Evidence integration and decision-confidence are modulated by stimulus consistency

Moshe Glickman¹, Rani Moran^{2,3} and Marius Usher^{1,4}

¹School of Psychology, University of Tel Aviv, Tel Aviv, Israel; ²Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, 10-12 Russell Square, London WC1B 5EH, United Kingdom; ³Wellcome Centre for Human Neuroimaging, University College London, London WC1N 3BG, United Kingdom;

⁴Sagol School of Neuroscience, University of Tel Aviv, Tel Aviv, Israel

Abstract

Evidence-integration is a normative algorithm for choosing between alternatives with noisy evidence, which has been successful in accounting for a vast amount of behavioral and neural data. However, this mechanism has been challenged as tracking integration boundaries sub-serving choice has proven elusive. Here we first show that the decision boundary can be monitored using a novel, model-free behavioral method, termed *Decision-Classification Boundary*. This method allowed us to both provide direct support for evidence-integration contributions and to identify a novel integration-bias, whereby incoming evidence is modulated based on its consistency with evidence from preceding time-frames. This consistency bias was supported in three cross-domain experiments, involving decisions with perceptual and numerical evidence, which showed that choice-accuracy and decision confidence are modulated by stimulus consistency. Strikingly, despite its seeming sub-optimality, this bias fosters performance by enhancing robustness to integration noise. We argue this bias constitutes a new form of micro-level, within-trial, confirmation bias and discuss implications to broad aspects of decision making.

Corresponding Author:

Moshe Glickman

School of Psychological Science, Tel Aviv University, Tel Aviv, Israel

Email: mosheglickman345@gmail.com

Introduction

Evidence integration to boundary is a normative mechanism for evidence-based decisions, which provides the fastest mean response-time (RT) for a target accuracy rate (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Gold & Shadlen, 2002; Moran, 2015; Wald, 1947) and accounts for an impressive amount of behavioral and neural choice data (see Forstmann, Ratcliff, & Wagenmakers, 2016 for a review). For example, integration-to-boundary models (Bogacz et al., 2006; Brown & Heathcote, 2008; Gold & Shadlen, 2001; Ratcliff & McKoon, 2008; Teodorescu & Usher, 2013; Usher & McClelland, 2001; Vickers, 1970) provide a parsimonious account for the shape of choice-RT distributions of correct and incorrect responses as a function of stimulus difficulty, as well as for the well-known speed accuracy trade-off (Wickelgren, 1977). Moreover, integration-to-boundary models are supported by the monitoring of neural activation in brain decision areas during choice tasks (Gold & Shadlen, 2007; Mulder, van Maanen, & Forstmann, 2014; but see Latimer, Yates, Meister, Huk, & Pillow, 2015).

Despite this strong support, however, in behavioral studies the decision-boundary remains a theoretical concept, which is usually inferred from the data via model fitting and depends on specific model assumptions. Moreover, this framework can be challenged by alternative non-integration mechanisms, such as heuristics based on the detection of a single high-value sample, which can account for many of these choice-patterns (Ditterich, 2006; Watson, 1979; see also discussion in Stine, Zylberberg, Ditterich, & Shadlen, 2020). It is thus imperative to validate the *integration* assumption and to directly monitor the decision boundary. While one can support integration to boundary via quantitative model comparison (Stine et al., 2020), such comparisons are limited by specific and somewhat arbitrary parametric assumptions (e.g., a particular temporal collapsing regimen of a choice threshold). Furthermore, recent findings indicate a variety of biases in the sampling and weighting of evidence (Bronfman et al., 2015; Cheadle et al., 2014; Dotan, Meyniel, & Dehaene, 2018; Glickman, Tsetsos & Usher, 2018; Glickman et al., 2019; Gluth, Kern, Kortmann, & Vitali, 2020; Rollwage et al., 2020; Talluri et al., 2018; Urai, de Gee, Tsetsos & Donner, 2019). Potential biases of this kind should be taken into account when assessing whether evidence is accumulated. Here, we addressed these issues using a novel behavioral and model-agnostic method, which we term *Decision-Classification Boundary* (DCB).

When evidence integration is perfect (i.e., free of distortions and biases), the DCB recovers the decision boundary. More broadly, however, DCBs provide a novel behavioral signature — a benchmark for evaluating biases in evidence-integration. Applying this method to data from three experiments across choice domains (numerical cognition and perception), we find strong support for evidence-integration over heuristic non-integration models. Furthermore, we demonstrate an important new factor modulating evidence accumulation — *stimulus-consistency*, corresponding to an increased relative weighting of pieces of evidence preceded by information supporting the same choice-alternative, resulting in a type of *momentary confirmation-bias* (Bronfman et al., 2015; Rollwage et al., 2020; Talluri et al., 2018), which operates during (rather than after) a decision. Importantly, this mechanism achieves an important computational advantage: it makes the decision robust to late (non-encoding) noise (Spitzer et al., 2017; Tsetsos et al., 2016).

We start with a description of our experimental design (showing it is possible to extract behavioral signatures of integration to boundary), followed by a computational section that presents the model-free DCB method. We then apply this method to the data from two experiments, focusing on the dependency of the DCB on stimulus consistency and resort to computational modeling to specify the stimulus-consistency mechanism. Finally, we present the results of a third experiment designed to test this mechanism.

Results

Experimental design and behavioral signature of integration to boundary

In two experiments (Exp. 1, reported in Glickman & Usher, 2019 and Exp. 2, novel data), participants were presented with sequences of pairs of stimuli (two digit numbers or bars, respectively; see Fig. 1A-B), which were sampled from two overlapping Normal distributions (see Experimental Methods). The sequences were presented at a rate of 2 pairs/sec (numbers) and 5 pairs/sec (bars), and were terminated by the participant's response. The task was to select the sequence that corresponded to the higher generating distribution (red Gaussian in Fig 1A-B). Between trials, we manipulated stimulus difficulty by varying the separation between the Normal distributions (see Experimental Method for further details). Twenty-seven participants performed 500 trials in Exp.1 and thirty participants performed 480 trials in Exp.2.

Figure 1. (A-B) Experimental paradigms; participants are presented with pairs of numerical values (Exp. 1) or bars (Exp. 2) sampled from two overlapping Normal distributions, and are asked to choose which sequence was drawn from a distribution with a higher mean. The presentation is terminated by the decision of the participant. (C) The accumulated evidence of all participants (blue line) as a function of decision-time; grey lines correspond to accumulated evidence of individual participants. The black, green and magenta dashed lines correspond to the random time model, and to the value and difference

heuristics, respectively. (D) Same as (C) for Exp. 2. (E-F) The collapsing-boundaries obtained in Exp. 1 (E) and 2 (F); Solid purple lines correspond to the boundaries generated using the group mean parameters, grey lines correspond to the boundaries of the individual participants.

In previous work, we showed that the choices in Exp. 1 (numerical evidence) support integration of evidence to a collapsing boundary and excluded a set of non-integration models (Glickman & Usher, 2019). Here, we extend this analysis to the data in the perceptual domain (Exp. 2). The blue lines (group data) and thin grey lines (individual subjects) in Fig 1 C-D are obtained by integrating the trial-by-trial stimulus evidence until the decision moment (i.e., the cumulative sum of differences between the sequence), and averaging across trials for each RT (see Glickman & Usher, 2019 and Suppl. A for further details). These lines show a mildly decreasing pattern (Fig 1C: $b = -2.86$, $t = -7.08$, $p < .001$, and Fig 1D: $b = -0.86$, $t = -3.18$, $p = .001$), which is the behavioral signature of integration to a collapsing boundary (see Suppl. A). These results are consistent with model fitting using a Weibull parametrization of the boundary (Hawkins et al., 2015), which also indicates a collapsing boundary (Fig. 1E-F). Critically, the descending slopes of integrated evidence rule out non-integration strategies, such as: i) random-timer – response time is determined by a process that is exogenous to the integration of evidence (black dashed line), ii) value-cutoff – observers choose the sequence in which a number exceeding some predetermined threshold first appears (green dashed line), and iii) difference-cutoff – observers choose based on the first frame in which the difference between the numbers exceeds a predetermined threshold (magenta dashed line). All of these non-integration models predict that the integrated evidence increases (rather than decreases) with the number of samples (see Computational Methods for further details about these strategies). This is because, if the stopping rule is independent of the evidence integrated so far, longer decision trials necessarily accumulate more evidence. While these results provide support for integration to boundary, the actual shape of the boundary trajectory (Fig. 1E-F) is only extracted via model fitting, as the evidence-integration lines (Fig 1 C-D, blue and thin grey lines) are systematically biased by the accumulation of noise during the trial. In the following, we present a novel model-free method to estimate the decision boundary via a

classification boundary curve, which is more robust to accumulation noise and reconstructs the actual shape of the decision boundary.

Model-Free extraction of decision boundaries: the boundary classification curve

We simulated synthetic data based on the experimental task we used in our experiments (see Fig. 1A-B and Computational Methods), in which sequences of values are sampled from two overlapping Normal distributions. It is assumed that the subjects integrate a noisy version of the evidence at each frame, according to the following difference equation:

$$X(t) = X(t - 1) + \mu(t) + \varepsilon(t), \varepsilon \sim N(0, \sigma^2) \quad \text{Eq. (1)}$$

where $X(t)$ is the accumulated differences between the sequences at time t , $\mu(t)$ is the difference between the samples at time t , and $\varepsilon(t)$ is a temporally-independent random internal Gaussian noise generated by the encoding system, which is independent from the evidence-sampling noise (Fig 1A-B). We also assumed that a response is triggered when the integrated noisy evidence reaches one of two symmetric decision boundaries. Two types of boundaries were used in these simulations: fixed and collapsing (Hawkins et al., 2015). For the collapsing boundary, we follow Hawkins et al. (2015) to parameterize the boundary using Weibull functions (see Computational Methods). In each simulation, we ran 10,000 trials, in which (after sampling values from the two distributions and integrating subject to noise), we record the choice, as well as the input sequences (without the internal noise) sampled until the decision was made.

Based on this data, we reconstruct the boundary at each time frame using a method based on Linear Discriminant Analysis (LDA; Fisher, 1936; McLachlan, 2004), which generates boundary classification curves. These curves are obtained by applying LDA to the integrated-evidence excluding the random internal noise, defined as:

$$Y(t) = Y(t - 1) + \mu(t) \quad \text{Eq. (2)}$$

The LDA was applied to the integrated-evidence ($Y(t)$) of the observer at each time frame, across all trials in the experiment, so as to classify the action at each time frame to one of three categories: choose alternative A (Fig. 2A-B, blue distributions), choose alternative B (Fig. 2A-B, green distributions), or continue sampling (Fig. 2A-B, orange distributions).

The classification boundary curve (Fig. 2A-B, red line) best separates the different classes (see Computational Methods). Note that in the absence of internal noise, the true decision boundary would separate perfectly these categories. However, in the presence of internal noise, which may increase the values of integrated evidence that lie below the boundary (blue areas), or vice versa (orange areas above the boundary), there is some unavoidable overlap.

The results are illustrated in Fig. 2A-B, which shows that the DCBs (red solid lines) recover quite accurately the generating boundaries (black dashed lines). In particular, the extracted boundary is temporally constant or decreases as a function of time, when the generating boundary is flat or collapsing, respectively. Notably, the quantitative agreement between the generating and recovered boundaries is high. Note that the model-free boundary extraction method makes no *a priori* parametric assumptions about the shape of the model boundary. As shown in Fig. 2C-D, for synthetic data generated using the non-integration-to-boundary value (or difference) cutoff heuristics, the three distributions of evidence – trials in which the model continues sampling (orange) and trials in which response has been made (blue and green) – totally overlap. Consequently, the decision classification curve algorithm fails to classify the three classes of trials based on integrated evidence. Thus, the presence of stable boundary classification curves (Fig. 1A-B) provides strong support against (non-integration) cutoff models. In both our data sets, we find stable classification curves (see Fig. 3E-F, red lines, for representative example subjects) that imply a collapsing boundary, which are consistent with model fitting (Fig. 3E-F, black dotted line), but make no parametric assumption on the form of this boundary (see Suppl. C for DCB of all participants).

Figure 2. (A-B) Illustration of the boundary extraction for a decision simulated using a Diffusion model with fixed (A) or collapsing (B) boundary. The black dashed line corresponds to the original boundary with which the model was simulated and the red lines correspond to the model-free best-fitted DCB. The orange distributions correspond to data points within each trial in which the simulated participant continued sampling, while the blue and green distributions correspond

to frames in which the decision was terminated. The total area under the orange, blue, and green distributions was normalized for each frame. (E-F) Experimental data of a representative subject that participated in the numerical (E) and perceptual (F) experiments. Black dashed line corresponds to the model-based best-fitted boundaries, and the red lines correspond to the model-free best-fitted boundaries of the experiments, respectively.

Evidence integration is modulated by stimulus consistency

A more detailed examination of the choice data in both experiments shows that *stimulus consistency* – operationalized as the absolute value of the difference between the number of frames with evidence favoring each alternative divided by the total number of frames – has a critical impact on choices and RT, above and beyond the effect of total evidence. To illustrate, consider two trials with the same total evidence: Trial 1: (2, 3, 1, 4, total evidence = 10) and Trial 2: (6, -1, 8, -3, total evidence = 10). While both trials have a total evidence of 10 in favor of one of the alternatives, the evidence stream is more consistent in Trial 1 (consistency measure = $\frac{4-0}{4} = 1$) compared with Trial 2, where half of the evidence favoring one alternative whereas the other half favors the other (consistency measure = $\frac{2-2}{4} = 0$). Previous research has shown that stimulus consistency modulates decision confidence (Dotan et al., 2018). Here, we examine its impact also on accuracy, RT, as well as on the shape of decision boundaries. To this end, we conducted several mixed-model regression analyses (logistic for accuracy and linear for RT and confidence), in which we predicted trial by trial choice accuracy, RT, and confidence, using accumulated evidence and stimulus consistency as fixed factors and participants as random intercepts. The results, which are shown in Table 1, indicate that stimulus consistency improves accuracy and confidence and reduces RT, independent of the accumulated evidence.

Note that the "Difference in Evidence Directions" (DED) is only one measure of stimulus consistency. More complex forms that include temporal factors can also be construed. For example, the consistency bias may correspond to the size of the larger temporal cluster (LTC) with evidence in the same direction (i.e., the largest cluster of evidence; see Suppl. B for analysis showing that such an LTC-measure also predicts differences in accuracy, RT, and confidence independent from total evidence).

Table 1. Beta coefficients for predicting Accuracy, RT, and confidence in Exp. 1 & 2.

	Exp. 1 (Numerical Stimuli)		Exp. 2 (Perceptual Stimuli)		
	Accuracy	RT	Accuracy	RT	Confidence
Evidence	$\beta = 1.71^{***}$ $t = 27.90$	$\beta = 0.26^{***}$ $t = -28.21$	$\beta = 0.55^{***}$ $t = 16.91$	$\beta = -0.16^{***}$ $t = -20.42$	$\beta = 0.13^{***}$ $t = 12.14$
Stimulus Consistency	$\beta = 0.26^{***}$ $t = 5.69$	$\beta = -0.13^{***}$ $t = -14.41$	$\beta = 0.27^{***}$ $t = -9.40$	$\beta = -0.13^{***}$ $t = -16.78$	$\beta = 0.14^{***}$ $t = 13.96$

*** $p < .001$

Decision classification curves modulation by stimulus consistency and model comparison

Motivated by the results above, we examined if stimulus-consistency modulates the decision classification curve. Towards this aim, we extracted the DCB of high consistency trials (DED measure higher than 0.5) using the model-free method, and compared it with the DCB of all trials, which correspond to the most accurate estimation of the boundary (in the absence of a consistency-modulation). As illustrated below (Fig. 3 A-B; for two representative subjects; see Suppl. C for all subjects), in both experiments, the DCB is lower for the high consistency trials (Fig. 3 A-B; red lines) than for all trials (Fig. 3 A-B; blue lines; Exp. 1: $t(26) = 2.67$, $p = .01$, Exp. 2: $t(29) = 3.91$, $p < .001$; see Suppl. D for further details). Critically, we tested that applying the same boundary reconstruction procedure to synthetic data simulated from a full integration model, does not result in DCB differences based on the consistency of the evidence (see Suppl. D).

Figure 3. Results of Exp. 1 & 2. (A-B) Decision boundaries for all trials (blue) and high consistency trials (red) of representative participants in (A) Exp. 1 - numerical stimuli and (B) Exp. 2 - perceptual stimuli. (C-D) The stimulus-consistency model outperformed the full integration, preference consistency and selective integration models in Exp. 1 - numerical stimuli (C), as well as in Exp. 2 - perceptual stimuli (D).

These findings could suggest either that under high stimulus consistency the participants require less evidence to commit to a decision (i.e., choice boundaries are truly lower) or alternatively, that evidence integration itself may be modulated towards the same decision boundary (for high consistency trials); an evidence sample might be boosted when it comes after another sample supporting the same alternative, but evidence is integrated to the same boundary. As our interest is on the functional mechanism, we do not aim to distinguish these two variants here. Rather, we see the DCB as an empirical measure (which is data and not model-driven), against which models can be evaluated.

To specify the evidence-integration mechanism, we carried out a quantitative model comparison on the likelihood to account for the data (response and number of samples taken) for each trial of each participant in the two experiments (we used AIC as a measure of fit to include a model-complexity cost). We started with the perfect integration model (which assumes no systematic distortion of the evidence a subject integrates in each frame, i.e., integration is only corrupted by additive noise). Next, we examined a selective-integration (SI) mechanism that amplifies or diminishes, respectively, the stronger or weaker evidence within each frame across the two evidence streams (Tsetsos et al., 2016; Usher, Tsetsos, Glickman, & Chater, 2020; see Computational Methods). Critically, we also examined several variants of a “stimulus-consistency” model. In the simplest variant, the evidence is modulated solely based on whether evidence from the preceding frame is consistent (i.e., it points in the same direction) with the current frame. In a slightly more complex variant, the modulation magnitude increases linearly with the number of consistent frames and resets to baseline (no modulation) with every swap (we report here only the results of the more complex version, which provided better fit for the data). Finally, we also examined a *preference-consistency* model, wherein incoming evidence is modulated based on consistency with the total integrated evidence (thus reflecting momentary preference) up to that time point (see Computational Methods for details).

In each model, we allowed for collapsing boundaries, which, as illustrated in Fig 2C-D, capture well the shape of the decision boundary (and provide much better fits to the data compared to fixed boundaries; Glickman & Usher, 2019). Fig. 3C-D shows group and individual participant fit measures for four models: full integration, selective-integration, preference-consistency, and stimulus-consistency in Exp. 1 & 2. As illustrated, the most successful of the models was the stimulus-consistency version in which evidence increased at each consecutive frame in the same direction, followed by the SI-model.

Experiment-3: testing the stimulus-consistency effects

Because in our first two experiments, participants’ choices terminated the information stream, trial-consistency depended, at least partially, on participants’ responses. To address this limitation, and to directly test the impact of stimulus consistency on evidence-based choice and on decision confidence, a third experiment was designed. In this experiment,

the number of samples was fixed, and the stimulus consistency was manipulated independently from the total evidence. Sequences of eight number pairs were presented at a rate of 2/sec (as in Exp. 1) and participants were instructed to choose (at the end of the presentation) the higher average sequence. For each set of eight value-pairs, we generated two yoked trials, which consisted of the very same evidence content (across the eight time frames) for each sequence and differed only in temporal order of the values (i.e., the yoked trials varied in how the values on each side were shuffled). In the “consistent” trial condition (Fig. 4A, bottom panel), one evidence stream provided stronger evidence in 7 out of the 8 frames, whereas in the “inconsistent” trial condition (Fig. 4A, upper panel), each stream provided stronger evidence in 4 frames.

As shown in Fig. 4C, participants were more accurate in consistent than inconsistent trials, both in the easy, $t(21) = 6.15, p < .001$, and difficult conditions, $t(21) = 6.90, p < .001$. A similar pattern was also obtained for confidence, $t(21) = 4.76, p < .001$, for the easy condition and $t(21) = 4.19, p < .001$, for the difficult condition (Fig. 4D). Interestingly, the modulation of the confidence responses was different for correct and incorrect responses. Whereas for correct responses, confidence increases with stimulus consistency, $t(21) = 4.05, p < .001$ and $t(21) = 3.59, p = .002$, for the easy and difficult conditions, respectively, this pattern was not obtained for incorrect trials, $t(21) = -0.64, p = .52$ and $t(21) = .06, p = .95$. These findings indicate that metacognitive accuracy (confidence-resolution = $\text{confidence}_{\text{corrects}} - \text{confidence}_{\text{errors}}$) increases with stimulus consistency, both for easy trials: $t(21) = 3.68, p = .002$ and difficult trials: $t(21) = 2.59, p = .02$. As we show in Suppl. F, this effect can also be accounted for by the stimulus-consistency model by applying a signal-detection confidence approach.

Figure 4. The results of Exp. 3. (A) Illustration of consistent and inconsistent trials. Note that both trials have the same amount of evidence. L and R symbols correspond to momentary advantages of the left or right sequences, respectively. (B) The impact of the temporal position (evidence) on choice was assessed using logistic regression (see Computational Methods), and revealed a recency effect, which necessitated the use of a leak term in the model. (C) Choice accuracy as a function of difficulty (separation between the sampling distributions: easy vs. difficult) and consistency (difference in the directions of the evidence: consistent vs. inconsistent). The blue, red and green lines corresponds to the predictions of the leaky-integration, selective integration and stimulus consistency models, respectively. (D) Confidence as a function of difficulty and consistency for correct (blue lines) and incorrect (red lines) responses. Data are shown with solid lines and circle symbols and model predictions are shown with dashed lines and diamond symbols. Error bars are within-subjects SE.

We conducted quantitative model comparison for three types of models using the data from Exp. 3. Since the data also shows a recency pattern (see Fig. 4B and Computational Methods), we used leaky-integration (Ossmy, Moran, Donner & Usher, 2013; Roe et al., 2001; Usher & Mclelland, 2001) instead of full integration as our default model (with leak as a model parameter, accounting for decaying temporal weights). We compared a leaky-integration model (in which there is no distortion of the integrated evidence, other than the decaying temporal weights), an SI-model (which, additional to integration leak, gives higher weight to high values compared to low values, within each frame) and a stimulus-consistency model (from Exp. 1-2; Fig 3C-D). The results show that the stimulus-consistency model outperforms the full and selective integration models, and provides the best account for the data (Fig. 5).

Figure 5. Model Comparison for Exp. 3. The stimulus consistency model outperformed the leaky and selective integration models.

Interestingly, whereas both of the bias models are able to account for the modulation of accuracy by stimulus consistency in Fig 4C, the stimulus consistency bias model accounts for subtler temporal clustering effects in the data. For example, unlike the SI model, the stimulus-consistency model predicts that accuracy is modulated by the size of the largest cluster of evidence consistent with the correct choice (LTC; see Suppl. B for the impact of this measure in Exp. 1 & 2 and Suppl. E for data showing an association between the LTC-enhancement and the advantage of stimulus-consistency over the SI model in Exp. 3).

Stimulus-consistency and normativity

Normative considerations

Why do participants increase the relative weighting of pieces of evidence that are consistent with previous ones? At face value, this distortion introduces a gap between the accumulated and ‘real’ evidence and should reduce task performance. Indeed, this is the case in the absence of (or for low) integration noise. As illustrated in Fig. 6A-B, however, in the presence of high integration noise, the consistency modulation makes the mechanism more robust to the corrupting impact of this noise¹ (see cross-over between red and light blue lines in Fig 6A, so that for each level of noise, there is a consistency-modulation that optimizes performance; black dots in Fig 6B and dark blue line in Fig 6A). A similar robustness effect due to selective-integration was reported by Tsetos et al. (2016) for the SI-model (see also Spitzer et al., 2017 and Usher, Tsetos, Glickman & Chater, 2019). In both cases, the integration mechanism distorts the actual evidence by shifting the distribution of accumulated evidence towards the correct side (see center of the blue/red Gaussians in Fig 4C-D, while also making this distribution broader). While for low noise this is detrimental to performance, for high noise it is beneficial, as the shift helps to make the effect of additional noise less pronounced.

¹ Since we are looking at normative considerations, those simulations exclude integration-leak; adding it to all models does not change any of the results.

Figure 6. Consistency based modulation and normativity (A) accuracy (i.e., robustness to noise) as a function of noise, separately for the full integration (red line) and stimulus consistency models (blue lines). The accuracy of the stimulus-consistency model is presented for the model simulated using a fixed consistency parameter value of 10 (light blue) and for the model simulated using the optimal consistency parameter for each noise level (dark blue; see also panel B). (B) Accuracy of the stimulus-consistency model as a function of the consistency-modulation strength, for different levels of noise (σ curves). Black circles indicate consistency-values that maximize accuracy for a given level of accumulation noise. One can see that as the level of the noise increases (blue to red lines), so does the level of the consistency modulation needed to achieve optimal performance (black circles). (C-D) The distributions of the total accumulated evidence at the moment of response for the full integration model (blue) and stimulus-consistency model (red). The accuracy of the full integration model is higher than that of the stimulus-consistency model for low noise simulations (C), however the stimulus-consistency model is more accurate for high noise simulations (D).

Discussion

In the present study we examined the mechanism that human observers deploy when making decisions on rapid streams of stochastic evidence (perceptual and numerical). Using a novel behavioral model-agnostic method — the DCB curve — we showed that decision making is well characterized by integration to boundary rather than by non-integration heuristics, and that the boundary collapses with the passage of time (Fig. 2E-F; see also Balsdon, Wyart & Mamassian, 2019). Furthermore, we found that DCB curves constitute an informative behavioral benchmark for evaluating biases in evidence-integration. Strikingly, the full integration model failed to account for the lowered DCB for high consistency trials, indicating that integration of evidence was subject to systematic biases.

Motivated by previous studies, we focused on two types of evidence integration biases. The first is an attention bias, which affects the relative weight of evidence given to temporally simultaneous sources of evidence (Krajbich et al., 2010). In the selective-integration model, for example, on each frame the higher of the two values receives a higher weight than the lower one (Glickman et al., 2018; Luyckx et al., 2020; Tsetsos et al., 2012). The second bias involves the sequential impact of a frame on subsequent frames, whereby evidence that is consistent with previous frames receives higher weight than inconsistent evidence. Formal model comparisons, as well as analysis based on qualitative behavioral patterns, supported the superiority of consistency biases in accounting for our data. Notably, consistency affected not only decision accuracy, but also decision confidence (see also Dotan et al., 2018), such that consistent evidence facilitated high decision confidence even after controlling for the total amount of evidence. Notably, consistency did not merely exert a biasing influence on confidence but improved participant's meta-cognitive performance as measured by the resolution of confidence (i.e., the correlation between confidence and choice correctness). Indeed, confidence as a function of consistency increased for correct choices but remained constant for incorrect choices. An open question for future studies is whether the effects that consistency exerts on choice accuracy and meta-cognition are dissociable.

Previous research has reported sequential effects operating at the trial level. For example, a choice biases the interpretation of evidence in subsequent trials (Abrahamyan et al., 2016; Braun et al., 2018; Urai et al., 2019; Yu & Cohen, 2009). Similarly, a preliminary decision biases processing of additional post-choice evidence towards confirming the initial decision (Bronfman et al., 2015; Talluri et al., 2018; Rollwage et al., 2020), and decisions bias the strength-evaluation of pro-choice evidence that led to it (Jazayeri & Movshon, 2007; Luu & Stocker, 2018; Stocker & Simoncelli, 2008). Common to these studies is the assumption that these biases are triggered by the formation of a decision. Our findings, however, extend this notion by suggesting that a similar micro-level evidence integration bias operates during decision formation, prior to committing to a choice (see also Cheadle et al., 2014; Patai et al., 2020). In particular, we found that the best-fitting model was one in which the evidence is boosted for consistent evidence frames (this boost increasing with the number of consecutive consistent frames) and is reset to baseline when the first inconsistent piece of evidence is encountered. Despite its evidence distortion, we have shown that this mechanism has an adaptive function in the presence of integration noise. Since evidence samples that are consistent with their predecessors are more likely to carry stronger evidence in favor of the correct alternative, inflating their weight provides extra protection from the corrupting effect of accumulation noise, resulting in increased decision accuracy (Fig. 6). This ‘normativity hypothesis’ predicts that to the extent that consistency-based evidence integration is a controlled and adjustable strategy, consistency sensitivity effects will increase as a function of integration noise, for example, when one is performing a dual task or when working memory is loaded. We leave this interesting question for future studies (see Spitzer et al., 2017; for a parallel investigation in relating to selective-integration).

Our findings raise the intriguing hypothesis that confirmation biases are a form of consistency bias, whereby post-choice evidence inconsistent with pre-choice evidence is integrated less effectively than post-choice consistent evidence. Future studies should investigate whether, to what extent, and how, consistency and confirmation biases are related, by measuring both biases within participants and using a unified paradigm. Another interesting possibility is that the consistency of evidence supporting a decision might affect the extent of a confirmation bias. For example, choices that are based on more consistent

evidence may likely be more prone to confirmation bias, for example, due to the mediating effect of decision confidence (Rollwage et al., 2020). Future studies could also beneficially examine whether and how consistency bias is related to a broad range of individual traits such as the need for cognitive closure (Kruglanski & Webster, 1996), need for consistency, political radicalism (Rollwage et al., 2018), and dogmatism, or to psychiatric conditions such as OCD (Cavedini et al., 2006). We believe our paradigm provides an important advantage over current confirmation bias paradigms by addressing these questions. In current confirmation bias paradigms, which probe one's ability to revise initially-wrong decisions, validity might be jeopardized by demand characteristics (e.g., presenting oneself in a self-consistent manner). In contrast, the current approach eschews these concerns, since participants are not required to contradict or confirm their previous decisions.

In conclusion, our novel methods allowed us to validate critical aspects of the evidence accumulation process and to unravel the biases that affect it. Our findings contribute to a growing literature speaking to the notion that self-inflicted distortions of evidence are ironically adaptive in that they act to increase choice veracity by making it robust to noise. A critical next step is to study how these strategies are acquired and how they relate to puzzling behaviors such as confirmation bias, and to better characterize the environmental and psychological variables that affect strategy selection.

Methods

Experimental methods

Participants. Twenty-seven undergraduate from Tel-Aviv University (22 females) participated in Exp. 1², thirty in Experiment 2 (22 females), and twenty-three in Exp. 3 (17 females), all of which reported having normal or corrected-to-normal vision. The participants received course credit in exchange for taking part in the experiments, as well as a bonus fee ranging from 15 – 25 ILS, which was determined by their task performance. The experiment was approved the ethics committee at TAU.

Stimuli. The stimuli were consisted of pairs of numerical values (Exp. 1 & 3) or bars (Exp. 2) which were presented simultaneously (see Fig. 1A-B for illustration), with a presentation

² The data of Exp. 1 was taken from our previously published paper (Glickman & Usher, 2019).

rate of 2 Hz (Exp. 1 & 3) or 5 Hz (Exp. 2). Displays in Exp. 1 & 3 were generated by an Intel I7 personal computer attached to a 24" Asus 248qe monitor with a 144 Hz refresh rate, using 1920×1080 resolution graphics mode. Displays in Exp. 2 were generated by an Intel I3 personal computer attached to a 19" ViewSonic Graphics Series G90fB CRT monitor with a 60 Hz refresh rate, using 1024×768 resolution graphics mode. Responses were collected via the computer keyboard. Viewing distance was approximately 60 cm from the monitor.

Task and Design. Each trial in the experiments began with a fixation display which consisted of a black $0.2^\circ \times 0.2^\circ$ fixation cross (+) that remained on the screen for 1s. Then, pairs of numerical values (Exp. 1 & 3) or bars (Exp. 2) were presented sequentially to the participants, who were asked to decide which of the sequences has a higher mean. The presentation in Exp. 1 & 2 was terminated by the participants' response (free response protocol), while the presentation in Exp. 3 was terminated after eight samples (interrogation protocol). Responses were given by pressing the arrow keys (left/right arrow keys for the left/right sequences, respectively). In Exp. 2 & 3, after indicating the sequence with the higher mean, participants were also asked to indicate the confidence in their choice on a 1-6 scale.

Experimental Conditions. In all the experiments the samples were drawn from Gaussians distributions. Exp. 1 included two difficulty levels which were manipulated by varying the separation between the Gaussians: in the easy trials the means of the Gaussians were: $\mu_1=52$ vs. $\mu_2=44$, $\sigma=10$, and in the difficult trials the means were: $\mu_1=52$ vs. $\mu_2=48$, $\sigma=10$. Exp. 2 also consisted of two difficulty levels $\mu_1 = 52.5$ vs. $\mu_2 = 47.5$ and $\mu_1 = 51.5$ vs. $\mu_2 = 48.5$, as well as an orthogonal manipulation of the sequences variance: $\sigma_1 = 0.1167$ vs. $\sigma_2 = 0.07$. In Exp. 3 we orthogonally manipulated the difficulty and the consistency of the evidence. Difficulty was manipulated by increasing the separation between the Gaussians from $\mu_1=52$ vs. $\mu_2=48$, $\sigma=10$ (difficult trials) to $\mu_1=52$ vs. $\mu_2=44$, $\sigma=10$ (easy trials). Consistency was manipulated by sampling eight values from the high as well as from the low mean distributions. Then, to generate consistent trials we paired these values, so that in seven out of the eight pairs the stronger evidence was in favor of the higher mean distribution (see Fig. 4A, lower panel). To generate inconsistent trials, we shuffled the temporal order of the

same values and repaired them, so that only in four out of the eight pairs the stronger evidence was in favor of the higher mean distribution (Fig. 4A, upper panel).

Computational methods

Mixed effects models. The effect of accumulated evidence and stimulus-consistency on accuracy, RT confidence in Exp. 1 & 2 (Table 1) were estimated on a trial-by-trial basis using mixed model regression analyses. The regressions were implemented using Matlab's 'fitlme' and 'fitglm' functions with participants serving as random effects and with a free covariance matrix. The fixed effects variables were: i) accumulated evidence - the sum of differences between the two streams of evidence at the moment of response, and ii) stimulus-consistency - the absolute value of the difference between the number of frames with evidence favoring the two alternatives, normalized by the length of trial. Both variables were normalized using z-score transformations.

Choices in Exp. 1 & 2 (coded as 1-correct and 0-error) were predicted using logistic regressions, which in Wilkinson notation was:

$$\text{logit}(P_{\text{Choices}}) \sim (\text{accumulated evidence}) + (\text{stimulus consistency}) + (1|\text{subject})$$

RTs (Exp. 1 & 2) and confidence (Exp. 2) were predicted using linear regressions, which in Wilkinson notation were:

$$\begin{aligned} \text{RT} &\sim (\text{accumulated evidence}) + (\text{stimulus consistency}) + (1|\text{subject}) \\ \text{confidence} &\sim (\text{accumulated evidence}) + (\text{stimulus consistency}) + (1|\text{subject}) \end{aligned}$$

The exact same pattern of results reported in Table 1 was obtained if the accumulated evidence and stimulus-consistency were also included as random effects.

Model-free. The validity of the model-free method was tested by simulating 10,000 synthetic decisions using known (fixed or collapsing) boundaries, and examining the ability of the model-free method to accurately recover them. The values in all simulations (Fig. 2A-D) were sampled from $X \sim N(52, 15^2)$ and $Y \sim N(46, 15^2)$. The decision process in Fig. 2A-B was based on Eq. 1 with either fixed (Fig. 2A) or collapsing boundary (Fig. 2B). The fixed-boundary was characterized by a single boundary parameter ($c=50$), and the

collapsing-boundary was characterized by four parameters: intercept, shape, scale and asymptote parameters ($a=100$, $k=3$, $\lambda=4$, $a'=20$; see Computational Methods and Hawkins et al., 2015). The decision process in Fig. 2C-D was based on the value and difference cut-off heuristics with cut-offs of 70 and 20, respectively (see Computational Methods and Glickman & Usher, 2019).

The decision boundaries in fig. 2A-D were extracted by applying the LDA algorithm (Fisher, 1936; McLachlan, 2004) to the integrated evidence excluding internal noise (i.e., $Y(t)$, see Eq. 2) of each frame of all 10,000 synthetic trials. For each frame, each trial was classified as one of the following categories: choose alternative-A, choose alternative-B or continue sampling. Then, using the LDA, we extracted the planes that optimize the separation between different classes for each frame. We assumed that the upper and lower boundaries are symmetrical and therefore averaged both.

As mentioned in the main text, the internal noise causes an unavoidable overlap between the different classes. This overlap impairs the LDA ability to correctly extract the decision boundary, and is particularly evident in slow trials due to the accumulation of internal noise across time (Teodorescu, Moran, & Usher, 2016). Thus, in order to increase the robustness of the model-free method to internal noise, we constrained the boundary extraction of each frame by previous ones. To this end, we extracted the boundary based on two features: t and $Y(t)$, $1 \leq t \leq n$, as illustrated in the table below:

$Y(t)$	t	Response classification
$Y_{\text{trial 1}}(t = n)$	n	Choose B
$Y_{\text{trial 2}}(t = n)$	n	Continue sampling
...	...	...
$Y_{\text{trial 1}}(t = n-1)$	$n-1$	Continue sampling
$Y_{\text{trial 2}}(t = n-1)$	$n-1$	Continue sampling
$Y_{\text{trial 3}}(t = n-1)$	$n-1$	Choose A
...	...	...
$Y_{\text{trial 1}}(t = 1)$	1	Continue sampling
$Y_{\text{trial 2}}(t = 1)$	1	Continue sampling
$Y_{\text{trial 3}}(t = 1)$	1	Continue sampling
...	...	...

The LDA algorithm provide linear functions that separate the different classes from each other. To obtain the value of the boundary for the n -th frame, we computed the value of the separating linear functions for this frame.

Modelling Details

Integration to boundary models. We examined several integration-to-boundary models, all of them assumed integration of evidence based on the following formula:

$$X(t) = X(t-1) + \mu(t) + \varepsilon(t), \varepsilon \sim N(0, \sigma^2) \quad \text{Eq. (1)}$$

where $X(t)$ is the accumulated differences at time t , $\varepsilon(t)$ is a random Gaussian noise, which is independent from the evidence-sampling noise. The term $\mu(t)$ varied between the different models as follows:

i) **Full integration**

$$\mu(t) = V_L(t) - V_R(t) \quad \text{Eq. (2)}$$

where $V_L(t)$ and $V_R(t)$ are the samples drawn from the left and right distributions at time t , respectively (note that $V_L(t)$ and $V_R(t)$ include the sampling noise).

ii) **Stimulus consistency**

$$\mu(t) = V_L(t) - V_R(t) + \text{sign}(V_L(t) - V_R(t)) \cdot \theta \cdot i \quad \text{Eq. (3)}$$

where θ is a free parameter representing the enhancement given to pieces of evidence that are consistent with previous ones, and i counts the run of consistent values (starting at 0). For example, if the differences between the values are: 15, 20, 8, -15 and -25, then $i_{t=1} = 0, i_{t=2} = 1, i_{t=3} = 2, i_{t=4} = 0$ and $i_{t=5} = 1$.

iii) **Preference consistency**

$$\mu(t) = \begin{cases} (V_L(t) - V_R(t)) \cdot \theta, & \text{if } \frac{\text{sign}(V_L(t) - V_R(t))}{\text{sign}(X(t-1))} = 1 \\ V_L(t) - V_R(t), & \text{otherwise} \end{cases} \quad \text{Eq. (4)}$$

where θ is a free parameter representing the enhancement given to pieces of evidence consistent with the total accumulated evidence at time $t - 1$.

iv) **Selective integration**

$$\mu(t) = V_L(t) \cdot \frac{1}{1 + e^{-\theta(V_L(t)-V_R(t))}} - V_R(t) \cdot \frac{1}{1 + e^{-\theta(V_R(t)-V_L(t))}} \quad \text{Eq. (5)}$$

where θ is a free parameter affecting the magnitude of the selective gating (see Glickman et al., 2018).

All the integration models in Exp. 1 & 2 assume integration to a collapsing-boundary (see Glickman & Usher, 2019), modeled using a Weibull cumulative distribution function (Hawkins et al., 2015):

$$u(t) = a - [1 - \exp\left(-\left(\frac{t}{\lambda}\right)^k\right)] \cdot (a - a') \quad \text{Eq. (6)}$$

where $\pm u(t)$ are the upper/lower thresholds at time t , a/a' are the initial (intercept) and asymptotic values of the boundary, respectively, and λ/k are the scale and shape parameters of the Weibull function, respectively.

In Exp. 3 we used an interrogation paradigm, in which the probability of choosing each alternative was calculated using an exponential version of Luce's choice rule (Luce, 1959):

$$P(\text{Left}) = \frac{1}{1 + e^{-(\beta_0 + \beta_1(\sum_{t=1}^n \mu(t)))}} \quad \text{Eq. (7)}$$

$$P(\text{right}) = 1 - P(\text{Left})$$

where β_1 indicates the sensitivity of the model to the accumulated evidence, with an intercept of β_0 .

In addition, we examined whether the participants in Exp. 3 showed a recency bias, as reported in several previous studies which used an interrogation paradigm (Tsetsos et al., 2012; 2016). To this end, for each participant, we performed a temporal logistic regression analysis, in which we predicted the response of each trial based on the differences between

V_L and V_R at each frame, ranked by their temporal order. As shown in Figure 4B, the mean weight of samples 5-8 was higher than that of samples 1-4 ($t(21) = 9.97, p < .001$), indicating a recency bias. This motivated us to include a leak term that controls the extent to which earlier values are given less weight (Usher & McClelland, 2001). Thus, Eq. 7 was extended to the following form:

$$P(Left) = \frac{1}{1 + e^{-(\beta_0 + \beta_1((1-\lambda)^{n-t} \cdot \sum_{t=1}^n \mu(t)))}} \quad \text{Eq. (8)}$$

$$P(right) = 1 - P(Left)$$

where λ is the leak term.

Non integration to boundary models. We examine three models that did not assume integration of evidence decision boundary (Fig. 1C-D and Fig. 2C-D). The first model is the value-cutoff heuristic, which assumes that observers choose based on the detection of a single high-value sample. For example, if a participant uses a cut-off value of 70, then s/he will choose the sequences in which a value higher than 70 first appears. The second heuristic is the difference-cutoff heuristic, which assumes that observers choose based on the first frame in which the difference between the numbers exceeds a predetermined threshold (Glickman & Usher, 2019). In addition to these two heuristics, we examined a third model which we labeled a random-timer model. This model assumes integration of evidence based on Eq. 1, however, response-time is determined by an exogenous process. The value cut-off heuristic was simulated using thresholds of 70 (Exp. 1, Fig. 1C), 80 (Exp. 2, Fig. 1D) and 70 (model-free, Fig. 2C). The difference cut-off heuristic was simulated using thresholds of 20 (Exp. 1, Fig. 1C), 25 (Exp. 2, Fig. 1D) and 20 (model-free, Fig. 2D). Response-times of the random timer model were sampled from an ex-Gaussian distribution with $\mu = 3, \sigma = 0.5$ and $\lambda = 2/3$ (Exp. 1, Fig. 1C) and $\mu = 6, \sigma = 1$ and $\lambda = 3$ (Exp. 2, Fig. 2C). These values were chosen because they provided accuracy and response-times similar to the ones observed in the data of Exp. 1 & 2.

Optimization procedure. The free parameters of the computational models were fitted to the data (choices and decision-times) of each participant in Exp. 1 & 2 separately, using maximum likelihood estimation (MLE). For each trial, we simulated the different models

1,000 times for a given set of proposal parameters and calculated the proportion of trials in which the model choice and decision-time matched the empirical data. Denoting the proportion of match between the simulated and empirical data by p_i , we maximized the likelihood function $L(D|\theta)$ of the data (D) given a set of proposal parameters (θ), by:

$$L(D|\theta) = \prod_{i=1}^N p_i$$

To find the best set of proposal parameters we first used an adaptive grid search algorithm (see Tavares, Perona, & Rangel, 2017 for details) and then used the three best sets of proposal parameters as starting points to a Simplex minimization routine (Nelder & Mead, 1965).

Acknowledgments

This research was supported by grants to Marius Usher from the United States-Israel Binational Science Foundation (CNCRS = 2014612). We wish to thank Dr. Douglas Lee for his critical reading of the manuscript and helpful comments.

References

- Abrahamyan, A., Silva, L. L., Dakin, S. C., Carandini, M., & Gardner, J. L. (2016). Adaptable history biases in human perceptual decisions. *Proceedings of the National Academy of Sciences*, 113(25), E3548-E3557.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4), 700-765.
- Balsdon, T., Wyart, V., & Mamassian, P. (2020). Confidence controls perceptual evidence accumulation. *Nature communications*, 11(1), 1-11.
- Braun, A., Urai, A. E., & Donner, T. H. (2018). Adaptive history biases result from confidence-weighted accumulation of past choices. *Journal of Neuroscience*, 38(10), 2418-2429.

- Bronfman, Z. Z., Brezis, N., Moran, R., Tsetsos, K., Donner, T., & Usher, M. (2015). Decisions reduce sensitivity to subsequent information. *Proceedings of the Royal Society B: Biological Sciences*, 282(1810), 20150228.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive psychology*, 57(3), 153-178.
- Cavedini, P., Gorini, A., & Bellodi, L. (2006). Understanding obsessive-compulsive disorder: focus on decision making. *Neuropsychology review*, 16(1), 3-15.
- Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., De Gardelle, V., Castañón, S. H., & Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron*, 81(6), 1429-1441.
- Ditterich, J. (2006). Evidence for time-variant decision making. *European Journal of Neuroscience*, 24(12), 3628-3641.
- Dotan, D., Meyniel, F., & Dehaene, S. (2018). On-line confidence monitoring during decision making. *Cognition*, 171, 112-121.
- Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2), 179-188.
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E. J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual review of psychology*, 67, 641-666.
- Glickman, M., & Usher, M. (2019). Integration to boundary in decisions between numerical sequences. *Cognition*, 193, 104022.
- Glickman, M., Shari, O., Levy, D. J., Niebur, E., Stuphorn, V., & Usher, M. (2019). The formation of preference in risky choice. *PLoS computational biology*, 15(8), e1007201.

- Glickman, M., Tsetsos, K., & Usher, M. (2018). Attentional selection mediates framing and risk-bias effects. *Psychological science*, 29(12), 2010-2019.
- Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in cognitive sciences*, 5(1), 10-16.
- Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36(2), 299-308.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual review of neuroscience*, 30, 535-574.
- Gluth, S., Kern, N., Kortmann, M., & Vitali, C. L. (2020). Value-based attention but not divisive normalization influences decisions with multiple alternatives. *Nature Human Behaviour*, 1-12.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E. J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, 35(6), 2476-2484.
- Jazayeri, M., & Movshon, J. A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature*, 446(7138), 912-915.
- Krajibich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience*, 13(10), 1292-1298.
- Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: "Seizing" and "freezing.". *Psychological review*, 103(2), 263-283.
- Latimer, K. W., Yates, J. L., Meister, M. L., Huk, A. C., & Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244), 184-187.

- Luu, L., & Stocker, A. A. (2018). Post-decision biases reveal a self-consistency principle in perceptual inference. *Elife*, 7, e33334.
- Luyckx, F., Spitzer, B., Blangero, A., Tsetsos, K., & Summerfield, C. (2020). Selective integration during sequential sampling in posterior neural signals. *Cerebral Cortex*, 30(8), 4454-4464.
- Luce, R. D. (1959). Individual choice behavior. Oxford: Wiley.
- McLachlan, G. J. (2004). Discriminant analysis and statistical pattern recognition (Vol. 544). John Wiley & Sons.
- Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. *Psychonomic bulletin & review*, 22(1), 38-53.
- Mulder, M. J., Van Maanen, L., & Forstmann, B. U. (2014). Perceptual decision neurosciences—a model-based review. *Neuroscience*, 277, 872-884.
- Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The computer journal*, 7(4), 308-313.
- Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., & Donner, T. H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. *Current Biology*, 23(11), 981-986.
- Patai, E. Z., Foltynie, T., Limousin, P., Zrinzo, L., Bogacz, R., & Litvak, V. (2020). Conflict detection in a sequential decision task is associated with increased cortico-subthalamic coherence and prolonged subthalamic oscillatory response in the beta band. bioRxiv.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, 20(4), 873-922.

- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological review*, 108(2), 370-392.
- Rollwage, M., Dolan, R. J., & Fleming, S. M. (2018). Metacognitive failure as a feature of those holding radical beliefs. *Current Biology*, 28(24), 4014-4021.
- Rollwage, M., Loosen, A., Hauser, T. U., Moran, R., Dolan, R. J., & Fleming, S. M. (2020). Confidence drives a neural confirmation bias. *Nature communications*, 11(1), 1-11.
- Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour*, 1(8), 0145.
- Stine, G. M., Zylberberg, A., Ditterich, J., & Shadlen, M. N. (2020). Differentiating between integration and non-integration strategies in perceptual decision making. *Elife*, 9, e55365.
- Stocker, A. A., & Simoncelli, E. P. (2008). A Bayesian model of conditioned perception. In *Advances in neural information processing systems* (pp. 1409-1416).
- Tavares, G., Perona, P., & Rangel, A. (2017). The attentional drift diffusion model of simple perceptual decision-making. *Frontiers in neuroscience*, 11, 468.
- Talluri, B. C., Urai, A. E., Tsetsos, K., Usher, M., & Donner, T. H. (2018). Confirmation bias through selective overweighting of choice-consistent evidence. *Current Biology*, 28(19), 3128-3135.
- Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychological review*, 120(1), 1-38.
- Teodorescu, A. R., Moran, R., & Usher, M. (2016). Absolutely relative or relatively absolute: violations of value invariance in human decision making. *Psychonomic bulletin & review*, 23(1), 22-38.

- Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109(24), 9659-9664.
- Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, 113(11), 3102-3107.
- Urai, A. E., De Gee, J. W., Tsetsos, K., & Donner, T. H. (2019). Choice history biases subsequent evidence accumulation. *Elife*, 8, e46331.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological review*, 108(3), 550-592.
- Usher, M., Tsetsos, K., Glickman, M., & Chater, N. (2019). Selective Integration: An Attentional Theory of Choice Biases and Adaptive Choice. *Current Directions in Psychological Science*, 28(6), 552-559.
- Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, 13(1), 37-58.
- Wald, A. (1947). Foundations of a general theory of sequential decision functions. *Econometrica, Journal of the Econometric Society*, 279-313.
- Watson, A. B. (1979). Probability summation over time. *Vision research*, 19(5), 515-522.
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta psychologica*, 41(1), 67-85.
- Angela, J. Y., & Cohen, J. D. (2009). Sequential effects: superstition or rational behavior?. In *Advances in neural information processing systems* (pp. 1873-1880).

Supplementary Information

A) Behavioral signature of integration to boundary

We compared integration and a non-integration models by using the method described in Glickman & Usher (2019). Each model was simulated 100,000 times and the mean accumulated evidence was plotted as a function of RT (see black, green and magenta dashed lines in Fig. 1C-D). Next, for each participant in Exp. 1 & 2, the accumulated evidence excluding internal noise (both of correct and incorrect responses) was plotted as a function of response-time (see Fig. S1A-B). A mixed-effect linear regression models were fitted to the data of Exp. 1 & 2, with RT and participants serving as random effects and with a free covariance matrix (see thin gray and blue lines in Fig. 1C-D). The model in Wilkinson notation was:

$$\text{Accumulated evidence at response} \sim RT + (RT|subject)$$

Fig. S1. Scatter-plot showing the accumulated evidence in individual trials (blue circles) as a function of decision time of a representative participant in Exp. 1 (A) and Exp. 2 (B). Red circles correspond to the average accumulated evidence at each RT, the grey solid line corresponds to linear regression fitted to the data, and the dashed gray line corresponds to the prediction of the random timer model (see Computational Methods).

B) Regression analysis - Largest temporal cluster of the evidence

Table S1. Beta coefficients for predicting Accuracy, RT and confidence in Exp. 1&2.

	Exp. 1		Exp. 2		
	(Numerical Stimuli)		(Perceptual Stimuli)		
	Accuracy	RT	Accuracy	RT	Confidence
Evidence	$\beta = 1.79^{***}$ $t = 32.04$	$\beta = -0.29^{***}$ $t = -35.57$	$\beta = 0.53^{***}$ $t = 17.58$	$\beta = -0.19^{***}$ $t = -28.12$	$\beta = 0.18^{***}$ $t = 20.95$
Stimulus Consistency (LTC)	$\beta = 0.20^{***}$ $t = 4.68$	$\beta = -0.14^{***}$ $t = -16.60$	$\beta = 0.97^{***}$ $t = 27.43$	$\beta = -0.16^{***}$ $t = -24.25$	$\beta = 0.10^{***}$ $t = 10.99$

*** $p < .001$

The choice-accuracy, RT and confidence in Exp. 1 & 2 were analyzed using mixed model regression analyses (logistic for accuracy and linear for RT and confidence), using the accumulated evidence and normalized stimulus-consistency as fixed factors and participants as random intercepts. The analysis was similar to one reported in Table 1 (see also Computational models), except here we used the more complex measure of stimulus-consistency – larger temporal cluster of the evidence (LTC), in which for each trial we computed the largest temporal cluster that goes with the evidence and divide it by the number of samples (e.g., if the evidence are: 4, -5, 3, 4, 5, 2, -2, -6, then the LTC measure will be equal to: $\frac{4}{8} = 0.5$). The results (Table S1), are fully consistent with the ones obtained using the simpler consistency measure (Difference in Evidence-Directions), and show that LTC improves accuracy and confidence and reduces RT, even while controlling for the accumulated evidence.

C) Model based vs. Model free - Exp. 1

Fig. S2. Decision Classification Boundaries (BCD) of the participants in Exp. 1. Black dashed lines correspond to the model-based DCB, and the red ones correspond to the model-free DCB.

Model based vs. Model free - Exp. 2

Fig. S3. Decision Classification Boundaries (BCD) of the participants in Exp. 2. Black dashed lines correspond to the model-based DCB, and the red ones correspond to the model-free DCB.

D) DCB modulation by stimulus consistency – Exp. 1

Fig. S4. Decision boundaries for all trials (blue) and high consistency trials (red) of all participants in Exp. 1 (numerical stimuli). The boundaries were extracted using the model-free method.

DCB modulation by stimulus consistency – Exp. 2

Fig. S5. Decision boundaries for all trials (blue) and high consistency trials (red) of all participants in Exp. 2 (numerical stimuli). The boundaries were extracted using the model-free method.

DCB modulation by stimulus consistency – Quantitative comparison

In order to quantitatively compare the Decision-Classification Boundary (DCB) of high consistency trials with that of all trials, we computed the difference of the area below the boundary (i.e., the area between the upper boundary and the x-axis) for the two conditions and normalized it by dividing by the sum of the areas:

$$\begin{aligned} & \text{Normalized area difference} \\ &= \frac{\text{Area}_{\text{all trials}} - \text{Area}_{\text{high consistency trials}}}{\text{Area}_{\text{all trials}} + \text{Area}_{\text{high consistency trials}}} \quad \text{Eq. S1.} \end{aligned}$$

If participants require less evidence to reach a decision in high consistency trials, then the decision boundary of these trials should be lower than the decision boundary of all trials, resulting in normalized area difference higher than 0. A single-sample *t*-tests (against 0) confirmed that this was the case in Exp. 1, $t(26) = 2.67$, $p = .01$ (Fig. S6A), as well as in Exp. 2, $t(26) = 3.91$, $p < .001$ (Fig. S6B).

Fig. S6. (A) Normalized area differences of the participants in (A) Exp. 1 and (B) Exp. 2.

We then tested that this DCB separation (consistent trials vs. all trials) takes place in synthetic data for which choices are generated by the stimulus-consistency model, but not when generating data from the full-integration model. Toward this aim, we simulated the

full-integration and the stimulus-consistency models using sequences with the same statistical properties as the ones in Exp. 1 & 2, and with the best fitted parameters of each participant. The DCB of each synthetic participant was extracted using the same procedure as the one used for the empirical data. Fig. S7A-B shows that the DCB separation effect is only predicted by the stimulus-consistency model, $t(26) = 2.45$, $p = .02$ (Exp. 1) and $t(26) = 2.77$, $p = .009$ (Exp. 2), and not by the full-integration model, $t(26) = 0.75$, $p = .45$ (Exp. 1) and $t(26) = -0.61$, $p = .54$ (Exp. 2).

Fig. S7. Mean normalized area difference of the stimulus consistency and the full integration models in (A) Exp. 1 and (B) Exp. 2. Error bars correspond to standard-error of the mean.

E) Temporal clustering effects

In Exp. 3 we generated consistent and inconsistent trials (Fig. 4A), by manipulating the difference between the number of frames with evidence favoring the two alternatives. This manipulation, however, is correlated with the number of frequent winners, which modulates choices in the selective-integration (SI) model (Tsetsos et al., 2016). Thus, although the stimulus-consistency (SC) model provided better fit for the data (Fig. 5), the consistency effect in Exp-3 can also be accounted for by the SI model (Fig. 4C). To distinguish between the two models, we used another measure of consistency – LTC (larger temporal cluster of the evidence, see Table S1). For each participant, we calculated the LTC measure of each trial, and divided the trials into high/low LTC trials by performing a median split. Then, to quantify the sensitivity of the participants to temporal clustering, we computed the difference in mean accuracy of high and low LTC trials ($\text{accuracy}_{\text{high LTC}} - \text{accuracy}_{\text{low LTC}}$). This difference should be higher for a participant whose bias mechanism is modulated by the consistency of each frame with previous ones, as presenting the positive evidence in an as large cluster as possible enhances the overall integrated evidence.

Fig. S8A shows the differences in goodness of fit ($\text{deviance} = -2 \cdot \text{LogLikelihood}$) between the SI and the SC models (positive values indicate a better fit of the SC model). As can be seen, the SC model provides better fit than the SI-model (for all except four participants). Fig. S8B shows the correlation between differences in the goodness of fit of the models and the effect of temporal clustering (each point represents a single participant). As can be seen, the SC model provide a better fit for the data, specifically for participants who showed higher sensitivity to temporal clustering.

Fig. S8. (A) Comparison between the SC and SI models across the 22 participants in Exp. 3. (B) Correlation between differences in goodness of fit of the SI and the SC models and the effect of temporal clustering; each point represents a single participant.

F) Confidence-resolution

In order to account for the differences in confidence resolution in Exp. 3 (Fig 4D), we simulated the stimulus-consistency model using pairs of consistent and inconsistent trials, that were created using the same generating distributions as in Exp. 3 (i.e., with the same evidence content for each pair). Fig. S9A shows 25 pairs of trials (blue – consistent vs. red – inconsistent), which were simulated using the mean fitted parameters: $\theta_{consistency\ parameter} = 1.75$ and $\lambda = 0.21$. The stimulus-consistency model predicts that the bias in favor of consistent evidence would increase the accumulated evidence of consistent compared with inconsistent trials (even though both types of trials have the same evidence). Fig. S9B shows the distributions of accumulated evidence of 100,000 consistent and inconsistent trials simulated using the stimulus-consistency model (using the mean internal noise obtained in the empirical data, $\beta = .05$). The mean distance from the criterion (0) of correct consistent trials is higher than for correct inconsistent trials $t = 3.55, p < .001$. For inconsistent trials, however, the differences in the mean distances from the criterion of the incorrect consistent and incorrect inconsistent trials did not reach statistical, $t = 0.25, p = .79$. Based on the distributions presented in Fig. S9B we predicted the mean confidence response for each condition of the stimulus consistency. To this end, we computed the absolute mean value of the accumulated evidence (i.e., distance from 0) of the correct and incorrect responses separately for consistent and inconsistent response. Then, using linear regression we mapped these values to the mean confidence level of each condition (Fig. S9C, dashed lines).

Fig. S9. Confidence-resolution in Exp. 3. (A) Twenty-five pairs of trials (consistent and inconsistent) were simulated using the same generating distributions as in Exp. 3 (without internal noise; blue – consistent trials, red - inconsistent). The stimulus-consistency model predicts that the bias in favor of consistent evidence would increase the accumulated evidence of consistent trials compared with inconsistent ones. (B) The distributions of accumulated evidence of consistent (blue) and inconsistent (red) trials simulated using the stimulus-consistency model (with internal noise). Solid blue and red vertical bars correspond to the mean of each distribution. (C) Confidence as a function of difficulty and consistency for correct (blue lines) and incorrect (red lines) responses. Data are shown with solid lines and circle symbols and model predictions are shown with dashed lines and diamond symbols. Error bars are within-subjects SE.