Selective integration during sequential sampling in posterior neural signals

Abbreviated title: Selective integration in posterior neural signals

Fabrice Luyckx[1], Bernhard Spitzer[1,2], Annabelle Blangero[1], Konstantinos Tsetsos[2*], Christopher Summerfield[1*]

[1] Department of Experimental Psychology, University of Oxford, Oxford, OX2 6GG, UK.

[2] Center for Adaptive Rationality, Max Planck Institute for Human Development, 14195 Berlin, Germany.

[3] Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, 20246 Hamburg, Germany.

*Equal author contribution.

Corresponding author: Fabrice Luyckx, fabrice.luyckx@psy.ox.ac.uk, Christopher Summerfield, christopher.summerfield@psy.ox.ac.uk

Number of pages: 27

Number of figures: 4

Number of words: 230 (Abstract), 646 (Introduction), 559 (Discussion)

Conflict of Interest: The authors declare no competing financial interests.

Acknowledgements: This research was funded by a British Academy/Leverhulme Small Research Grant (grant number SG141565). FL is funded by the Clarendon Fund, Department of Experimental Psychology and New College Graduate Studentship. CS is funded by a European Research Council Consolidator Award. KT is funded by a European Research council Starting Grant. The authors thank Nick Yeung for providing access to the EEG equipment and Keno Jüchems for useful feedback on earlier versions of the manuscript.

Abstract

Decisions are typically made after integrating information about multiple attributes of alternatives in a choice set. The computational mechanisms by which this integration occurs have been a focus of extensive research in humans and other animals. Where observers are obliged to consider attributes in turn, a framework known as "selective integration" can capture salient biases in human choices. The model proposes that successive attributes compete for processing resources and integration is biased towards the alternative with the locally preferred attribute. Quantitative analysis shows that this model, although it discards choice-relevant information, is optimal when the observers' decisions are corrupted by noise that occurs beyond the sensory stage. Here, we used scalp electroencephalographic (EEG) recordings to test a neural prediction of the model: that locally preferred attributes should be encoded with higher gain in neural signals over posterior cortex. Over two sessions, human observers (of either sex) judged which of two simultaneous streams of bars had the higher (or lower) average height. The selective integration model fit the data better than a rival model without bias. Single-trial analysis showed that neural signals contralateral to the preferred attribute covaried more steeply with the decision information conferred by locally preferred attributes. Moreover, both broadband and time-frequency signals before response onset were best explained by evidence integration under selective integration. These findings provide neural evidence in support of selective integration, complementing existing behavioural work.

Significance Statement

We often make choices about stimuli with multiple attributes, such as when deciding which car to buy on the basis of price, performance and fuel economy. A model of the choice process, known as selective integration, proposes that rather than taking all of the decision-relevant information equally into account when making choices, we discard or overlook a portion of it. Although information is discarded, this strategy can lead to better decisions when memory is limited. Here, we test and confirm predictions of the model about the brain signals that occur when different stimulus attributes of stimulus are being evaluated. Our work provides the first neural support for the selective integration model.

INTRODUCTION

Biological brains evolved to be both precise and efficient. Responding accurately to external stimulation requires noisy sensory signals x to be transduced such that internal estimates $\hat{x}$ are as close as possible to their generative counterparts. Sequential sampling and integration allow the precision of sensory estimates to grow with the number of independent observations obtained (Wald and Wolfowitz, 1949), and neural signals in multiple brain regions implement a basic form of memory that allows optimal estimates to be gradually integrated during inference (Gold and Shadlen, 2007; Hanks and Summerfield, 2017). Forming more precise estimates of the sensory world is often consequential for reinforcement, both in the lab (e.g. when a monkey categorises a stream of noisy sensory signals in return for liquid reward) and in the real world (e.g. when a consumer evaluates the quality of multiple relevant attributes of a product).

However, where information arrives in high volumes, the carrying capacity of the neural system can limit the precision of sensory estimates. Consider a binary discrimination judgment between two stimuli A and B, where information about their worth x_k^A and x_k^B arrives in parallel via a sequence of n discrete sample pairs k. Let us assume that decisions are made by integrating and comparing transduced sensory estimates under the corrupting influence of "late" noise:

$$\sum_{k=1}^{n}[f(x_k^A - x_k^B)] + \mathcal{N}(0,\sigma) > 0$$

With limitless capacity ($\sigma = 0$) the best judgments will be made by simply comparing the summed estimates for A and B – in other words, the optimal $f(\cdot)$ is the identity function. Intuitively, if samples x are numbers and you are in possession of a calculator, the best thing to do is to simply compute their relative sum. However, under imperfect memory, comparative judgments of ground truth quality can paradoxically be promoted by more selective transduction that discards part of the sensory information (Tsetsos et al., 2016; Li et al., 2017; Spitzer et al., 2017; Moran and Tsetsos, 2018). For example, when $\sigma > 0$, a reward-maximising strategy is to selectively integrate the locally highest valued sample (i.e. x_k^A when $x_k^A > x_k^B$) by partially down-weighting the less valued sample by a factor w $(1 > w > 0)$. For illustration, imagine that deciding between two apartments to rent involves evaluating the candidates along multiple continuously-valued dimensions (price, location, size, etc). As memory demands become prohibitive

(e.g. as σ grows), a strategy that is partly based on tallying the number of dimensions on which an alternative is higher valued (e.g. $w < 1$) will lead to more accurate choices. This is because ordinal (rank-based) strategies confer robustness on inference, just as nonparametric methods allow statistical calculations to be more robust to outlying data. Selective integration can be shown to be "optimal" because it maximises reward given irreducible constraints on capacity. At the same time, however, selective integration is "irrational" in that it discards information and can lead to inconsistent choices, including violations of the normative axioms of transitivity and independence from irrelevant alternatives (Von Neumann and Morgenstern, 1944). Several models embodying this principle provide a good fit to human data in tasks involving sequential integration and comparison of discrete samples of information (Tsetsos et al., 2012; Bhatia, 2013; Summerfield and Tsetsos, 2015; Tsetsos et al., 2016; Glickman et al., 2018; Gluth et al., 2018).

The framework of selective integration also makes predictions about neural signals, but these remain as yet untested. For example, several studies have shown that during sequential integration of decision information, the amplitude of single-trial EEG signals contralateral to the side of stimulation correlates with the strength of evidence conferred by each sample (Wyart et al., 2015). Thus, we might expect that those samples that are selectively integrated by virtue of being the local "winner" would be encoded with higher gain, i.e. a steeper linear relationship between x and relevant EEG signals. Here, we tested this view by recording scalp electroencephalographic (EEG) data whilst participants performed a task that involved averaging and comparing bar height within two parallel streams.

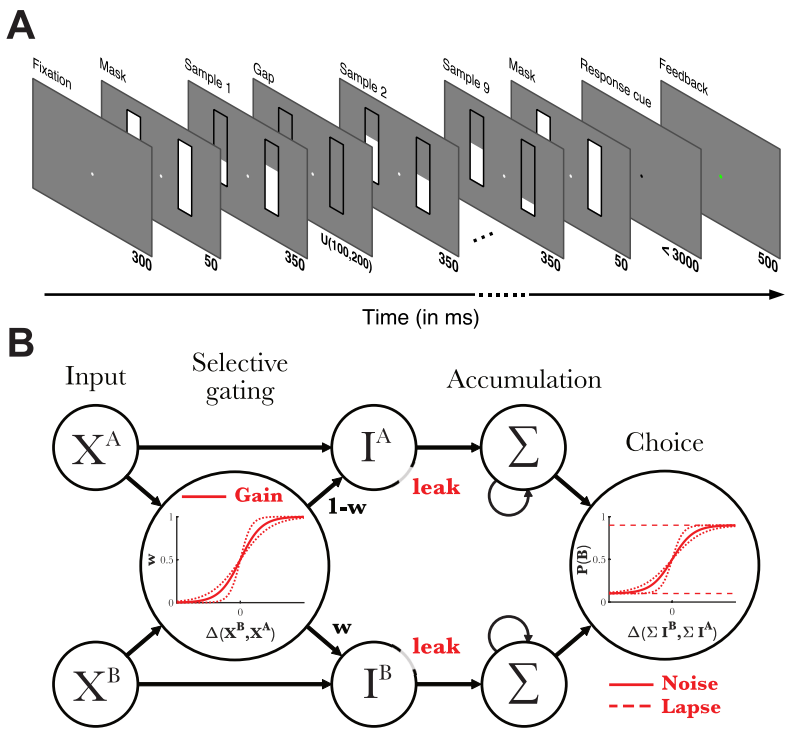

Figure 1. Task design and model schematic. **(A)** Participants viewed a sequence of nine pairs of bars varying in height and compared the average bar height of the left and right stream. In two separate recording sessions, participants indicated the sequence with either the highest or lowest average bar height. **(B)** The selective integration model assumes a separate accumulation process for each option, here the left and right bar sequence. Two concurrent input values X^A and X^B are transformed with a factor w based on their relative difference, down-weighting the lower input value and up-weighting the higher input value. All transduced values I^A and I^B are then integrated over time, with greater loss of information for earlier samples. Finally, the model generates a choice probability based on the difference in integrated values between the two streams.

MATERIALS AND METHODS

Participants

Participants were 18 healthy adult volunteers with no history of neurological or psychiatric illness. Those who failed to complete both recording sessions (n = 2) or who performed at chance level (n = 1) were excluded. Analyses were conducted on the remaining 30 EEG sessions from 15 volunteers (female = 8; left-handed = 2; M_{age} = 24.29 ± 4.48). The study was approved by the Oxford University Medical Sciences Division Ethics Committee (MSD/IDREC/C1/2009/1) and informed consent was given at the start of each recording session. Monetary compensation was based on performance during the experiment (approximately £25 per participant).

Task design and stimuli

All stimuli were created in the Psychophysics 3 Toolbox (Brainard, 1997; Kleiner et al., 2007) for Matlab (MathWorks). The experiment was presented on a 17-inch CRT monitor with 60 Hz presentation rate at a viewing distance of ~70 cm.

On each trial participants viewed a sequence of 9 pairs of bars of variable height (Fig. 1A) that appeared in sequence, left and right of a central fixation dot. The task framing differed across sessions. In one recording session participants were asked to indicate with a button press the stream with the *highest* average bar height ('high frame') and in the other session they were asked to indicate the stream with the *lowest* average height ('low frame'). The order of the two sessions was counterbalanced over participants and both sessions were separated by at least one week. We excluded participants who did not complete both sessions, because the framing manipulation allowed us to cleanly orthogonalize perceptual value (the raw bar height in pixels) from decision value (the level of evidence in favour of either response). The instructions stated that the bars indicated the value of two stock options whose prices fluctuated over time. They were asked to buy the most favourable (high frame) or sell the least favourable (low frame) stock option in different sessions.

Each trial started with a white fixation dot (radius = 5 pixels [px]) presented centrally on a grey background. The dot remained on screen for the whole duration of the trial. All bar stimuli were presented within two black rectangular placeholders (60 by 200 px) on either side of the fixation dot at a horizontal distance of 160 px. Each pair of bars remained on screen for 350 ms, and between successive pairs of bars, there was a gap of empty black frames, lasting approximately 150 ms (uniformly drawn between 100 – 200 ms). The first bar was a forward mask (300ms after fixation onset) of maximal bar height that occurred on both sides and was not included in the analysis. Bars were thus presented at a rate of ~2 Hz with a jitter to minimise steady state neural responses. When all 9 bar pairs had been presented, a backward mask appeared again on screen for 50 ms. The fixation dot then turned black, indicating to the participants that they could respond by pressing the A and L keys on a QWERTY standard keyboard to choose left and right stream respectively. If participants failed to respond within 3 seconds, the fixation dot would turn red for 1000 ms and the words 'Too late' appeared in red above the fixation dot. If they responded before the deadline, the fixation dot would change colour for 500 ms: green for correct response and red for incorrect. The

next trial started after an inter-trial interval (ITI) of 300 ms. A trial could thus last for a maximum of 9350 ms (assuming 3 seconds with no response).

Trials were drawn from one of four bar height distributions (conditions), randomly intermixed throughout the experiment. In two conditions ('Low variance' and 'High variance'), samples for the left and right stream were drawn pseudorandomly from a Gaussian distribution whose mean μ was in turn drawn from a uniform distribution between 80 px and 120 px, and whose standard deviation σ set to 10 (low variance) or 20 (high variance conditions). We used a resampling method to ensure that sample mean and variance closely matched the generative means and variance. To produce a correct answer, 6 px were subtracted from one of the two streams.

Samples for the remaining two conditions ('Frequent winner' and 'Infrequent winner') were generated as follows. In the frequent winner condition, the correct stream was manipulated to contain the winning sample in 2/3 of the sample pairs, while in the infrequent condition the correct stream only contained the winning sample 1/3 of the time. To achieve this, the mean bar length (M) of each trial was first drawn from a uniform distribution ranging between 110 and 130 px. Next, values were calculated separately for every three sample pairs ('triplets' T) in the trial. A deviation (δ) from the trial mean for each triplet pair T was drawn from a uniform distribution $\partial \sim U(15,25)$. Additional noise ε_1 and ε_2 was drawn per triplet T from a normal distribution $\varepsilon \sim N(0,3)$. For every triplet T, the two streams T_A and T_B would have the following form:

$$T_A = \{M + \varepsilon_1, \quad M + \delta + \varepsilon_2, \quad M + 2\delta + (\varepsilon_2 - \varepsilon_1)\}$$
$$T_B = \{M + \delta, \quad M + 2\delta, \quad M\}$$

In this form, T_B always contains two winning samples ($M + \delta$ and $M + 2\delta$), while the means of both triplets are identical. To produce a correct answer, 6 px were subtracted from one of the two streams. In the frequent winner condition the mean of stream A, which contained most of the losing samples, was adjusted to make it the incorrect answer, while in the infrequent winner condition the mean of stream B, containing the most winning samples, was adjusted. Adjustments in low and high frames were inversions of one another.

To prevent participants from shifting gaze directly at one of the two streams, an area around the fixation dot (256 x 205 px) was defined in which participants had to keep fixation. Eye movements were monitored using a Tobii EyeX eye-tracker (Tobii Technology, Stockholm, Sweden) and when the eyes moved outside the fixation area, the trial was classified as incorrect and the message 'Eyes moved!' appeared above the fixation dot. These trials were also omitted from analysis (0.35%).

Each recording session consisted of 600 trials, divided in 10 blocks, resulting in a maximum of 1200 trials (10800 samples) for each participant. Participants could take self-timed breaks in between blocks. One session lasted around 2.5 h, including preparation of the cap, 60 minutes of task performance and removal of the cap at the end of the experiment.

Selective integration model

The model decision is based on the output of two accumulators Y^A and Y^B that integrate the input values of the left and right stream respectively (Fig. 1B). Input values X^A and X^B are the raw pixel heights H of each sample (bar), inverted for the low frame $(200 - H)$. Each integrator is updated separately for each sample k according to the following formula:

$$\begin{aligned} Y_k^A &= Y_{k-1}^A + I_k^A \cdot \lambda \\ Y_k^B &= Y_{k-1}^B + I_k^B \cdot \lambda \end{aligned} \tag{1}$$

Y^A and Y^B are both initialized to zero. I^A and I^B are the transformed values of the input after graded selective integration:

$$\begin{matrix} I_k^A = (1-w) \cdot X_k^A \\ I_k^B = w \cdot X_k^B \end{matrix} \text{ with } \begin{cases} w < 0.5 \text{ where } X^A > X^B \text{ for sample } k \\ w > 0.5 \text{ where } X^A < X^B \text{ for sample } k \end{cases} \tag{2}$$

Where w is a gating variable that is determined by passing the difference between input values $(\Delta X = X^B - X^A)$ through a logistic function

$$w = \frac{1}{1 + e^{-m(\Delta X)}} \tag{3}$$

with the slope of m determining the extent to which the difference between the two input values attenuates the input values X. The transfer function ensures that for very large differences between the bars, where it is clear which of the two bars wins, the winning input value is integrated close to its original value, while the losing sample value is almost entirely suppressed. On the other hand, for input values that lie close together and comparison is difficult, both values carry approximately equal weight to their accumulators. A recent study demonstrated that a 'graded' selective integration such as the one proposed here outperformed a 'binary' selective integration that works independent of the size of the sample difference, both in predicting human choices and its robustness to noise (Glickman et al., 2018).

The model assumes a leaky accumulation, with earlier samples carrying less weight on the final decision. Each transduced sample is therefore multiplied with the leak function value:

$$\lambda = \frac{(1-l)^{n-k}}{(1-l)} \tag{4}$$

with l the amount of information loss and n the total number of samples.

Finally, the model outputs a choice probability based on the difference between the accumulators after all evidence has been accumulated:

$$P(B) = \kappa + \frac{1-2\kappa}{1-e^{-s(Y^B-Y^A)}} \tag{5}$$

s is the slope of the response function, often referred to as 'late' or 'integration noise', a factor that adds uncertainty to the final choice. Finally, a lapse rate parameter κ was added to the response function to allow for higher error rate in the model.

Model fitting procedure

The model has four free parameters: selective gating parameter m, leak l, late noise s and lapse rate κ. Best fitting parameter sets for each participant were obtained by

minimizing the negative log-likelihood function using a scatter-search based global optimization solver in Matlab. The parameter search space was constrained as follows:

$$m, \kappa \in \{0.001, \dots, 0.2\},$$
$$l, s \in \{0.01, \dots, 0.5\}$$

To test whether the gating parameter m significantly contributed to the model performance, two models were compared: the full SI model and a fixed gating model. For the fixed gating model m was set to 0, which effectively gave equal weight to each input ($w = 0.5$), whereas for the full model m was unrestricted, allowing w to take any value between 0.5 and 1. Model fits were compared through split-half cross-validation, a method that allows comparing models with different numbers of free parameters through their generalizability to new data. Both models were trained on one recording session and the best fitting parameters from the training set were then used to estimate how well they predicted responses in the other session. Negative log-likelihoods were summed per participant and fed into a Bayesian Model Selection procedure (Stephan et al., 2009) to compute posterior evidence for one model over another. Best fitting parameters for the model used in subsequent analyses were estimated from the collapsed data over both recording sessions.

Parameter recovery

To ensure there was no trade-off between the parameters of our model within our dataset, we tested whether simulated parameter combinations could be recovered using the same model-fitting procedure as above. We therefore generated 100 random combinations of our four parameters, with the search space of each parameter limited by the minimum and maximum estimated values from the real data. Model responses were then generated for each parameter combination using the input data from a randomly selected subject each time (1200 trials). Finally, the input data and model generated responses were used to estimate best-fitting parameters again and test how well we could recover the original parameter combinations.

Behavioural analysis: regression models

All analyses of human data were performed on the collapsed data from both recording sessions, with input values for the 'low frame' session inverted $(200 - H)$ to reflect the input value rather than raw bar height. The impact of winning and losing samples on choice was tested through a generalized linear model (GLM), modelled with a binomial distribution and a logit link function, predicting the probability of choosing the right stream, as follows:

$$P(R) = \beta_0 + \beta_1 \cdot \sum X^L(L > R) + \beta_2 \cdot \sum X^L(L < R) + \beta_3 \cdot \sum X^R(L < R) + \beta_4 \cdot \sum X^R(L > R) + \varepsilon \quad (6)$$

The four parametric regressors in the model coded for the cumulative sum of input values for all (1) winning samples on the left side $[X^L(L > R)]$, (2) losing samples on the left side $[X^L(L < R)]$, (3) winning samples on the right side $[X^R(L < R)]$ and (4) losing samples on the right side $[X^R(L > R)]$. The regressors were always standardized before parameter estimation. Parameter estimates for the left stream were then sign-flipped and averaged with the estimates of the right stream, in order to obtain an average modulation of winning and losing samples (Fig. 2A).

The SI model further allows for a 'recency effect' through the leak parameter, whereby earlier samples carry less weight in the final choice than later samples. We tested this assumption through a new GLM predicting choices for the right stream based on the signed difference of each sample pair (right minus left).

$$P(R) = \beta_0 + \beta_1 \cdot \Delta(X_1^R, X_1^L) + \cdots + \beta_9 \cdot \Delta(X_9^R, X_9^L) + \varepsilon \quad (7)$$

The model thus consisted of 9 parametric regressors representing the serial position in the stream. If choices were driven more by more recent samples, coefficient estimates should be higher for regressors coding for later samples (Fig. 2C).

All statistical tests were performed at the group level. Given the relatively low number of subjects, we opted for non-parametric tests that do not assume a normal distribution of the data.

EEG acquisition and pre-processing

EEG signals were recorded using a Neuroscan system with SynAmps-2 digital amplifiers with 60 Ag/AgCl electrodes located at FP1, FPz, FP2, F7, F5, F3, F1, Fz, F2, F4, F6, F8, FT7, FC5, FC3, FC1, FCz, FC2, FC4, FC6, FT8, T7, C5, C3, Cz, C2, C4, C6, T8, TP7, CP5, CP3, CP1, CPz, CP2, CP4, CP6, TP8, P7, P5, P3, P1, Pz, P2, P4, P6, P8, PO7, PO5, PO3, POz, PO4, PO6, PO8, O1, Oz, and O2. Four additional EOG electrodes in bipolar montage (two horizontal, two vertical) were recorded, together with one mastoid for reference. Electrode impedance was brought below 10 kΩ before recording. Data was collected at a sampling rate of 1 kHz and high-pass filtered at .01 Hz.

Preprocessing was done in Matlab using functions from the EEGLAB toolbox (Delorme and Makeig, 2004) and custom scripts. First, data was downsampled to 250 Hz, subsequently low-pass filtered at 40 Hz and then high-pass filtered at 0.5 Hz. Excessively noisy channels were identified through visual inspection for each participant and interpolated based on the weighted average of the surrounding electrodes. Next the data was re-referenced offline to average reference (excluding EOG channels). Trial epochs were extracted spanning 1 second prior to the fixation dot onset until 7 seconds after. Epochs were subsequently baselined relative to the pre-fixation time window of -500 to 0 ms. Artefacts related to eye-blinks and other sources of consistent noise were identified through Independent Component Analysis (ICA) and removed from the data after visual inspection. Finally, the data were epoched again at different times for various analyses. For the sample-based regressions, the data were epoched relative to each sample pair onset, starting at 100 ms before until 750 ms after sample onset and baselined again relative to the full pre-stimulus window, to exclude any systematic offset during the length of the trial. For response-locked analyses, the epoch was set to 3 seconds prior to response onset to 300 ms after. The baseline window was chosen at 3 to 2.5 seconds before response onset.

Time-frequency transformation

The pre-processed epochs spanning 8 seconds were transformed into the time-frequency domain using functions from the Fieldtrip toolbox (Oostenveld et al., 2011) for Matlab and custom scripts. Power was calculated every 25 ms within an epoch by convolving a sliding Morlet wavelet (width = 7 cycles), for frequencies between 8 and 38 Hz, incremented in steps of 3 Hz. To allow for a comparison of power estimates

over frequency bands, each frequency within an epoch was dB transformed relative to the baseline window [-300 300] ms, a 'silent' period including the ITI and start of a new trial, using the following formula:

$$dB = 10 \cdot log_{10}\left(\frac{signal}{baseline}\right). \tag{8}$$

For the sample regression on TF data, the data was subsequently epoched into smaller epochs between [-100, 750] ms relative to sample pair onset. For visualization purposes the time-frequency results were interpolated, while cluster statistics were performed on the raw data.

EEG analysis: regression models

To understand how EEG signals were modulated by the stimulus information, we used a regression-based approach. Single-trial EEG data was regressed against parametric predictors for each time-point and electrode independently. This analysis tests for a linear relationship between the magnitude of decision values and the momentary amplitude of the EEG signal at each timepoint following each sample. Coefficient estimates obtained from the regression model reflect the slope of the modulation, i.e. the strength of the linear relationship. Group-level statistics were subsequently computed over the obtained time courses of parameter estimates over all participants.

First, we constructed a sample-based linear regression model that allowed us to dissociate the neural modulation of winning and losing samples:

$$EEG = \beta_0(L > R) + \beta_1(R > L) + \beta_2 \cdot |X_k^L|(L > R) + \beta_3 \cdot |X_k^L|(L < R) + \beta_4 \cdot |X_k^R|(R > L) + \beta_5 \cdot |X_k^R|(R < L) + \varepsilon \tag{9}$$

with β_0 and β_1 as binary indicator variables for samples where the left (L) or right (R) bar in the pair had the highest decision value respectively. The four parametric regressors β_3 - β_6 code for the absolute decision evidence (X_k) of the bar when the left bar won (β_3), left lost (β_4), right won (β_5) and right lost (β_6). Effectively, each sample pair related to four regressors (β_0, β_1 and two of the four parametric regressors). The trial mean was subtracted from all bars to eliminate any differences related to the

absolute size of the bars. Single-trial EEG signal locked to stimulus onset was subsequently regressed against the full model with 6 regressors, independently for each sample pair. Estimated time courses were averaged over all sample pairs before statistical inference.

We defined two regions of interest (ROI) based on previous research (Wyart et al., 2015) to test for early modulation of sample evidence: left (O1, PO5, PO7, P5, P7, CP5, TP7) and right (O2, PO6, PO8, P6, P8, CP6, TP8) occipito-parietal (OP) electrodes. The coefficients β_0 and β_1 in our regression model encode the average EEG signal for sample pairs where the left (β_0) or right sample won (β_1). To compare categorical modulation of stimulus-evoked activity, we averaged the activity over those electrodes contralateral to the winning sample, namely right OP electrodes for β_0 (trials where the left sample won) and left OP electrodes for β_1(trials where the right sample won). This is equivalent to flipping the scalp maps and treating each sample pair as if the right sample won. The opposite hemispheres then contained the neural signals modulated by the losing sample. Time series for a parametric modulation of winning samples (X_{win}) were obtained in a similar way: averaging over the right OP electrodes when left won (β_3) and left OP electrodes when right won (β_5). Time series for parametric modulation of losing samples (X_{lose}) were obtained at right OP electrodes when left lost (β_4) and left OP electrodes when right lost (β_6). Significance of adjacent time points in the averaged time series were tested using nonparametric cluster-based permutation tests (cluster-defining threshold and corrected significance level at $p < 0.05$, 5000 iterations) (Maris and Oostenveld, 2007). Estimated time series were smoothed per subject using a low-pass filter (Butterworth filter, cut-off = 25 Hz), removing small fluctuations in the signal that could obscure larger clusters from the permutation test.

The same regression model and contrasts were used to study differences in lateralization of alpha band activity, with an additional iteration over frequency bands. The averaged time window was temporally smoothed with a 100 ms full-width at half-maximum (FWHM) Gaussian kernel before cluster-based permutation test (cluster-defining threshold and corrected significance level at $p < 0.025$).

Next, we sought to identify neural signals related to evidence accumulation. We formalized 'accumulated evidence' as the signed difference between the cumulative sum of evidence of the right minus left stream.

$$EEG = \beta_0 + \beta_1 \cdot \sum_{k=1}^{9} \Delta(X_k^R, X_k^L)_k + \varepsilon \quad (10)$$

Under the assumption that information is not integrated as such, but undergoes a weighted transformation before integration, we predicted that the accumulated evidence given by the SI model (Eq. 1) should be better at capturing the signal variance than the actual presented input. Since both accumulated evidences are highly correlated, we partialled one out from the other before estimating parameter coefficients to identify if either the two inputs captured unique variance in the EEG signal. Single-trial EEG data locked to the response was then regressed onto the residual accumulated evidence at each time point and electrode. Contrasts of activity at central electrodes were then calculated (C3 – C4) to test the differences in temporal fluctuations of parameter estimates from both regressions. An identical procedure was used for the time-frequency data, with an additional iteration over frequency bands. The time window was temporally smoothed with a 100 ms full-width at half-maximum (FWHM) Gaussian kernel before cluster-based permutation test (cluster-defining threshold and corrected significance level at $p < 0.025$).

RESULTS

Behavioural results

We first examined the behavioural data to test whether participants weighted the samples as described by the selective integration model. We fitted a logistic regression model that separately estimated the influence of winning and losing samples on right-hand choices (Eq. 6). The null hypothesis is that there will be no difference in how much influence winning samples (e.g. X^A when $X^A > X^B$) and losing samples (e.g. X^B when $X^A > X^B$) carry on choice. However, we found that parameter estimates for winning samples (Mdn β_{win}= 3.12) were significantly higher than for losing samples (Mdn β_{lose} = 3.06; W = -120, $p < 0.001$, two-sided Wilcoxon paired signed rank test), suggesting there is indeed a difference in how samples were weighted according to their relative decision value (Fig. 2A).

If participants give more weight to winning samples, their performance might be poorer in a condition where the incorrect stream contains the most winning samples, a so-called 'frequent-winner' effect (Tsetsos et al., 2016). To maximize the opportunity to compare trials where the winning stream contained the most versus the fewest winning samples, we included two conditions that explicitly contained these types of trials ('frequent winner' and 'infrequent winner') as well as two conditions where the height variance was either low or high. A non-parametric omnibus test (Friedman test) first confirmed there were differences in choice accuracy between the four conditions (X^2 = 13.34, p = 0.004; d.f. = 3, N=15) (Fig. 2B). We then ran a post hoc pairwise comparison between all conditions (two-sided Wilcoxon signed rank test) and corrected for multiple comparison using FDR correction (Benjamini and Hochberg, 2009). Unlike previous reports (Tsetsos et al., 2016), we did not observe a statistically significant difference between the 'infrequent winner' (Mdn = 0.66) and 'frequent winner' conditions (Mdn = 0.76; Z = 1.62, p = 0.1054), although the results did show a numerical trend in the predicted direction and a selective integration model fitted to the human choice data predicted a difference between the (in)frequent winner conditions (see below). Participants performed significantly worse in the 'infrequent winner' condition, but only in comparison to the 'low variance' condition (Mdn = 0.77; Z = 2.24, p = 0.0248). Participants also performed significantly worse in the more difficult high variance condition (Mdn = 0.72) compared to the low variance condition (W = -120, p < 0.001).

Finally, the SI model predicts a 'recency effect', where samples presented later in a trial should carry more weight on the final decision, because there was less time for this information to be lost. A new logistic regression model predicting right-hand choices (Eq. 7) showed that indeed, when sample pairs were assessed based on their serial position in a trial, parameter estimates increased over time, indicated by a significantly increasing slope fitted to each subject's parameter estimates (Mdn slope = 0.13; W = -120, p < 0.001).

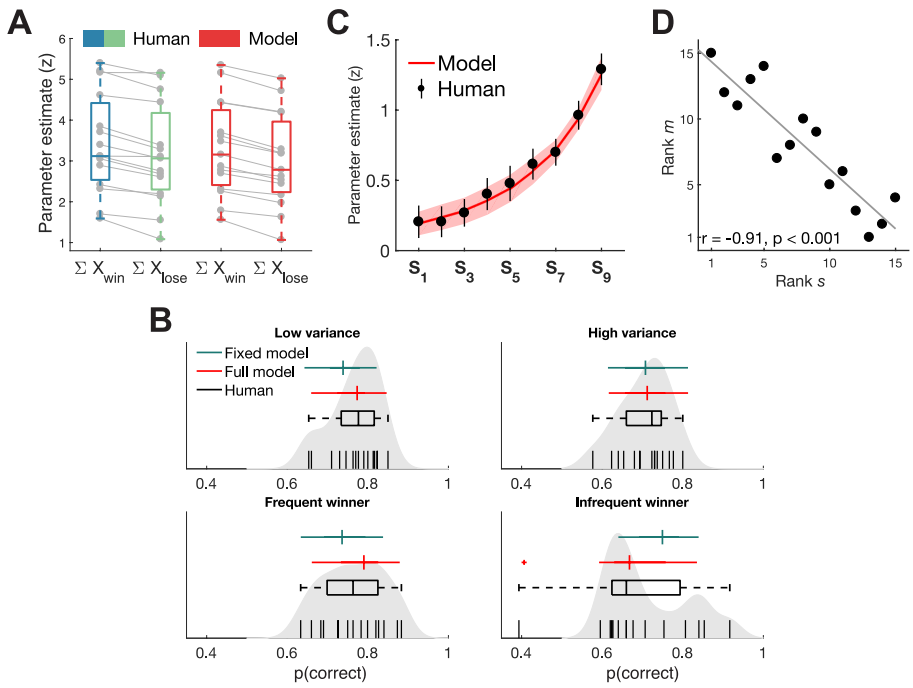

Figure 2. Behavioural results and model fits. **(A)** Parameter estimates from a regression model predicting the probability of choosing the right-hand sequence, both in humans (blue/green) and the fitted SI model (red). In line with predictions of the SI model, sample values were weighted higher on average when decision evidence was greater relative to the concurrent sample. A SI model fitted to the choice data was able to recreate this pattern. **(B)** SI predicts that a tendency of tallying winning samples will lead to lower accuracy in trials where the incorrect sequence contains most winners. Performance was indeed lower in the 'infrequent winner' condition, but not significantly so from the 'frequent winner' condition. Qualitatively, the full SI model (red) resembled human accuracy better in all four conditions compared to a model where the gating parameter was fixed to 0. Black vertical lines show individual participant's accuracy and grey shaded area represents non-parametric kernel estimate (bandwidth = 0.03) of the accuracy distribution in our sample. **(C)** Parameter estimates from a new regression model showed a 'recency effect', where earlier samples carried less weight on participants' final choice. The fitted SI model closely followed the same trend (red). Error bars and shaded area represent 95% confidence interval. **(D)** The slope of the response function (s) in the fitted SI model correlated significantly with the slope of the selective gating function (m), suggesting participants with greater levels of late noise compensate through stronger gating of information. Parameter recovery showed that this pattern was not due to a trade-off between the two parameters. Grey line represents least-squares fit line.

Model fits to human data

A more formal test for selective integration can be achieved by fitting the SI model to human data and comparing it to an equivalent alternative model that does not show selective integration. A full SI model with 4 free parameters (gating m, leak l, late noise s and lapse rate κ) was fitted to each participant's trialwise choice data (m = 0.0377 ±

512 0.03; l = 0.234 ± 0.1; s = 0.0414 ± 0.03; κ = 0.0176 ± 0.02). To assess whether the
513 gating parameter meaningfully contributed to the model fits, we compared the full
514 model to a fixed gating model where the gating parameter was fixed to 0, eliminating
515 the unequal weighting of sample inputs. Both models were tested through two-fold
516 cross-validation, estimating parameter fits on one recording session and testing model
517 predictions on the other session. Negative log-likelihood (LL) estimates of the test sets
518 were summed per participant and aggregate LL were fed into a Bayesian Model
519 Selection procedure (Stephan et al., 2009) to determine the best fitting model. The full
520 model (mean LL = -602.05) had a protected exceedance probability of 0.9983 compared
521 to the fixed model (mean LL = -615.54), suggesting the gating parameter of the SI
522 model was necessary to explain choice behaviour. To determine qualitative fits of both
523 models, we reran our behavioural analyses on model choices. The full model was able
524 to reproduce the difference in parameter estimates for winning and losing samples (Fig.
525 2A) and the recency effect (Fig. 2C), even though the model was not specifically fit to
526 these data points. Only the full model was further able to capture the lower performance
527 in the 'infrequent winner' condition (Fig. 2B), and in general captured the patterns in
528 human performance per condition better than the fixed model.
529 Finally, Tsetsos et al. (2016) showed that the fitted parameters for gating and late
530 noise were highly correlated in their data set (but not for simulated data), suggesting
531 that the strength of the gating process compensates for higher levels of integration
532 noise. We replicated this finding in our data set (Spearman ρ = -0.91, $p < 0.0001$) (Fig.
533 2D). To ensure that this correlation was not due to a trade-off between the two
534 parameters in our data set, we performed parameter recovery on simulated parameter
535 combinations and assessed how well our fitting procedure could recover the original
536 parameters. For all four parameters, the mean absolute squared difference between
537 original and recovered parameters was smaller than 10^{-10} (100 simulations).
538 Furthermore, the recovered parameters for gating m and late noise s did not correlate
539 after the fitting procedure (Spearman ρ = -0.07, p = 0.4602).
540
541 **Early modulation of posterior EEG activity**
542 The SI model states that selective gating occurs at the level of the individual samples
543 before evidence is passed to a subsequent integration stage. We thus predicted that
544 posterior EEG signals would encode locally "winning" samples with higher gain than

545 “losing” samples, i.e. that the slope of regression linking decision information to EEG
546 amplitudes would be steeper for winners than losers. In addition to this multiplicative
547 effect, we also tested for an additive bias, i.e. that EEG signals were higher contralateral
548 to the winning sample, irrespective of its input value.

549 Since samples were presented parafoveally, we expected these effects to occur
550 contralateral to the location of each sample. We therefore focused our analyses on two
551 a priori defined posterior regions of interest (ROI): left and right occipito-parietal
552 electrodes. We constructed a linear regression model (Eq. 9) with two intercepts, coding
553 whether the left or right sample won within a sample pair (to test the additive effect),
554 and four parametric regressors coding the sample evidence separately for left and right
555 and whether the sample won or not (to test the multiplicative effect).

556 We first examined the coefficients associated with the parametric regressors that, at
557 each time point and electrode, reflect the slope of the relationship between the
558 magnitude of the decision evidence and EEG amplitude, separated for winning and
559 losing samples. After collapsing over posterior electrodes contralateral to winning or
560 losing samples, we found that the sample evidence of both winning and losing samples
561 was encoded in the EEG signal around 250 ms, but more importantly, this was
562 significantly stronger for winning samples ($p_{cluster} < 0.05$; Fig. 3A). This suggests that,
563 as predicted, samples carrying equal decision evidence are encoded more strongly in
564 the EEG signal when they are the winning sample in a sample pair.

565 The intercepts of the regression model then reflect the average signal in trials where
566 either the left or the right sample won, independent of their decision evidence. At each
567 time point and electrode, the parameter estimates reflect an additive effect of sample
568 identity (winner or loser) on EEG signals. To test a difference in the additive effect of
569 winning and losing samples, we averaged parameter estimates at ROI electrodes
570 contralateral to all winning samples and compared them to the averaged parameter
571 estimates at ROI electrodes contralateral to all losing samples. We found that between
572 ~200 and 300 ms after sample onset the EEG signal in posterior electrodes encoded
573 winning samples more strongly than losing samples ($p_{cluster} < 0.05$; Fig. 3B), indicating
574 samples were encoded differently based on the sample’s relative identity (winner/loser
575 status), independent of the decision evidence they carried.

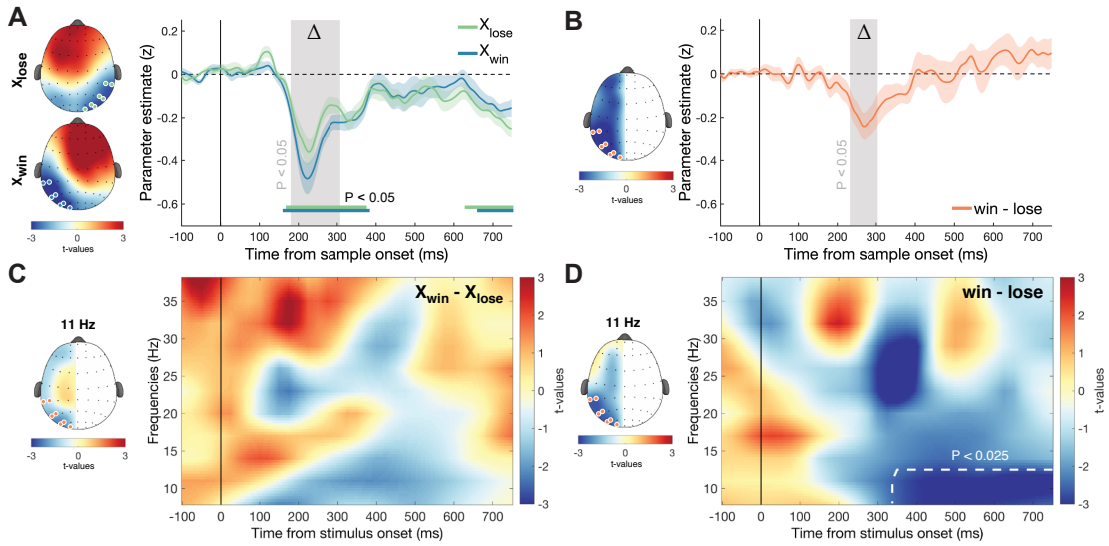

Figure 3. Early modulation by relative sample evidence in posterior electrodes. A linear regression model was used to separately assess the influence of winning and losing samples on EEG signals at each time point, electrode and sample pair. Coefficients for left sample regressors in our model were horizontally flipped, so winning samples were projected to the left hemisphere and losing samples to the right hemisphere. Plots represent coefficient estimates averaged for all nine sample pairs. **(A)** Coefficient estimates for the parametric regressors in our model showed that winning and losing samples were both significantly encoded in the EEG signal around 250 ms (bottom colored lines, $p_{cluster} < 0.05$) in posterior electrodes (coloured dots on the scalp plots) contralateral to the sample location. More importantly, winning samples were encoded significantly stronger compared to losing samples (grey shaded area, $p_{cluster} < 0.05$). Shaded coloured area represents SEM. Top and bottom scalp plot show significance of coefficient estimates for losing and winning samples respectively at the identified time window of significant dispersion. **(B)** By contrasting the coefficients for the intercepts of the regression model, depicting the categorical encoding of winning and losing samples, we observed an additive effect around 250 ms after stimulus onset in posterior electrodes that was stronger for winning samples compared to losing samples (grey shaded area, $p_{cluster} < 0.05$). The orange line represents the difference between the time series of the two intercepts at the posterior electrodes indicated on the scalp plot. Scalp plot shows the left minus right hemispheric difference of coefficient estimates. Shaded coloured area represents SEM. **(C)** An identical regression model was used to explain time-frequency data. We did not find a significant difference in coefficient estimates between winning and losing samples based on their parametric values. Colour map shows difference between winning and losing samples at the electrodes indicated on the scalp plot. Scalp plot shows difference between winning and losing samples and hemispheres at same time window as (D) for the 11 Hz frequency band. Time-frequency data was interpolated for visualization purposes, while cluster statistics were performed on the raw data. **(D)** However, we did observe an additive effect in alpha frequencies (8 – 12 Hz), with greater suppression for winning samples compared to losing samples in the same contralateral posterior electrodes starting around 375 ms after sample onset ($p_{cluster} < 0.025$). Colour map shows the difference between winning and losing samples at the electrodes indicated on the scalp plot. Scalp plot shows difference between

winning and losing samples and hemispheres at the time of the identified cluster for the 11 Hz frequency band.

Alpha suppression

It has previously been suggested that SI could occur when selective attention is oriented to the winning sample (Glickman et al., 2018). Previous work has reported that a lateralized suppression in alpha (8 – 12Hz) power occurs when participants orient visuospatial attention towards a contralateral target (Sauseng et al., 2005; Bacigalupo and Luck, 2019). We used the same sample-based regression model as above (Eq. 9), but now predicting sample-wise fluctuations in power over frequency bands between 8 and 38 Hz. The same occipito-parietal contrasts for the intercepts of the model revealed a significant cluster of greater alpha suppression contralateral to the winning sample compared to the losing sample from ~375 ms onwards ($p_{cluster} < 0.025$) (Fig. 3D). This might imply that covert attention was dynamically shifting to the winning sample. We could not find any statistical evidence for a parametric modulation of alpha suppression by input value ($p_{cluster} > 0.05$) (Fig. 3C).

Evidence accumulation

In the SI model, samples are first weighted according to their relative value, and then integrated into a cumulative signal that eventually determines choice. Previous studies have observed in both perceptual and value-based decision making tasks that central EEG signals build up in concert with the strength of accumulated evidence (Donner et al., 2009; van Vugt et al., 2012; Gluth et al., 2013; Wyart et al., 2015; Tickle et al., 2016; Pisauro et al., 2017; von Lautz et al., 2019). We used a time-window several hundreds of milliseconds before the response onset to test for encoding of the accumulated evidence, as predicted by the SI model. The response-locked EEG signal was regressed onto the total signed cumulative difference (Eq. 10), first calculated according to the true input and second according to the transduced input from the individually fitted SI model. To test which of the two inputs explained unique variance, inputs were partialled out from one another before regressing EEG signal onto the residual input. Only the input values predicted by the SI model significantly explained unique variance in the lateralized broadband EEG signal of central electrodes (C3 and C4) from ~120 ms before response onset ($p_{cluster} < 0.05$; Fig. 4A). We then repeated the same analysis in the time-frequency domain, with a focus on the beta band (15 – 30

Hz). Again, the SI model was better in explaining the neural data than the original input, with a significant cluster in beta band frequencies starting around 300 ms before response onset ($p_{cluster} < 0.025$; Fig. 4B).

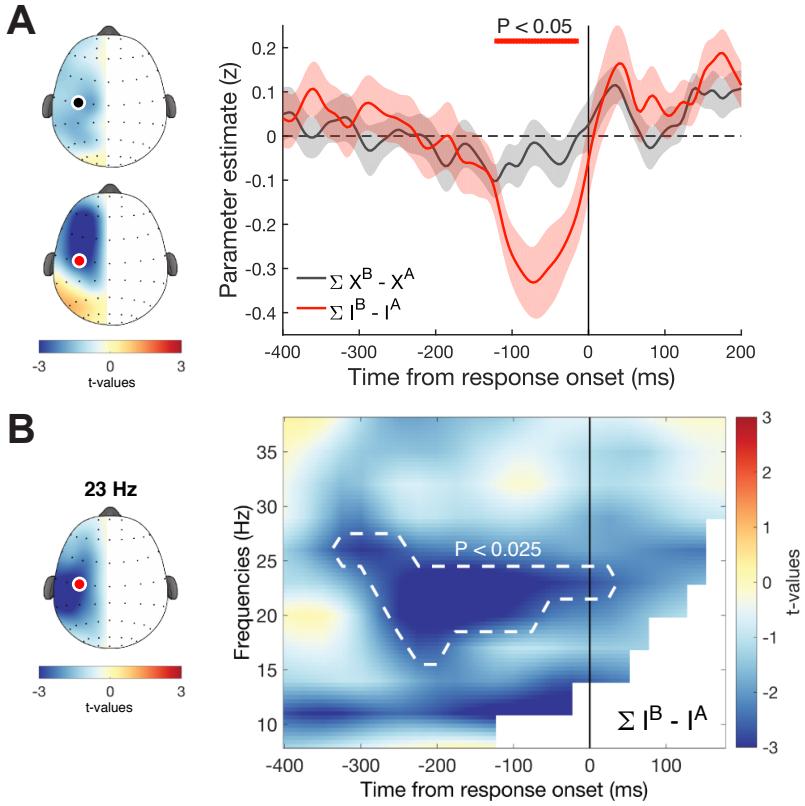

Figure 4. Evidence accumulation is explained by SI model in central electrodes **(A)** Response-locked EEG signal was regressed onto the accumulated values of the untransformed input values (black) or as estimated by the SI model (red). To find which of the two highly correlated regressors explained unique variance in the signal, both regressors were first partialled out from one another and the residual values were used as a new regressor in the regression model explaining EEG signals. Only the estimated accumulation values from the SI model explained unique variance, with a significant negative modulation contralateral to the response hand from around 120 ms before response onset. Scalp plots show the difference in parameter estimates of left minus right hemisphere between -200 and 0 ms. Time series show the same hemispheric difference at motor electrodes C3 and C4 (indicated as coloured dots on the scalp plots). Shaded coloured area represents SEM. Top coloured line depicts time of cluster significance ($p_{cluster} < 0.05$)**. (B)** The same analysis pipeline was applied to each frequency band of the time-frequency transformed data. We identified a significant cluster of contralateral decrease in beta power (15-30 Hz) starting around 300 ms before response onset ($p_{cluster} < 0.025$). Scalp plots show the difference of parameter estimates of left minus right hemisphere between -200 and 0 ms at 23 Hz. Time-frequency data was interpolated for visualization purposes, while cluster statistics were performed on the raw data.

DISCUSSION

In this report we tested for, and obtained, neural evidence in support of a "selective integration" policy during human decision-making. Input values conferred by each of two samples (bars) was weighted differently according to its local rank, and this behavioural effect was accompanied by modulation of neural signals over posterior electrodes. This modulation arose early after sample onset (~ 250 ms) and thus possibly before evidence integration could occur. Furthermore, input values that had been adjusted as the SI model predicts (I^A and I^B) explained more variance in neural signals than the untransformed input values (X^A and X^B), both in broadband signals and beta-band power.

Selective attention has been hypothesized to be an important driving mechanism, for selective integration models (Glickman et al., 2018) and related theoretical accounts (Bhatia, 2013; Gluth et al., 2018). Our results seem to suggest an important contribution of selective attention in two ways. First, contralateral alpha power suppression has long been related to visuospatial attention towards a target (Sauseng et al., 2005; Bacigalupo and Luck, 2019). We found that 375 ms after each stimulus onset there was a suppression of alpha power contralateral to the winning sample. Although later in time than the modulation we found in broadband signals, previous studies have found similar 'late' suppression of alpha. Second, the timing, location and pattern of the parametric modulation resemble the findings of a previous study that measured fluctuations in human perceptual choice under focused and divided attention (Wyart et al., 2015). In that experiment, participants integrated tilt information from two simultaneous streams of Gabor patches under focussed attention (where the decision-relevant stream was signalled in advance) and divided attention (where it was not). In the focused condition, posterior electrodes contralateral to the attended stream, but not the unattended stream, showed a negative modulation around 250 ms. The reduced encoding of the losing sample in our study could similarly point to a focussing of attention towards the winning sample. Here we additionally show that this putative neural correlate of attention can be shifted rapidly between the two sides of the screen as samples arrive.

Our neural findings support both selective integration as described by Tsetsos et al. (2012) and models making similar theoretical claims. For example, the Associations and Accumulation Model (AAM) (Bhatia, 2013) similarly predicts that attention should

be drawn to options where the attribute value is high, which in turn asymmetrically drives choices. The SI model also shares a lot of features with the Multi-alternative Decision by Sampling (MDbS) model (Noguchi and Stewart, 2018). One characterizing feature of the MDbS model is that attribute values are integrated through ordinal pairwise comparison, simply tallying attribute ranks. This is in contrast to the SI model, where the magnitude of this difference influences the strength of the evidence weighting. Both our model fits and the neural results suggest that a comparison between attribute values is not solely a process of ordinal comparison. We did indeed find an additive effect by rank (winner vs loser), encoded both in broadband and time-frequency signals. However, there was a further parametric modulation between winning and losing samples that depended on the magnitude difference of the samples, which supports an additional multiplicative step not posited by MDbS.

In conclusion, our study provides a first step in understanding the neural mechanisms which accompany parallel integration and comparison of discrete samples of information.

REFERENCES

Bacigalupo F, Luck SJ (2019) Lateralized suppression of alpha-band EEG activity as a mechanism of target processing. J Neurosci 39:900–917.

Benjamini Y, Hochberg Y (2009) Controlling the False Discovery Rate : A practical and powerful approach to multiple testing. J R Stat Soc Ser B 57:289–300.

Bhatia S (2013) Associations and the accumulation of preference. Psychol Rev 120:522–543.

Brainard DH (1997) The psychophysics toolbox. Spat Vis:433–436.

Delorme A, Makeig S (2004) EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J Neurosci Methods 134:9–21.

Donner TH, Siegel M, Fries P, Engel AK (2009) Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. Curr Biol 19:1581–1585.

Glickman M, Tsetsos K, Usher M (2018) Attentional Selection Mediates Framing and Risk-Bias Effects. Psychol Sci 29:2010–2019.

Gluth S, Rieskamp J, Büchel C (2013) Classic EEG motor potentials track the emergence of value-based decisions. NeuroImage 79:394–403.

Gluth S, Spektor MS, Rieskamp J (2018) Value-based attentional capture affects multi-alternative decision making. eLife:1–36.

Gold JI, Shadlen MN (2007) The Neural Basis of Decision Making. Annu Rev Neurosci 30:535–574.

Hanks TD, Summerfield C (2017) Perceptual Decision Making in Rodents, Monkeys, and Humans. Neuron 93:15–31.

Kleiner M, Brainard D, Pelli D (2007) “What’s new in Psychtoolbox-3?” Perception 36.

Li V, Castañón SH, Solomon JA, Vandormael H, Summerfield C (2017) Robust Averaging Protects Decisions from Noise in Neural Computations. PLoS Comput Biol 13:e1005723.

Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. J Neurosci Methods 164:177–190.

Moran R, Tsetsos K (2018) The standard Bayesian model is normatively invalid for biological brains. Behav Brain Sci 41.

Noguchi T, Stewart N (2018) Multialternative decision by sampling: A model of decision making constrained by process data. Psychol Rev 125:512–544.

Oostenveld R, Fries P, Maris E, Schoffelen JM (2011) FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Comput Intell Neurosci 2011.

Pisauro MA, Fouragnan E, Retzler C, Philiastides MG (2017) Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous EEG-fMRI. Nat Commun 8:15808.

Sauseng P, Klimesch W, Stadler W, Schabus M, Doppelmayr M, Hanslmayr S, Gruber WR, Birbaumer N (2005) A shift of visual spatial attention is selectively associated with human EEG alpha activity. Eur J Neurosci 22:2917–2926.

Spitzer B, Waschke L, Summerfield C (2017) Selective overweighting of larger magnitudes during noisy numerical comparison. Nat Hum Behav 1:1–8.

Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Group Bayesian Model Selection for Group Studies. NeuroImage 46:1004–1017.

Summerfield C, Tsetsos K (2015) Do humans make good decisions? Trends Cogn Sci 19:27–34.

Tickle H, Speekenbrink M, Tsetsos K, Michael E, Summerfield C (2016) Near-optimal integration of magnitude in the human parietal cortex. J Cogn Neurosci 28:589–603.

Tsetsos K, Chater N, Usher M (2012) Salience driven value integration explains decision biases and preference reversal. Proc Natl Acad Sci 109:9659–9664.

Tsetsos K, Moran R, Moreland J, Chater N, Usher M, Summerfield C (2016) Economic irrationality is optimal during noisy decision making. Proc Natl Acad Sci 113:3102–3107.

van Vugt MK, Simen P, Nystrom LE, Holmes P, Cohen JD (2012) EEG oscillations reveal neural correlates of evidence accumulation. Front Neurosci 6:1–13.

von Lautz A, Herding J, Blankenburg F (2019) Neuronal signatures of a random-dot motion comparison task. NeuroImage 193:57–66.

Von Neumann J, Morgenstern O (1944) Theory of games and cconomic behavior. Princeton University Press.

Wald A, Wolfowitz J (1949) Bayes solutions to sequential decision problems. Proc Natl Acad Sci USA 35:99–102.

Wyart V, Myers NE, Summerfield C (2015) Neural mechanisms of human perceptual

choice under focused and divided attention. J Neurosci 35:3485–3498.