63$, $P(C|\{A,C\}) = 0.62$, and $P < 0.001$ for all comparisons to chance; significant frequent-winner effect in 17 out of 28 participants; *SI Results*].

On a further intermixed set of "increment" trials, we increased the average bar height in either the frequently winning (e.g., A when $A \rightarrow B$) or the frequently losing (e.g., B when $A \rightarrow B$) streams, breaking the tie without altering the relative proportion of winning attributes (2:1 in favor of A). Participants were more likely to choose the stream with the increment in both cases (Fig. 2*C*; $P < 0.001$)—being sensitive to the average height difference and not merely the difference in the number of winning attributes—but accuracy was higher when the increment occurred in the frequently winning than the frequently losing stream ($P < 0.001$). These "frequent-winner" effects in cyclic and increment trials were highly correlated across the cohort ($r = 0.75$, $P < 0.001$; Fig. 2*D*) and both correlated positively with participants' estimated w ($r = 0.66$ in cyclic; $r = 0.68$ in increment; $P < 0.001$ in both) (model-fitting procedures and results in *SI Methods* and Fig. S1).

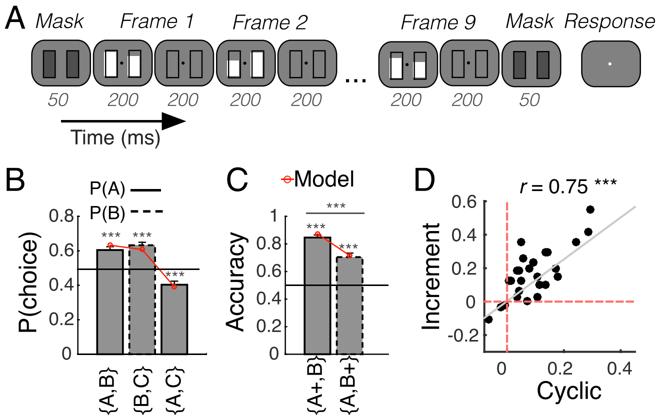

Fig. 2. Behavioral task and selective integration. (*A*) Trial schematic in experiment 1. Participants (n = 28) viewed two streams of bars and had to choose which stream was overall highest. (*B*) Mean choices in cyclic trials in experiment 1 (Fig. 1*B*) revealed a frequent-winner effect: *A* was chosen more often over *B*, *B* over *C*, and *C* over *A*. (*C*) Accuracy in increment trials, where the samples of either the frequently winning or the frequently losing stream were increased by a constant (*A*+ vs. *B* and *A* vs. *B*+, respectively). The difference in accuracy between the *A*+ vs. *B* and the *A* vs. *B*+ trials also revealed a frequent-winner effect. (*D*) The frequent-winner effect in cyclic ($[P(A|\{A,B\})+P(B|\{B, C\})+P(C|\{A,C\})]/3–0.5$) and increment trials $[P(A_+|\{A_+,B\}) - P(B_+|\{A,B_+\})]$ correlated positively to each other. Filled circles correspond to different participants. Grey curve is the linear regression line. Error bars are 2 SEM. ***P < 0.001.

A tendency to prefer the frequently winning stream could be explained by normalization theories (Table S3) that attribute some aspects of choice irrationality on efficient neural coding schemas [i.e., divisive (23) or range-normalization (24) applied on each pair of values before accumulation]. Similarly, a simpler "majority of confirming dimensions" (MCD) rule that decides based on the total number of local winners (25) could also explain the frequent-winner effect (Table S3). One unique signature of selective integration is that in choices between two alternatives with equal mean value but different variances, it predicts a higher preference for the high-variance alternative. This "provariance" effect has been empirically verified elsewhere (14) but was also observed in the current experiment. When the two alternatives had different variances, accuracy was higher when the high variance was assigned to the correct alternative compared with trials where the correct alternative had low variance [standard conditions 3 and 4 in *SI Methods*; accuracy difference between 3 and 4: mean ± SEM = 0.23 ± 0.03; $P < 0.001$]. As shown in Table S4, the provariance effect cannot be captured by normalization theories or by MCD (Fig. S1*B*).

Finally, in two further experiments, we examined whether participants adopted a $w > 0$ due to the lack of processing resources or the scarcity of information, as theories of bounded rationality would advocate (13). When we slowed the presentation rate to 1 Hz (experiment 2), we obtained a similar frequent-winner effect, suggesting that selective gating does not just reflect a processing bottleneck due to the rapid stimulus presentation (Fig. S2*A*). When we increased the sequence length from few (6) to many (12) samples (experiment 3), the frequent-winner effect increased in the latter (Fig. S2*B*), indicating that the tendency to discount losing values does not decrease when more information (samples) is available, contrary to the predictions of heuristic models (26).

Systematic Violations of WST in Human Observers. As shown in Fig. 1*B*, selective integration violates the principle of WST. The preference patterns in the cyclic trials in experiment 1 offer a widely used proxy for the degree of intransitivity; the conclusion that such patterns definitively violate WST can, however, be challenged in statistical grounds (20). [Cyclic trials were created using n different *A*, *B*, *C* triplets (*Methods*). Participants encountered the three pairwise comparisons for each triplet only once. We could count in how many triplets per participant an intransitive circle was obtained. However, the statistical interpretation of this metric of intransitivity, based on pattern counting, is limited and controversial as explained elsewhere (20). Thus the design of experiment 1 was not suitable for rigorous examination of WST violations.] We thus adjusted the experimental design to rigorously assess WST violations within individuals (experiment 4). Participants (N = 21) chose between pairs of alternatives—each corresponding to a job candidate—characterized by three sequentially presented pairs of bars (*Methods*). Each pair of bars was presented within a colored outline, with the color indicating an explicitly defined choice dimension (Fig. S2*C*). The presentation order of the different dimensions was randomized on each trial. The main departure from experiment 1 was that, for each participant, three unique cyclic trials (*A* vs. *B*, *B* vs. *C*, and *A* vs. *C*) were constructed based on a single *A*–*B*–*C* triplet and presented several times, as in multiattribute or risky choice studies of intransitivity (11, 20).

Accuracy in standard trials ranged from 85% to 98%, whereas significant frequent-winner effects were detected in both cyclic (Fig. S2*D*; P < 0.001 in all three comparisons of the frequent-winning option to chance; frequent-winner effect significant in 15 out of 21 participants; *SI Results*) and increment trials (Fig. S2*E*; P < 0.001). As in experiment 1, the two frequent-winner effects were correlated to each other (r = 0.86; P < 0.001) and to the selective gating parameter in the model (r = 0.89 in cyclic; r = 0.96 in increment; P < 0.001 in both). A significant provariance effect was observed in standard trials (0.04 ± 0.01; P < 0.003), ruling out MCD and normalization models. Finally, 11 out of 21 participants violated WST significantly (Table S5). The probability that all these 11 participants corresponded to a type I error is extremely low ($P = 1.1 \times 10^{-9}$). A detailed presentation of these individual-level analyses is given in *SI Methods* and *SI Results*. [It has been recently argued that WST violations can occur spuriously (20). As a remedy, a more stringent test, against the so-called "triangle inequality," has been prescribed. Three participants were intransitive according to this test. However, this test does not seem suitable for our study. First, the chances that the reported WST violations occurred spuriously in our psychophysical task are negligible (Fig. S3 and *SI Results*). Second, the test is conservative in the sense that it would fail to detect real intransitivity effects that, although substantial, are below a certain magnitude due to the presence of experimental noise.]

Selective Integration and Decision Accuracy in the Face of Late Noise. Why do humans discount locally weaker values, provoking intransitive choices (experiment 4) and other violations of economic rationality (14)? We next compared the accuracy (and consequent rewards) that is obtained under selective ($w > 0$) and lossless integration ($w = 0$), simulating an experimental setting in which the attribute values of the two alternatives are generated from two Gaussian distributions with the same variance (σ) and different means (Fig. 3*A*). This setting is equivalent to a two-alternative forced-choice paradigm (15, 17), where one-dimensional quantities (e.g., perceptual signals or economic values of magnitude I_A, I_B) are corrupted by noise, which arises early, before their accumulation (27) (Fig. 1*A*, lower box). In addition to variability at the input level (which can be both due to early internal noise and due to exogenous fluctuations in the stimuli values), we assumed that noise could also arise late, at the level of accumulation (i.e., decision noise ξ; Fig. 1*A*, upper box) (28).

When late noise is absent and early noise is present (Fig. 3*B*, top blue line), integrating samples with equal gain ($w = 0$) is optimal as postulated by statistical decision theory (15, 17, 28). Most surprisingly, however, when late noise is also present, maximum accuracy is achieved for $w > 0$, with the value of w that maximizes accuracy increasing with late noise (Fig. 3*B*, black circles). [The situation reverses and the optimal w regresses toward 0 as the level of early noise increases. This happens because, when early noise is heightened, selective integration is more costly operating almost at random, given that the "frequency of winning" becomes less predictive of the identity of the correct alternative (Fig. S4*A*).] Why does selective integration confer an improvement in accuracy in the face of heightened late noise? On average, discounting locally weaker attribute values via selective gating exaggerates the accumulated differences between higher-valued (typically, frequently winning) and lower-valued (typically, frequently losing) alternatives (Fig. 3*C* for a simulated trial example). This policy occasionally inflates a lower-valued alternative over its higher-valued rival (i.e., when the former is the local winner more often), leading to a slight cost in accuracy relative to lossless integration. When late noise is present, however, the benefit of inflating the accumulated differences offsets this cost.

In Fig. 3*D*, we illustrate this point by plotting the bivariate end-state distributions of the two accumulators, under no late noise (top panels) and high (bottom panels) late noise. Relative to lossless integration ($w = 0$), selective integration ($w = 0.5$) drives accumulator states away from the equality line (density right to the line corresponds to percentage of choice errors), yielding more robust preference states and thus higher performance under late noise (bottom left/right panels, percent accuracies signaled on each panel). This robustness is comparable to that observed in nonparametric statistics, where inferences depend on the ranked data, whereas the mechanism is related to the heightened signal-to-noise ratio in chains of neural populations when the gain of individual neurons increases (29). Thus, when noise arises at different stages of the processing hierarchy, selective integration—although it ignores part of the input and leads to violations of transitivity—can outperform the lossless integration algorithm (15, 17) by acting against late noise.

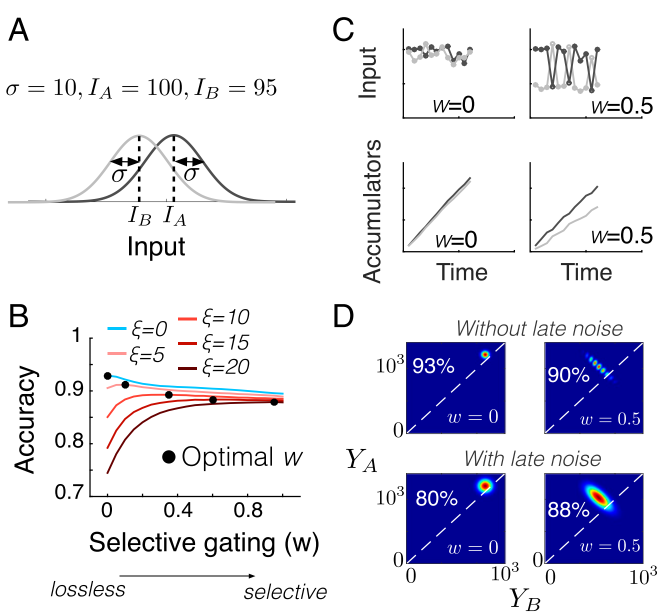

Fig. 3. Selective integration and decision accuracy. (*A*) The input distributions in a typical two-alternative forced-choice scenario. The SD of the distributions (σ) corresponds to early noise. (*B*) Decision accuracy in the model for the scenario in *A*, as a function of *w*, for different levels of late noise (ξ; curves) and after 12 accumulation steps (*t*). Black circles indicate the value of *w* that maximizes accuracy for a given level of late noise. (*C*) Example input (*Top*) and single-trial accumulator states (*Bottom*) for lossless (*Left*) and selective integration (*Right*). The input parameters are as in *A*, and late noise was absent. (*D*) Bivariate end-state ($t = 12$) accumulator distributions for the choice problem in *A*, for lossless (*Left*) and selective integration (*Right*). (*Top*) $\xi = 0$. (*Bottom*) $\xi = 15$. Density to the *Left* of dashed diagonal corresponds to accuracy (in percentage). Higher density is depicted with red, and lower density with blue.

Selective Gating and Late Noise Relationship in Human Observers. Finally, we interrogated the behavioral data to test whether selective integration might be an adaptation specifically evolved to counteract late noise during integration. If so, individuals with higher late noise should have a higher selective gating parameter as depicted in Fig. 3*B*. We fitted three variants of the selective integration model that differed in their assumptions about the source of internal noise: (*i*) a full model, having both early and late noise; (*ii*) a model with early noise only; and (*iii*) a model with late noise only (*SI Methods*). The early noise in the model corresponded to extra internal noise applied at the input representation stage, on top of the stimulus external variability. We factored in the latter by fitting the models using the actual stochastic input that participants saw in the experiments. In addition to noise and selective gating (*w*) parameters, all variants had a leak parameter to capture the recency effect (14) that was observed in all experiments (Fig. S1*A*).

In all experiments, model comparison favored the selective integration variant that omitted early noise (variant *iii*) (Fig. S1 *B–E*). Furthermore, examining the noise parameters in the full model (*i*) revealed that late noise was significantly higher than early noise, with the latter having a negligible magnitude (Fig. S1*F*). After verifying our fitting method with regards to parameter recovery (Fig. S4 *B* and *C*, and *SI Methods*), we found a positive correlation between participants' estimated late noise and selective gating in all experiments (experiments 1–4: $r = 0.63$, $P < 0.001$, $r = 0.80$; $P < 0.001$, $r = 0.52$; $P < 0.006$; $r = 0.81$; $P < 0.001$; Fig. 4*A*, scatterplot for all experiments). We ruled out the possibility that this correlation occurred as an artifact of the parameter estimation method, by reporting no relationship ($r = 0.03$; $P = 0.502$) between estimated late noise and *w* in simulated datasets (Fig. 4*B* and *SI Methods*). This finding indicates that selective gating has an adaptive role, being adjusted to each individual's late noise levels, in the service of reward-maximizing decisions.

Discussion

Violations of the axioms of choice rationality have been exhaustively documented in the decision-making literature (2, 7, 9–11, 24). Numerous studies have found that the subjective value of an economic prospect depends not only on its own attribute values but also on the irrelevant context provided by competing alternatives. Although this relative (rather than absolute) valuation schema is incorporated in descriptive theories of choice (3), and is reflected in neural signals recorded both from single cells (30) and whole-brain areas (31), it currently lacks a plausible normative explanation. Instead, violations of choice rationality appear to have negative repercussions, potentially leading to a continuous drain on resources, for example to what economic theory knows as a "money pump" (6).

Here, we argue that choice irrationality occurs because of selective integration, a policy that explicitly discards some information about the rival choice alternatives but paradoxically maximizes reward in the face of decision noise. Selective integration builds on an established framework for understanding both perceptual (15, 17) and economic decisions (16, 32), in which momentary decision values (e.g., sensory samples of a noisy stimulus, or attributes values for an economic prospect) are accumulated in parallel for two or more alternatives, corrupted by noise that could arise either during encoding or during information integration (28). The additional assumption of the model is that, where attributes compete locally, the winner can be integrated with relatively higher gain. Thus, when contemplating the (excellent) weather in Bali, the (reasonable) weather in rival Berlin appears poor by comparison, and does not drive a positive evaluation of a Berlin holiday as much as it should.

Selective integration predicts that violations of transitivity will occur when choice alternatives differ circularly in their number of winning attributes. Here, we verified this prediction empirically, showing that humans performing a magnitude discrimination task make intransitive choices about alternatives with equal cumulative value. Normalization or heuristic models could explain

PSYCHOLOGICAL AND COGNITIVE SCIENCES

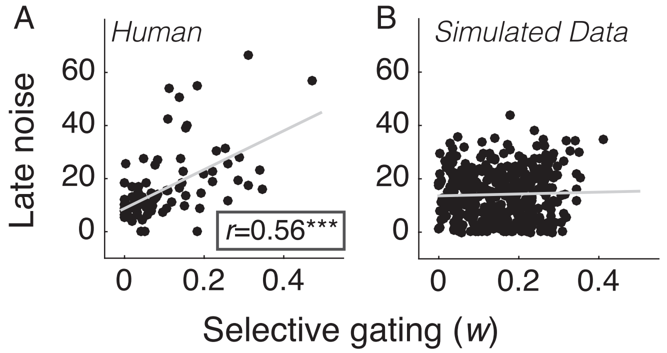

Fig. 4. Relationship between selective gating and late noise. (*A*) Estimated selective gating parameters for each individual (circles) in all four experiments (n = 93) plotted against estimated late noise parameters. (*B*) Same as *A*, but for simulated data (see Fig. S4 *B* and *C* and *SI Methods*). Gray curves are linear regression lines. ***P < 0.001.

a tendency to choose a frequently winning alternative and hence intransitive choices in our task, but they fail to explain other aspects of the data such as why participants prefer high-variance to low-variance alternatives. It has been recently shown that most of the past instances of intransitivity could have been caused by spurious factors and not necessarily by decision processes that violate economic rationality (20, 22). Our well-controlled psychophysical task together with computational modeling suggests that the systematic intransitivity reported here is not caused by spurious factors but by the tendency to integrate information selectively.

In sharp contrast to the mainstay in decision theory—which considers reward-rate optimal the algorithm that accumulates all available information without loss (17)—selective integration maximizes accuracy (and subsequent rewards) in the face of heightened late noise. The model achieves so by implicitly exaggerating monetary differences (i.e., via weakening the losing input), inflating the average accumulated difference between the correct and incorrect alternative. Fitting the model to human data revealed a strong positive correlation between late noise and the tendency to integrate selectively. This correlation suggests that observers use selective integration as a compensatory mechanism to alleviate the potentially negative impact of elevated accumulation noise. Violations of the axioms of choice rationality emerge as a side effect of this policy and occur when alternatives are structured in unusual ways [e.g., when the number of winning attributes differ circularly or when dominated decoys are introduced in the choice set (9)].

It is conceivable that explicit and equal increase of the gain of processing for both inputs, if strong enough, could outperform selective integration by cancelling out late noise. In our analyses, however, we assumed that such direct and unbounded gain amplification is not plausible because organisms operate within computational and metabolic constraints (33). Due to these constraints, even under conditions of increased vigilance behavioral and neural variability perseveres (34), indicating that a portion of internal noise is virtually irreducible. The way our model acts against this (otherwise-irreducible) noise presents a paradox for decision theory analogous to "less-is-more" effects in other domains (35), whereby ignoring part of the available information leads to better performance. Thus, in contrast to normalization theories (23, 24), in which accuracy and metabolic efficiency trade off against each other, selective integration increases choice accuracy while reducing the cost of information processing.

Our account of choice irrationality is not only normatively motivated but also builds on well-established psychological and neural principles. First, as in models of selective attention and visual search (36), our explanation incorporates selective processing. Although selective processing has been recently added to an influential evidence accumulation model to explain the increased accumulation rate for visually fixated alternatives (32), our approach differs in that gain modulation is determined by the value of the incoming information rather than merely by the (random) locus of fixation. Second, our account is biologically plausible, building upon two widely accepted neurobiological facts: that decisions are realized in a hierarchy of cortical layers and that processing at each layer is corrupted by independent neuronal noise (29, 37). It is the distributed and noisy nature of neural information processing that allows nonnormative choice algorithms, such as selective integration, to practically outperform the normative benchmark.

Why humans make irrational choices has puzzled economists and psychologists for decades. The findings described here suggest that violations of choice rationality are a natural consequence of selective gating—a processing bottleneck that discounts locally weaker samples when evidence is accumulated over time. We demonstrated that this bottleneck could protect decisions from the pernicious influence of late noise—that arising downstream from the input representation stage. Such late noise may be an indispensable feature of neural computation, perhaps because it promotes learning and exploratory behavior (38). Fitting selective integration to human choices, we indeed showed that selective gating was stronger in those individuals with higher late noise. This finding calls into question the longstanding argument that humans are irrational because they lack the computational resources to engage in effortful executive processes and fall back instead on less costly, intuitive strategies or heuristics (12). We suggest instead that apparently irrational choices may stem from an evolutionary pressure for reward-maximizing decisions, realized in a hierarchy of noisy cortical layers (29). This calls for a broader theory of ecological rationality (39) that is bounded by neurophysiological constraints.

Methods

Participants. Ninety-three participants (42 females; age range: 18–50; $N_1 = 28$, N_2 = 17, N_3 = 27, and N_4 = 21 in experiments 1–4, respectively) were recruited from Oxford University (experiments 1–3) and Warwick University (experiment 4) participant pools and gave informed consent to take part. All participants reported normal or corrected-to-normal vision and no history of neurological or psychiatric impairment. The experimental procedures were approved by the Oxford University Medical Sciences Division Ethics Committee (approval no. MSD/IDREC/C1/2009/1) and Warwick University Humanities and Social Sciences Research Ethics Sub-Committee (approval no. 83/14-15:DR@W). Participants received £8/h for their participation and a bonus of £15 that was subject to task performance.

Task. In all experiments, participants viewed two streams of bars of varying height presented simultaneously left and right from a central fixation point. Bar height was described as indicating the scores of two job candidates on different dimensions. In all experiments, participants were instructed that all dimensions are equally important. The dimensions were explicitly specified (i.e., intelligence, motivation, experience of a job candidate) and explicitly announced during stimuli presentation via changes in the color of a rectangular outline only in experiment 4 (*SI Methods*). After a fixed number of pairs of bars presented at a fixed rate, participants were asked to choose which stream (candidate), the left or the right, had on average higher bars (scores). Participants received partial feedback in experiments 1 and 4, and full feedback in experiments 2 and 3 (*SI Methods*) in 8 (experiments 1–3) or 18 (experiment 4) blocks (each lasting less than 10 min). At the end of the experiment, participants viewed their average accuracy on the screen, and if it fell within the 85th percentile of the cohort, they received a bonus of £15. A detailed description of the visual stimuli and trial time course is provided in *SI Methods*.

Outline of Experimental Conditions. Experiment 1 consisted of nine conditions that differed in the way the two sequences were constructed. We classify the different conditions into cyclic (three conditions), increment (two conditions), and standard (four conditions). Cyclic trials were constructed based on a set of sequences that resembled the *A*–*B*–*C* alternatives in Fig. 1*B*. Cyclic trials were divided in three conditions: (*i*) *A* vs. *B*, (*ii*) *B* vs. *C*, and (*iii*) *A* vs. *C*. There were n unique A_j–B_j–C_j triplets ($j = 1\ldots n$), with each triplet yielding one set of cyclic (*i*–*iii*) trials. Twelve participants performed a short version of the task with cyclic trials being generated by n = 40 unique triplets, whereas 16 participants performed a longer version with n = 60 (*SI Methods*). In each A_j–B_j–C_j triplet, the three sequences had identical values, but their order was reshuffled as per the example in Fig. 1*B* (i.e., B_j was created via a right circular shift of A_j, whereas C_j via a right circular shift of B_j ; see *SI Methods*). The two increment conditions were created by modifying *i* above, by increasing by 6 pixels either the height of all bars of A_j (*A*+ vs. *B*) or all bars of B_j (*A* vs. *B*+). Finally, in standard trials, the two alternatives consisted of normally distributed values. Two levels of mean and variance (high/low) were resampled with replacement, yielding overall four conditions. Experiments 2 and 3 consisted only of increment and standard trials (six conditions). Experiment 4 was similar to experiment 1 and had all nine conditions. However, cyclic trials in experiment 4 were constructed from a unique (per participant) *A*–*B*–*C* triplet, with each *A*–*B*–*C* sequence having the very same dimensional values in the whole experiment. The dimensions in experiment 4 were explicitly specified and announced alongside the presentation of the bars. Details about the experimental conditions, the sequences construction, and the number of trials per condition in the four experiments are provided in *SI Methods*.

Selective Integration. Two accumulators (Y_A and Y_B) integrate the attribute values (i.e., pixels representing the heights of the two bars) of the two sequences (*A* and *B*) according to the following difference equations:

$$\begin{aligned} Y_A(t) &= (1-\lambda)\cdot Y_A(t-1) + I_A(t) + \xi\cdot\zeta_A(t), \\ Y_B(t) &= (1-\lambda)\cdot Y_B(t-1) + I_B(t) + \xi\cdot\zeta_B(t). \end{aligned} \quad [1]$$

In the above, *t* denotes the current discrete time step, λ is integration leak, ξ is late noise, and $\zeta_{A,B}(t)$ are random standard Gaussian samples independent of each other and across *t*. The leak parameter was introduced to capture the recency-weighting profile (Fig. S1*A*) that was obtained in all experiments (see also ref. 14).

The two accumulators are initialized at 0:

$$Y_A(0) = Y_B(0) = 0.$$

The momentary inputs to the accumulators, $I_{A,B}(t)$, are defined as follows:

$$\begin{aligned} I_A(t) &= \theta(X_A(t) - X_B(t))\cdot X_A(t), \\ I_B(t) &= \theta(X_B(t) - X_A(t))\cdot X_B(t). \end{aligned} \quad [2]$$

The gain function θ is a step function defined as follows:

$$\theta(x) = \begin{cases} 1, & \text{if } x > 0 \\ w, & \text{if } x < 0 \end{cases}, \quad [3]$$

with *w* in [0, 1] being the selective gating parameter. Finally, $X_{A,B}(t)$ correspond to the incoming stimuli corrupted by internal (early) noise:

$$\begin{aligned} X_A(t) &= \sigma\cdot\rho_A(t) + s_A(t), \\ X_B(t) &= \sigma\cdot\rho_B(t) + s_B(t), \end{aligned} \quad [4]$$

where σ is the early noise parameter, $\rho_{A,B}(t)$ are random standard Gaussian samples independent of each other and across *t*, and $s_{A,B}(t)$ are the presented stimuli on time step *t* (bar heights, in pixels).

At the end of stimuli presentation (e.g., $t = 9$, in experiment 1), the model chooses sequence *A* if $Y_A(t) > Y_B(t)$, sequence *B* if $Y_B(t) > Y_A(t)$, and randomly between *A* and *B* if $Y_A(t) = Y_B(t)$. Overall, the full selective integration model has four free parameters: leak (λ), early noise (σ), late noise (ξ), and selective gating parameter (*w*). Noise parameters are expressed in pixels.

Model-fitting procedures and model parameters in the various simulations reported in the main text are given in *SI Methods*.

ACKNOWLEDGMENTS. This work was supported by a British Academy/Leverhulme Trust award (to K.T.); European Research Council (ERC), Economic and Social Research Council, Research Councils UK, and Leverhulme Trust awards (to N.C.); a German–Israeli Science Foundation grant and a Leverhulme Trust Visiting Professorship (to M.U.); and an ERC award (to C.S.).

1. Keeney RL, Raiffa H (1976) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (Wiley, New York).
2. Tversky A (1972) Elimination by aspects: A theory of choice. *Psychol Rev* 79(4): 281–299.
3. Vlaev I, Chater N, Stewart N, Brown GD (2011) Does the brain calculate value? *Trends Cogn Sci* 15(11):546–554.
4. Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci* 9(7):545–556.
5. Von Neumann J, Morgenstern O (1947) *Theory of Games and Economic Behavior* (Princeton Univ Press, Princeton), 2nd Ed.
6. Allingham M (2002) *Choice Theory: A Very Short Introduction* (Oxford Univ Press, Oxford).
7. Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* 211(4481):453–458.
8. Shafir E (1993) Choosing versus rejecting: Why some options are both better and worse than others. *Mem Cognit* 21(4):546–556.
9. Huber J, Payne JW, Puto C (1982) Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *J Consum Res* 9(1):90–98.
10. Trueblood JS, Brown SD, Heathcote A, Busemeyer JR (2013) Not just for consumers: Context effects are fundamental to decision making. *Psychol Sci* 24(6):901–908.
11. Tversky A (1969) Intransitivity of preferences. *Psychol Rev* 76(1):31–48.
12. Kahneman D (2011) *Thinking, Fast and Slow* (Allen Lane, London).
13. Simon HA (1982) *Models of Bounded Rationality* (MIT Press, Cambridge, MA).
14. Tsetsos K, Chater N, Usher M (2012) Salience driven value integration explains decision biases and preference reversal. *Proc Natl Acad Sci USA* 109(24):9659–9664.
15. Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
16. Busemeyer JR, Townsend JT (1993) Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychol Rev* 100(3):432–459.
17. 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(4):700–765.
18. Wald A (1973) *Sequential Analysis* (Dover Publications, New York).
19. Roe RM, Busemeyer JR, Townsend JT (2001) Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychol Rev* 108(2):370–392.
20. Regenwetter M, Dana J, Davis-Stober CP (2011) Transitivity of preferences. *Psychol Rev* 118(1):42–56.
21. Thurstone LL (1954) The measurement of values. *Psychol Rev* 61(1):47–58.
22. Regenwetter M, Dana J, Davis-Stober CP (2010) Testing transitivity of preferences on two-alternative forced choice data. *Front Psychol* 1:148.
23. Louie K, Khaw MW, Glimcher PW (2013) Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci USA* 110(15):6139–6144.
24. Soltani A, De Martino B, Camerer C (2012) A range-normalization model of context-dependent choice: A new model and evidence. *PLoS Comput Biol* 8(7):e1002607.
25. May KO (1954) Intransitivity, utility, and the aggregation of preference patterns. *Econometrica* 22(1):1–13.
26. Todd PM, Gigerenzer G (2000) Précis of simple heuristics that make us smart. *Behav Brain Sci* 23(5):727–741, discussion 742–780.
27. Renart A, Machens CK (2014) Variability in neural activity and behavior. *Curr Opin Neurobiol* 25:211–220.
28. Brunton BW, Botvinick MM, Brody CD (2013) Rats and humans can optimally accumulate evidence for decision-making. *Science* 340(6128):95–98.
29. Servan-Schreiber D, Printz H, Cohen JD (1990) A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science* 249(4971):892–895.
30. Peters J, Büchel C (2010) Neural representations of subjective reward value. *Behav Brain Res* 213(2):135–141.
31. Hunt LT, Dolan RJ, Behrens TE (2014) Hierarchical competitions subserving multi-attribute choice. *Nat Neurosci* 17(11):1613–1622.
32. Krajbich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13(10):1292–1298.
33. Summerfield C, Tsetsos K (2015) Do humans make good decisions? *Trends Cogn Sci* 19(1):27–34.
34. Hunt LT (2014) What are the neural origins of choice variability? *Trends Cogn Sci* 18(5):222–224.
35. Gigerenzer G, Goldstein DG (1996) Reasoning the fast and frugal way: Models of bounded rationality. *Psychol Rev* 103(4):650–669.
36. Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222.
37. Siegel M, Donner TH, Engel AK (2012) Spectral fingerprints of large-scale neuronal interactions. *Nat Rev Neurosci* 13(2):121–134.
38. Engel TA, Chaisangmongkon W, Freedman DJ, Wang XJ (2015) Choice-correlated activity fluctuations underlie learning of neuronal category representation. *Nat Commun* 6:6454.
39. Gigerenzer G, Brighton H (2009) Homo heuristicus: Why biased minds make better inferences. *Top Cogn Sci* 1(1):107–143.
40. Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10(4):433–436.
41. Regenwetter M, et al. (2014) QTest: Quantitative testing of theories of binary choice. *Decisions* 1(1):2–34.
42. Regenwetter M, Dana J, Davis-Stober CP, Guo Y (2011) Parsimonious testing of transitive or intransitive preferences: Reply to Birnbaum (2011). *Psychol Rev* 118(4): 684–688.
43. Bogacz R, Cohen JD (2004) Parameterization of connectionist models. *Behav Res Methods Instrum Comput* 36(4):732–741.