Probing the role of sequential sampling and integration in decisions about protracted, noiseless stimuli

Hadiseh Hajimohammadi[1], Kieran S Mohr[1], Redmond G O'Connell[2], and Simon P Kelly[1]

[1]School of Electrical and Electronic Engineering and UCD Centre for Biomedical Engineering, University College Dublin, Belfield, Dublin 4, Ireland;

[2] Trinity College Institute of Neuroscience and School of Psychology, Trinity College Dublin, Ireland

Correspondence: hadiseh.hajimohammadi@ucd.ie and simon.kelly@ucd.ie

Abstract

Perceptual decision behaviour is known to be well-captured by models based on sequential sampling and temporal integration, but doubts have been raised about the generality and identifiability of these operations. Here we used neurally-constrained modelling to probe their role and temporal extent in the uncertain case of perceptual judgments about long-duration stimuli with no physical noise but weak evidence. We found that accuracy steadily improved across four covertly-manipulated evidence durations, indicating protracted sampling, but these delayed behavioural reports alone were insufficient to establish the operation of integration or of a decision-terminating bound. We then elaborated the models to prescribe how they would generate decision variable signals as well as choices, and fit them additionally to the average dynamics of a centroparietal electroencephalographic signal ('CPP') that traces decision formation. This established the setting of a bound and ruled out some non-integration mechanisms. However, one extrema detection model, which evokes a stereotyped signal 'flagging' the first bound-exceeding sample, rivalled the integration model in reproducing the evidence-dependent buildup dynamics of the CPP, alongside behavioural data. Moreover, the two models captured different features of neural motor preparation signals but neither captured all of them. We discuss the implications for the generality of integration and the technical challenges of neurally-informed modelling given limited behavioural data.

Introduction

In the most prominent models of perceptual decision-making, evidence is sequentially sampled and integrated over time into a decision variable that can be subjected to a criterion to generate a choice (Gold & Shadlen, 2007). These models provide an intuitive means to achieve optimal performance despite perceptual noise and provide excellent fits to behaviour on a wide range of tasks (Ratcliff et al., 2016). Perfect integration over long timescales (tens of seconds) has been demonstrated using expanded judgment tasks where evidence is provided in just a few discrete, highly stochastic tokens (Hyafil et al., 2023; Waskom & Kiani, 2018), and extensive behavioural and neurophysiological work points to a key role of temporal integration in speeded judgments of continuously-presented, noisy stimuli like dot-motion (Gold & Shadlen 2007; Steinemann et al 2024). However, outside of these experimentally-tractable, stochastic and/or noisy task cases, an extended window of temporal integration is less clearly motivated theoretically (Uchida et al., 2006; Thura et al 2012) and can be more difficult to identify due to a lack of constraints (O'Connell et al 2018). Here we examine one such task situation, in which one must conduct subtle feature comparisons of stimuli without physical noise to guide actions to be made at a later time. The absence of noise in the physical stimulus, on one hand, would seem to minimise the benefits of time-averaging, but on the other hand integration may still be needed to overcome internal neural noise associated with encoding a subtle sensory feature.

Not only is the use of integration uncertain for noise-free stimuli, but the temporal extent of integration is uncertain even for noisy stimuli. In some cases sensory information is found to be strongly choice-predictive only over a restricted time range (e.g. Ludwig et al., 2005), and across experimenter-controlled, variable stimulus durations, improvements in accuracy are distinctly less than would be predicted if decision makers utilised all evidence through continued sampling (Kiani et al., 2008). Such early decision termination is also expressed in neurophysiological decision signals reaching their peaks long before the end of noisy stimulus durations of >1 sec (Twomey et al., 2016; McCone et al 2026). It is not clear whether this reflects a general tendency to time-limit integration due to capacity limits or effort-saving, or the strategic setting of a low bound because the prevalence of strong (Twomey et al 2016) or short (Kiani et al 2008) stimuli meant that there was relatively little to be gained by extended sampling.

Here, we wished to test whether decision makers generally tend towards extended integration even in the absence of stimulus noise, under conditions that in principle would encourage it, where the stimulus is always long and usually difficult. We thus devised a contrast-comparison task in which participants view a fixed 1.6 seconds of noise-free stimulation on all trials before being asked to make a response, and covertly manipulated evidence duration within this stimulus window by seamlessly stepping the contrast difference back to zero at a variable time without the participants' knowledge, and in a way that could not be confidently detected. Although this covert duration manipulation allowed us to probe the temporal extent of sampling without the participants tailoring their bound-settings to expected short stimulus durations, a drawback of the deferred action choices is that only the outcome and not the timing of decision termination is available in behavioural data, thus limiting inferential power and model constraints.

A key technical challenge that limits the ability to identify the operation of integration through behavioural modelling is the problem of model mimicry, where data generated from one mechanism can be quantitatively explained equally well by another mechanism (Evans & Wagenmakers, 2020; Khodadadi & Townsend, 2015; Trueblood et al., 2021). For example, even for dot-motion discrimination, non-integration strategies such as extrema detection have been shown to provide competitive fits to key behavioural data patterns classically associated with integration, including improvements in accuracy with duration and reduction of reaction time (RT) and error rate with increasing evidence strength (Stine et al., 2020). An emerging approach aimed at overcoming such mimicry issues is to further constrain the models with neurophysiological signals that reflect evidence-dependent decision formation

dynamics (Hanks et al 2014; Purcell & Palmeri, 2017; O'Connell et al., 2018). Such decision signals have been reported in multiple brain areas of multiple species, most commonly in neural signals associated with preparing the decision-reporting actions (Gold & Shadlen, 2001; Hanks & Summerfield, 2017; O'Connell & Kelly 2021). In human electrophysiology, evidence-dependent buildup dynamics are observed not only in signatures of motor preparation such as decreases in Mu/Beta (8-30 Hz) amplitude contralateral to the upcoming movement (Pfurtscheller & Da Silva, 1999; Donner et al 2009; de Lange et al 2013; Murphy et al 2015), but also in a distinct, motor-independent Centro-Parietal Positivity (CPP) in the event-related potential (ERP). The CPP onsets at a time locked to evidence onset and builds at a rate that scales with evidence strength (Kelly & O'Connell 2013), dynamically tracking perturbations in the evidence, even in the absence of motor requirements (O'Connell et al., 2012). When decisions require immediate responses, the CPP peaks and falls at a time closely aligned to single-trial RT, with an amplitude that co-varies with decision accuracy (Kelly et al., 2021; Steinemann et al., 2018), and when responses are deferred until after a protracted stimulus, the CPP's peak and fall during stimulus viewing marks early decision termination (Twomey et al. 2016; Pares-Pujolras et al., 2025; McCone et al., 2026). We thus used CPP dynamics alongside behavioural choice accuracy across evidence durations to examine the temporal extent of sampling in the present task. Further, we devised non-integration models based on extrema detection that could viably generate the observed buildup dynamics of decision formation as well as choice accuracy trends, to examine the extent to which joint CPP-behaviour fitting can uniquely distinguish integration from non-integration mechanisms under such conditions.

We found that behavioural accuracy steadily improved with evidence duration, including from 0.8 to 1.6 sec, and that the CPP remained elevated up to the end of the stimulus period, together indicating protracted sampling. Behavioural accuracy on its own could be quantitatively well-captured not only by integration (drift-diffusion) but also by several extrema detection models, and neither model class identified the setting of a decision bound. CPP dynamics in occasionally-presented, trivially easy trials indicated the setting of a bound, and models jointly fit to CPP waveforms and choice accuracy confirmed this, and ruled out some non-integration models. However, contrary to expectations, an extrema detection model that evoked a stereotyped signal 'flagging' decision termination at the first bound-crossing produced a fit that rivalled that of integration and reproduced the slow ramping build-up of the trial-averaged CPP and its modulation by evidence duration. All model classes benefitted from the inclusion of a collapsing decision bound and/or starting point variability to achieve a good fit, evidence for which was also observed in empirical motor preparation signals. While the degree of these features estimated by the Integration model was closer to that observed in the empirical signals, the Extremum-flagging model predicted drift rate modulations by difficulty that were closer to estimates from the empirical signals. We go on to discuss the relative plausibility of these predictions and potential implications of these findings for other task conditions such as those where immediate responses provide information on decision termination timing, and what might be required of such an extremum-flagging model in order to further explain such data.

Results

We recorded EEG from 16 human participants during a task where they viewed two oppositely oriented, overlaid gratings flickering in alternation with no noise or stochasticity. After a 1.2-sec baseline period of equal, 50% contrast, the gratings underwent an equal and opposite contrast change, and after an additional 1.6-sec stimulus-assessment period, participants reported by mouse button press which grating had higher contrast. In 80% of trials, the contrast differential was small (±10%) and lasted for 0.2, 0.4, 0.8, or 1.6 sec before stepping back to equal contrast and ultimately disappearing at 1.6 sec. For the remaining 20% of trials, a trivially easy (±40%) contrast difference was presented for 1.6 sec. The participants were not informed of the evidence duration manipulation, and this long, easy condition

both reinforced the belief that all trials contained 1.6 sec of evidence (Figure 1A) and provided a means to probe the setting of a decision bound causing early terminations.

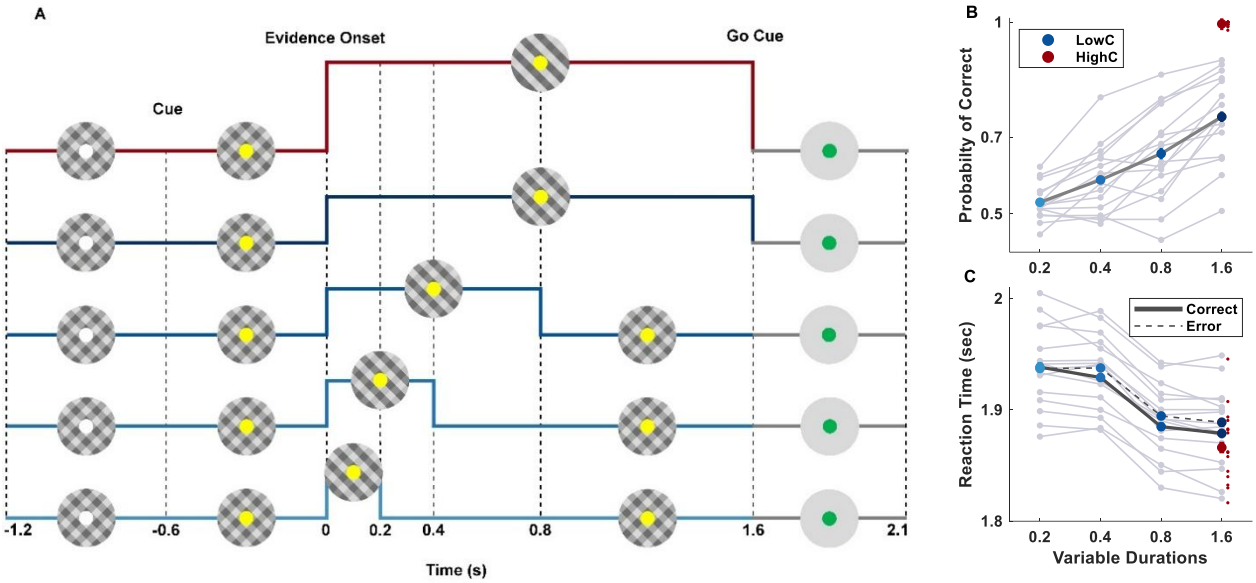

Figure 1. Contrast comparison task and behavioural results. (A) Stimulus and evidence time courses: The stimulus consisted of two interleaved gratings flickering on and off in alternation, whose relative contrast had to be judged. 0.6 sec after achieving fixation and triggering stimulus onset, the fixation point turned yellow to cue participants to prepare for the onset of contrast-difference evidence another 0.6 sec later. In all trials, the stimulus remained on for a full 1.6 sec following evidence onset, after which it disappeared and the fixation turned green to prompt the subject to make a left/right-hand button press to indicate that the left- or right-tilted grating had higher-contrast, within a 0.5-sec deadline. In 80% of trials, a weak (±10%) contrast increment/decrement was oppositely applied to the two gratings for 0.2, 0.4, 0.8, or 1.6 sec, returning to equal baseline contrast thereafter until stimulus offset. In the remaining 20% of trials, an easy contrast increment/decrement of ±40% was applied for the entire 1.6-sec period. (B) Accuracy significantly increased with the duration of the stimulus for low-contrast trials (blue dots) and was at the ceiling for the high-contrast condition (red dot). Individual participants' accuracies are shown in grey lines. (C) The mean reaction time (RT; relative to evidence onset) slightly decreased in trials with longer durations and higher contrast levels. Correct responses (solid, with individuals also shown in grey) were also slightly faster than error responses (dashed) in this delayed response task. For b and c, data are mean ± s.e.m. after between-participant variance was factored out.

Accuracy varied significantly across the five conditions (rmANOVA, $F(4,60) = 150.47$, $p < 4.4 \times 10^{-16}$, Figure 1B) and increased with duration within the low-contrast conditions (rmANOVA, $F(3,45) = 50.19$, increasing trend $p < 5 \times 10^{-11}$; all three consecutive duration pairs $p < 0.002$; Table S1). Despite responses being deferred to stimulus offset, response time (RT) also significantly varied across conditions (rmANOVA $F=12.35.31$, $p=2.15 \times 10^{-7}$), showing a mild downward trend with increasing duration in the low-contrast conditions ($p < 2.48 \times 10^{-2}$), suggesting that subjects had not always reached commitment before the end of the stimulus period (Figure 1C; Table S2).

Models fit to behaviour only

We began by comparing an integration model to two classes of non-integration model in terms of their ability to capture the improvements in accuracy with increasing evidence duration. The Integration model entailed a simple drift-diffusion process (Ratcliff & McKoon, 2008; Figure 2A, top row, see Methods). The first non-integration model was an Extrema detection model, which entails sequential sampling but with the decision determined by an individual evidence sample reaching an extreme value (Hyafil et al., 2023; Stine et al., 2020; Waskom & Kiani, 2018; Figure 2A middle row). The second non-integration model was a Snapshot model, which entails neither integration or sequential sampling, instead basing the decision on a single randomly-chosen sample within the 1.6-sec stimulus-evaluation period (Stine et al, 2020). For all three models, we implemented both bounded and unbounded versions

to test for early termination. In the Integration model, choices were determined by the sign of the first sample to hit a bound, or if the bound was absent/not hit, the last sample. In three variants of the bounded Extrema detection model, choices were similarly determined by the sign of the first bound-crossing sample, and if not hit, choices were determined by i) a random 50/50 guess, ii) the sign of the last sample, or iii) the sign of the most extreme recorded sample ('Extremum-tracking'). The latter extremum-tracking mechanism applied on all trials in the unbounded Extrema detection model. Finally, in the Snapshot model, the bound, when present, determined whether or not the randomly chosen sample would determine the choice, requiring a random 50/50 guess when the bound was not exceeded by the sample.

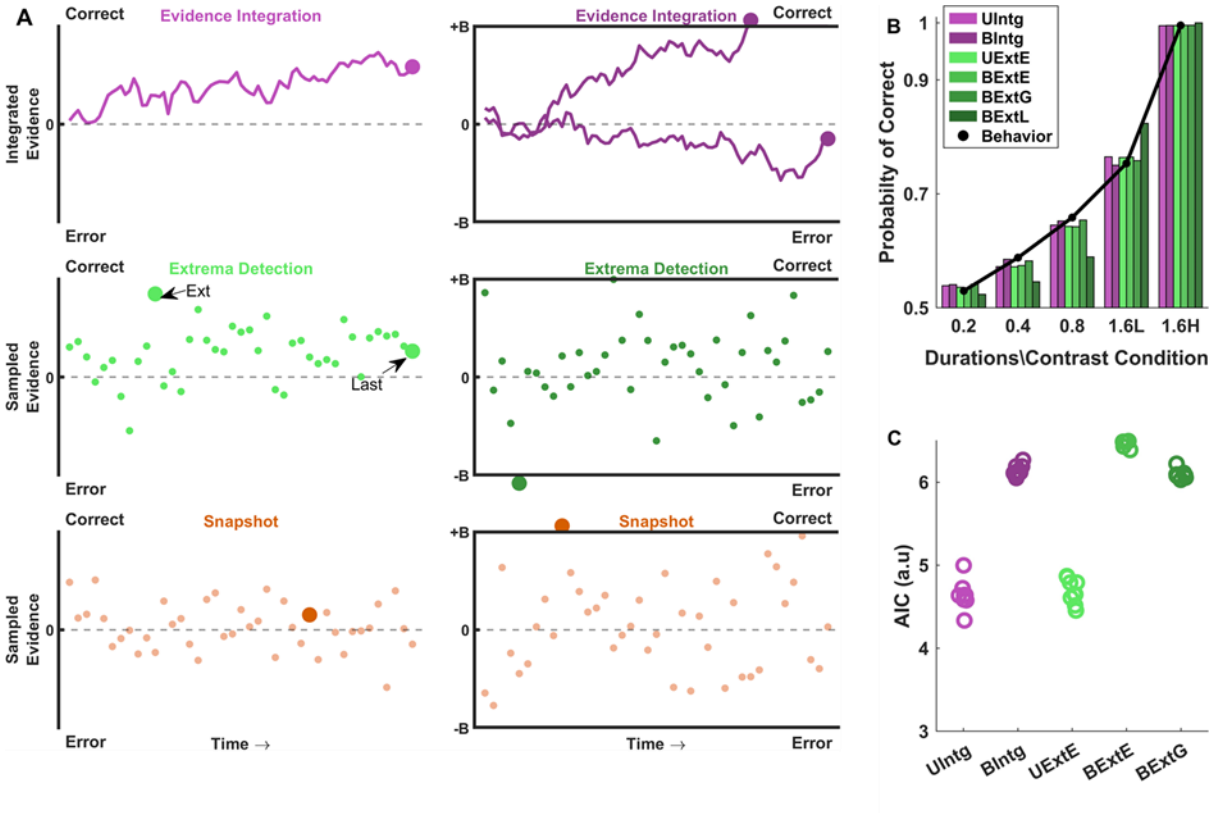

Figure 2. Alternative models and their fits to behavior only. (**A**) Model schematics. In the unbounded integration model ('UIntg;' top left), the sign of the accumulated evidence at stimulus offset determines the choice, while in bounded integration ('BIntg;' top right), the decision is either terminated early by reaching the correct or error bound or otherwise determined in the same way as in the unbounded model. In the unbounded Extremum-tracking model ('UExtE;' middle left), the decision is based on the sign of the evidence sample that attains the maximum absolute value, which by making use of the full stimulus-assessment period can produce accuracy improvements with increasing evidence duration. In the bounded Extrema detection model (middle right), a choice is determined as soon as a single sample exceeds a bound or, if neither bound is reached, the choice is determined by: i) a random guess ('BExtG'); ii) the sign of the last sample ('BExtL'); or iii) the sign of the sample that attains the most extreme value ('BExtE'). Finally, the Snapshot model, where the choice is based on a single sample chosen at random (uniform distribution over time), is shown in the bottom row. In the unbounded version ('USpst'), the decision is made based on the sign of the sample. In the bounded version ('BSpst'), the decision is based on whether the sample surpasses the bound; otherwise, the decision is guessed. (**B**) Observed (black dots) and model-predicted (bars) accuracy, showing that the five best-fitting models can capture the observed accuracy improvement with duration well. (**C**) AIC values across the 5 most competitive models show that models without bounds are favored due to parsimony, and selection among integration and non-integration mechanisms is largely inconclusive. Given that we computed G^2 values through Monte-Carlo simulation, we conducted the fits with 10 different instantiations of noise, each shown here as a separate point. These clustered points depict the variability of the G^2 values due to the stochastic nature of these simulations, underlining how the small margins among the best models in Table 1 are inconclusive (see Methods).

Table 1: Goodness of fit metrics and parameter estimates of models fit to behaviour only. Models beginning with U were unbounded, and beginning with B were bounded, with the bound parameter denoted by B. Drift rates for low and high contrast difference were denoted by D1 and D2, respectively; single-drift rate variants (denoted '_1d') of the unbounded models provided a worst-case benchmark for each mechanism. G^2 quantifies how closely the model quantitatively fits the data, while Akaike's Information Criteria (AIC) adds a penalty for complexity. For both metrics, lower values indicate better fits. Colour shading highlights the competitive models depicted in Figure 2C.

Models		AIC	G^2	D1	D2	B
Snap Shot	USpst_1d	96.02	96.02	3.97	-	-
	USpst	36.31	32.31	2.0	8.1	-
	BSpst	39.4	33.4	2.0	9.06	0.00007
Integration	UIntg_1d	64.4	64.4	0.088	-	-
	UIntg	4.53	0.53	0.057	0.20	-
	BIntg	6.14	0.14	0.056	0.226	0.128
Extrema Detection	BExtG	6.17	0.17	10.3	11.65	0.034
	BExtL	21.25	15.25	1.41	6.96	0.017
	UExtE_1d	65.01	65.01	0.569	-	-
	UExtE	4.54	0.54	0.382	1.59	-
	BExtE	6.49	0.49	0.385	1.59	0.045

The model fit estimates and metrics are shown in Table 1. The drift rate for the high-contrast condition was in all cases higher than for low-contrast, as expected, and non-integration models generally needed much higher drift rates than integration models, simply because the same accuracies must be achieved based on one rather than many samples (Hyafil et al., 2023; Stine et al., 2020). The Snapshot models provided far poorer fits than the two sequential sampling models. The Extrema detection model that relied on the last sample in the absence of bound crossing (BExtL was also not competitive, because it predicts a disproportionate step-up in accuracy from the 0.8 to 1.6 sec duration where the last sample becomes informative (Figure 2B). However, Extrema detection models with a default guessing mechanism (BExtG) and with the Extremum-tracking mechanism (UExtE and BExtE) fit the data just as well as the equivalent Integration models (Figure 2B,C; see Figure 2 - Figure Supplement 1 for bound crossing histograms). Further, the benefit of including a bound to fit the accuracy data was too small to warrant the additional complexity for any of the models (Figure 2C).

Decision-related Neural Signals

Since behaviour alone was insufficient to arbitrate between integration and multiple extrema detection mechanisms and was inconclusive regarding the setting of a bound, we turned to neural signatures of decision formation to further inform and constrain the models. We began by appraising qualitative features of the CPP (Figure 3B) alongside signals reflecting sensory evidence encoding (Figure 3A) and motor preparation (Figure 3C), which inform model structure. The differential steady-state visual evoked potential (SSVEP), driven by the contrast difference between the correct (higher-contrast) and

incorrect gratings, stepped from a zero baseline up to a level commensurate with the contrast difference, remaining stable for the duration of the evidence before falling back down to zero after evidence offset (Figure 3A). This stable contrast-difference encoding suggested that a constant (i.e. non-adapting) drift rate is a reasonable simplifying model assumption.

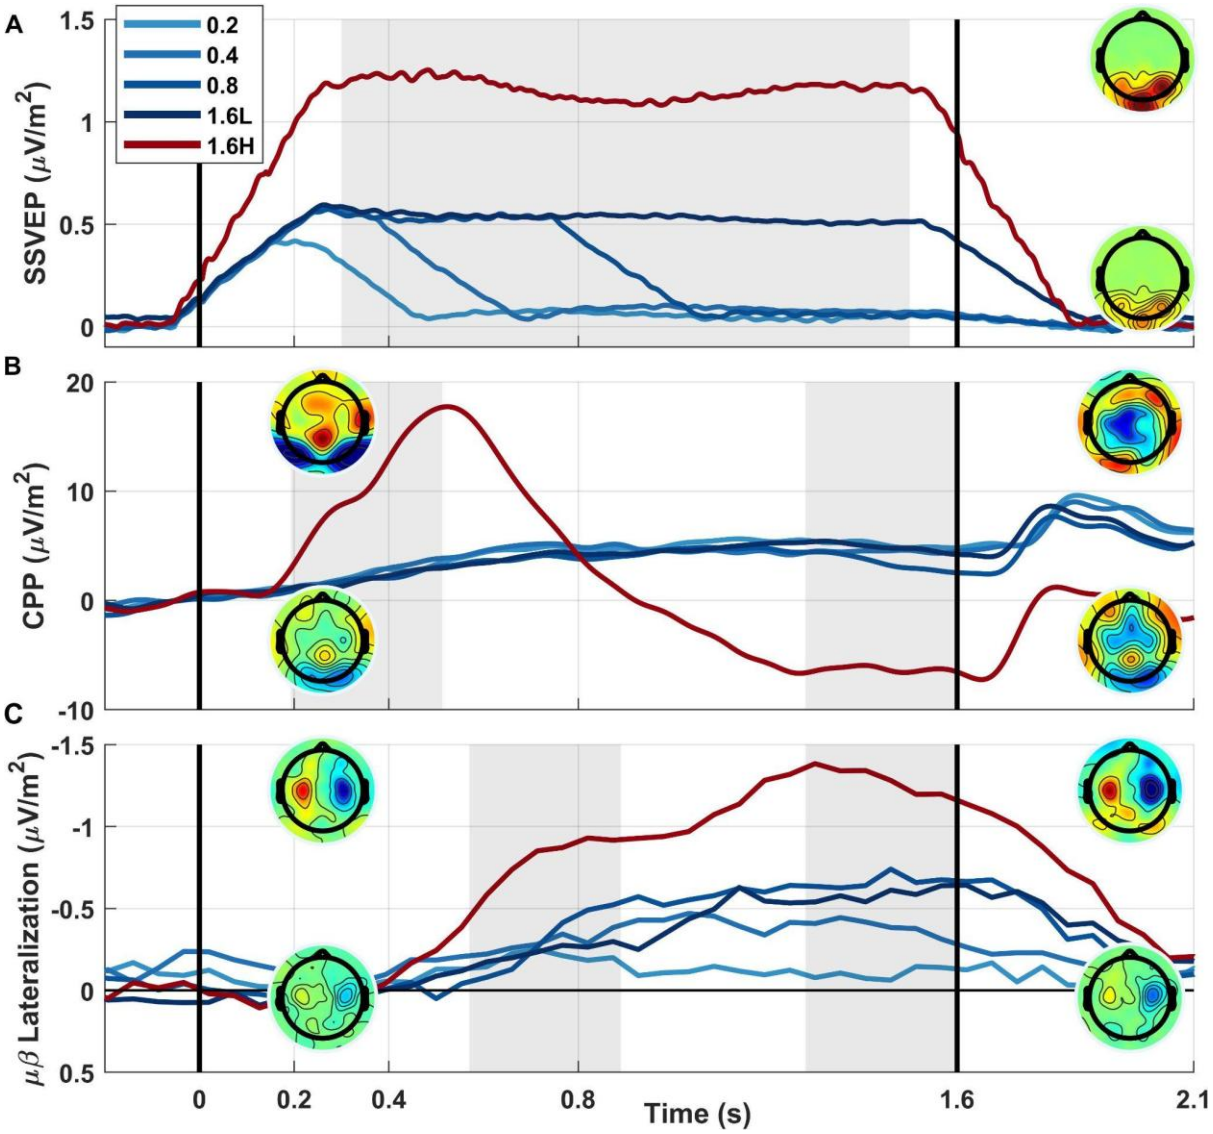

Figure 3. Neural Signals. (A) SSVEP showed a robust representation of sensory evidence encoding. Note that the signal ramps rather than suddenly steps simply due to the width of the window for spectral estimation (230 ms), and lags the physical contrast changes simply due to transmission delays to the visual cortex. Topographies show the average SSVEP amplitude from 300 ms to 1500 ms (shaded period) for the high contrast conditions (top) and the longest low contrast condition (bottom), with a primary focus at standard site Oz. (B) Event-related potential (ERP) at centroparietal sites averaged over trials (regardless of outcome) within each condition, showing that the centro-parietal positivity (CPP) undergoes a gradual buildup and sustained elevation for lower contrast conditions, and a peak at around 530 ms for the high contrast trials, at more than twice the amplitude. Topographies show ERP amplitude for the high contrast (top) and the average over low contrast conditions (bottom), for the shaded time ranges 210-530 ms (left) and 1280-1600 ms (right). (C) Motor preparation was measured by subtracting ipsilateral from contralateral mu/beta-band activity (MB; 8-30 Hz) with respect to the correct side. The topographies show the difference between trials favouring a left minus right response, again shown for both high contrast (top) and averaged low contrast (bottom) conditions in the shaded windows 570-890 ms (left) and 1280-1600 ms (right).

Turning to relative motor preparation reflected in Mu/Beta lateralisation (Figure 3C), we observed evidence-dependent buildup toward the correct alternative, remaining elevated until the delayed response as previously observed (Twomey et al 2016; McCone et al 2026), at a level that scaled with low-contrast evidence duration (rmANOVA, $F(3,45) = 5.8$, $p = 0.003$). At the level of decision formation, the CPP similarly exhibited an evidence-dependent buildup initially. However, as in previous studies (Twomey et al 2016; McCone et al 2026), the high-contrast condition did not sustain but rather peaked and fell to baseline (Figure 3B), indicating a bound setting that leads to early terminations in those trials, and enabling us to rule out unbounded models (Figure 2 - Figure Supplement 2). In the low-contrast conditions, the CPP exhibited a sustained elevation above baseline, indicating protracted sampling in line with the observed accuracy improvements. In contrast to Mu/Beta lateralisation, the final CPP amplitude, measured in the 320 ms window ending at evidence offset, did not vary significantly across the four low-contrast evidence durations (rmANOVA, $F(3,45) = 1.23$, $p = 0.3$). This lack of duration-dependence can be explained by the fact that whereas Mu/Beta lateralisation is a signed quantity whose polarity reflects the favoured alternative, the CPP always builds positively regardless of the favoured alternative, consistent with the sum of two underlying neural populations coding for the two alternatives (Afacan-Seref et al 2018). This means that noise variations do not cancel and the average CPP will exhibit a buildup even in conditions with weak or no evidence, blunting the distinction between noise-driven build-up and truly evidence-driven build-up for very weak evidence strengths (Steinemann et al., 2018). While all of these signal properties are consistent with previous observations, the data do not provide qualitative grounds to judge whether they arise from integration or non-integration mechanisms, which we next investigate through quantitative modelling.

Neurally-constrained models

We devised a decision-signal generating mechanism for each of the three competing bounded models that could not be resolved based on behavioural modelling alone, in order to constrain them further by comparing simulated trial-averaged decision variable (DV) waveforms with observed CPP waveforms. We chose the CPP to provide the neural constraints because one of our core interests was the temporal extent of sampling, and the CPP's peak and fall dynamics render it more diagnostic of early decision termination than Mu/beta, which sustains until response.

In the Integration model, the DV was computed as the perfect (lossless) integration of the momentary evidence samples over time with a fall-down back to zero following a bound crossing (Figure 4B). In the Extremum-tracking model, a running record of the most extreme evidence sample witnessed so far was encoded in the decision variable, again with a fall-down to zero following a bound crossing. In principle, this enables an early peak-and-fall in the high-contrast condition and a slower more sustained buildup in the low-contrast conditions (Figure 4C). Finally, for the Extrema detection model with guess default (BExtG), we simulated a DV involving no evidence-driven buildup dynamics on any single trial but rather a stereotyped signal marking each single-trial bound crossing, which when averaged across trials could approximate a gradual buildup. In this "Extremum-flagging" model, we assumed that bound-crossings triggered a Half-Sine-shaped (upper half; Weindel et al 2025) temporal function with an amplitude and width that was fixed across trials (Figure 4D), which would be expected to cluster at early latencies in high-contrast trials and spread over longer latencies in low-contrast trials, thus plausibly approximating the observed CPP dynamics. Though we use a neutral 'flagging' function here, the idea is consistent with classic theories of the closely-linked P300 or P3b component (O'Connell et al., 2012; Twomey et al., 2015), which posit that it reflects a process triggered by completion of a perceptual judgment, such as "context updating" (Donchin & Coles 1988) or response facilitation (Nieuwenhuis et al., 2005; Verleger, 2020). If no bound was reached by stimulus offset, there was no flag signal elicited, and the subject made a random guess. Since the half-sine flagging signal inherently incorporates a quarter-sine-shaped fall-down, we also used a quarter-sine-shaped fall-down in the Integration and Extremum-tracking models for consistency.

To constrain the models to fit CPP waveform dynamics in addition to accuracy, we added a penalty term quantifying the divergence between normalized real and simulated CPP waveforms (200-900 ms for high-contrast, 200-1600 ms for averaged low contrast; see Methods) to the goodness-of-fit metric to be minimised. We comprehensively examined the relative ability of the three model classes to jointly fit accuracy and CPP dynamics by fitting the models with not only the basic bound, drift rate and sine-width parameters, but also plausible additional parameters, tested one at a time. These included drift rate variability, starting point variability, sequential sampling duration (allowing premature quitting), sampling onset time (allowing sampling to begin before or after evidence is available), collapsing bound, and in the case of integration, leakage (see methods for full explanation of each). We fit all models with five neural penalty term weightings (w = 0.1, 1, 10, 100, 1000) to examine how the behavioural and neural fits change as a function of emphasis placed on reproducing the CPP dynamics.

In Figure 4 we compare simulations of the three model classes including both starting point variability and collapsing bounds, which helped achieve a good joint fit in all cases (Figure 4 - Figure Supplement 1). AIC derived from the behavioural fit element of the objective function generally increased with neural weighting (Figure 4E), demonstrating an expected tradeoff where the higher the emphasis on a close fit to CPP waveforms, the more the behavioural fit is compromised. Looking at the overall objective function (G^2 plus the neural penalty term; Figure 4F), the Extremum-tracking model clearly could not achieve a competitive fit, but the Integration and Extremum-flagging models achieve similarly close fits. Although the Integration model appeared to gain an advantage at extremely high neural constraint weightings (Figure 4F), this came with a greater sacrifice to behavioural fit (Figure 4E). We traced the neural-behavioural trade-off by plotting dissimilarity between simulated and observed data for the neural fit against the behavioural fit (Figure 4G). When endowed with starting point variability and collapsing bounds, both the Integration and Extremum-flagging model were able to achieve a better neural fit without sacrificing behavioural fit, compared to ‘basic’ versions of the models without those features. However, which of the two models gives the overall better fit depends on the relative importance one places on fitting behaviour versus neural dynamics. This is reflected in the fact that the trade-off curve for neither model lies fully inside (closer to the origin) that of the other model, which would have indicated a winner independent of neural constraint weighting. We accordingly evaluate the models in terms of their ability to capture key features of the neural data. The average DV dynamics simulated from each model illustrate that these two models could both qualitatively produce average CPPs with a shallow incline and continued elevation that varied little across low-contrast durations, while also producing a steep initial rise and early peak in the high contrast condition (Figure 4K,M; Figure 4 - Figure Supplement 2). By contrast, the Extremum-tracking model could produce a peak-and-fall morphology in high-contrast, and sustained elevation in low-contrast conditions, but could not do so while retaining a gradual evidence-dependent buildup (Figure 4L, Figure 4 - Figure Supplement 2). The DV waveform simulations in Figure 4, as well as the estimated parameters and fit metrics in Table 2, are for a neural constraint weight of w = 10, for which both competitive models achieved a good fit to CPP dynamics without substantively sacrificing behavioural fit.

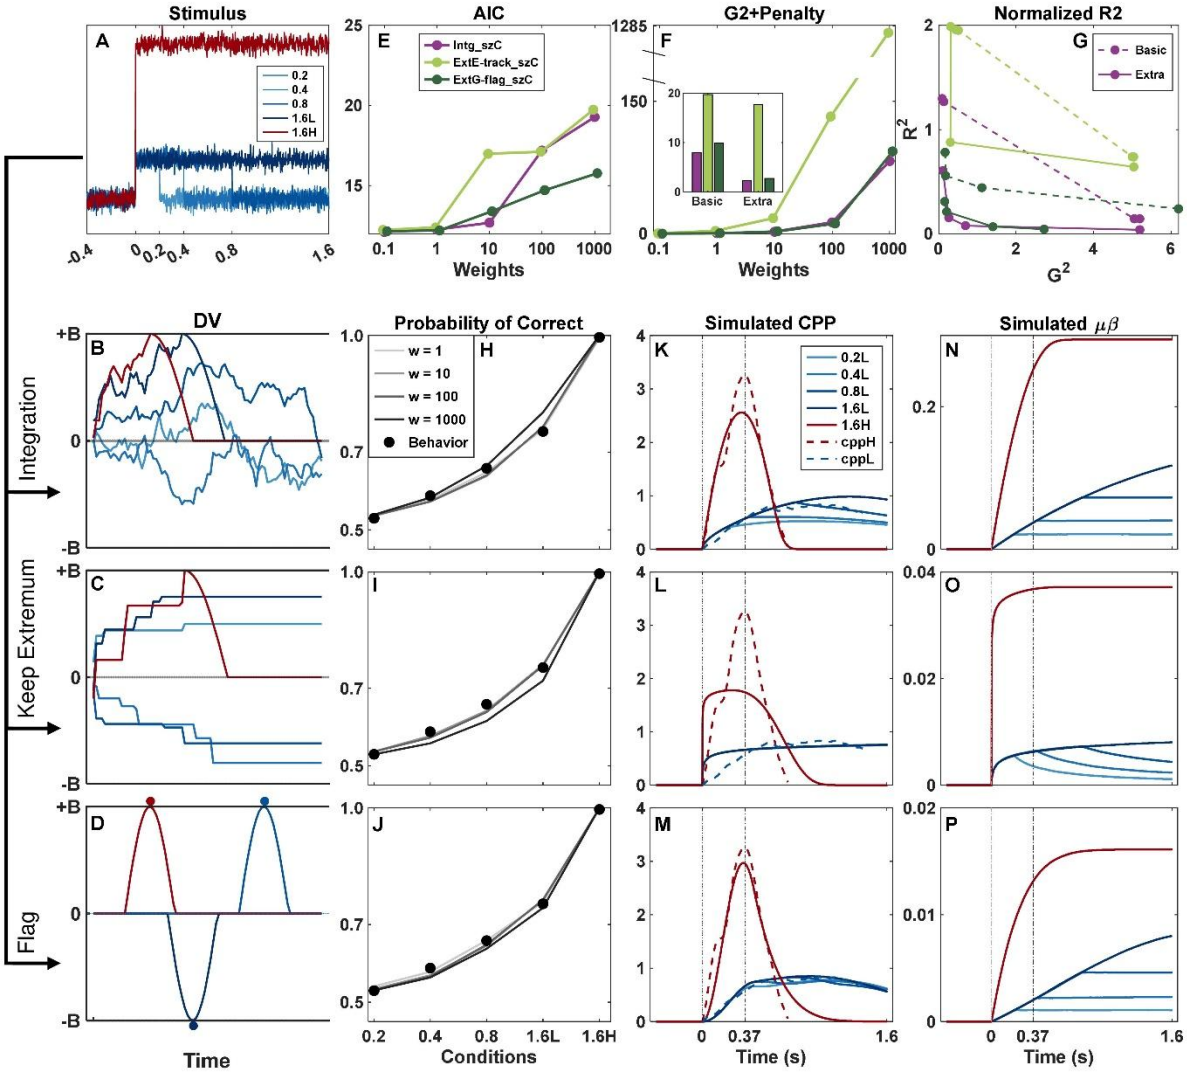

Figure 4. Neurally constrained models and their predictions. Among the models with various additional parameters, starting-point variability and collapsing bounds stood out as helping to achieve good CPP waveform fits without substantially sacrificing behavioural fits (Figure 4 - Figure Supplement 1), and given that both of these features were exhibited empirically in motor preparation signal effects (Figure 5), we show these models here (but see Figure 4 - Figure Supplements 4-6 for basic versions of the models without the extra parameters). (**A**) Single-trial sensory evidence traces were simulated for each of the five conditions using boxcar functions (stepping up during positive evidence) with added Gaussian noise as in behaviour-only models (one example for each condition shown). (**B-D**) single-trial examples of the decision variable in each of the three competing models, again simulating one trial per condition. (**B**) In the **Integration model,** the evidence was integrated over time until it reached a bound, after which it fell back to zero. Where the decision variable did not hit the bound the decision was made based on the sign of accumulated evidence at 1.6 sec. (**C**) In the **Extremum-tracking** model, the DV keeps track of the most extreme value so far, until it reaches the bound, after which it falls to zero. The sign of the extremum determined the choice, and the most extreme sub-bound value was used in the event that it did not cross the bound. (**D**) In the **Extremum-flagging** model, an all-or-nothing Half-Sine-profile signal marking choice termination onsetting at the bound crossing time (see main text). As in the other models, one example DV trace per condition is shown; for three of these (red and two dark blue traces) the bound was reached, while in the others it was not, so no signal was generated (flat light blue lines), and the choice was made as a random guess. In all 3 models in this figure, starting point variability and collapsing bound parameters are included, and they have an equal number of free parameters. (**E**) Akaike's Information Criterion (AIC) derived from only the behavioural component of the neurally-constrained fit, plotted as a function of the weighting of CPP-waveform constraint. With increasing emphasis on neural data in the fit, the fit to behaviour was compromised to varying extents, and the Extremum-flagging model was best at retaining good behavioural fits at higher neural weights. (**F**) Overall neurally-constrained

objective function (behavioural G^2 value plus weighted neural penalty quantifying divergence in observed vs simulated CPP) as a function of neural constraint weighting. The inset compares the objective function values with starting point variability and collapsing bound parameters (labelled 'Extra' in G) to the 'basic' models without either feature, at w = 10. (**G**) Neural fit (R^2) plotted against behavioural fit (G^2), illustrating a trade-off between how well the models can capture behaviour and the CPP up to w = 100. The closer the joint values come to the origin the better they are at simultaneously fitting both. (**H-J**) Predicted behavioural accuracy for a range of neural constraint weightings ('w'). (**K-M**) Simulated average CPP (where the DV is rectified on each single trial in keeping with CPP's always-positive nature; see Methods) for the best fits of each model at a neural weighting of w=10 (see Figure 4 - Figure Supplement 2 for simulations from all weightings, and Figure 4 - Figure Supplement 3 for predicted bound crossing frequencies). The simulated average CPP (solid lines) is shown for all five conditions. The empirical data is also shown (dashed lines) with the waveforms for the four low-contrast conditions averaged together, as they were for the neurally-constrained fitting. (**N-P**) Simulated differential motor preparation. Same as (K-M) except showing the trial-averaged differential DV (without taking the absolute value) and having the signal sustain at the bound level after it is crossed.

Table 2 Fit metrics and parameter estimates for neurally-constrained models with a neural-constraint weighting of w = 10. B refers to the bound; D1 and D2 refer to the drift rates in the low- and high-contrast conditions, respectively. FlagWidth refers to the width of the half-sine, Sz refers to starting-point variability, and K refers to the temporal slope of the collapsing bound (see Tables S3 for fits using a nonlinear collapse function, with similar findings). Models including starting point variability and collapsing bounds are distinguished from basic models without those features, by the model name ending '_szC.'

w = 10	G^2+Penalty	B	D1	D2	FlagWidth	sz	K
BIntg	7.9864	0.3591	0.0566	1.0991	0.6688	-	-
BIntg_szC	2.2533	0.4683	0.1607	1.4049	0.2502	0.9419	0.312
BExtE	19.7663	0.035	0.3815	10.865	0.9	-	-
BExtE_szC	17.7528	0.04	0.3813	12.3253	0.89	0.0004	0.3205
BExtG	9.9581	0.0145	0.5079	2.0663	0.52	-	-
BExtG_szC	2.7589	0.015	0.4387	2.1614	0.4462	0.000035	0.1062

Since model fits based on the CPP were equivocal, we evaluated the models further by simulating relative motor preparation from the best fitting parameters found in the CPP-constrained fits of all models (assumed to sustain at bound level once crossed). These exhibited dynamics that were qualitatively similar to the empirical MB lateralisation signal, including a steeper rise and higher sustained level for the higher contrast difference and greater amplitude variations across low-contrast durations than seen in the CPP. Temporal profiles varied across models, however: While both the Integration and Extremum-flagging models exhibited sustaining across the low contrast conditions (Figure 4N,P) similar to the empirical MB signal, the Extremum-tracking model showed a drop beginning at evidence offset in each condition, since many trials would find a negative extremum arise by chance during the zero-evidence lead-out period (Figure 4O). Between the Integration and Extremum-flagging models, a difference is apparent in the degree of average buildup in lateralised motor preparation for the low-contrast conditions relative to the maximal level reached by the high-contrast condition, with the Extremum-flagging model reaching higher levels closer to those observed in the empirical MB signals (Figure 3C). Indeed, with increasing neural constraint weighting, the Integration model tends to estimate a proportional difference in drift rates for high- vs low-contrast conditions that increasingly exceeds that of the physical contrast differences (40% vs 10%, i.e. a factor of 4); at w=10, this proportional difference

is a factor of 8.7 for the Integration model, compared to 4.9 for the Extremum-flagging model, and a factor of 3.5 measured from the empirical, relative motor preparation slopes in the window 500-700 ms (Figure 3C).

To further explore empirical Mu/Beta dynamics against model predictions, we plotted the contralateral and ipsilateral Mu/Beta signals separately, broken out by decision outcome (correct/error; Figure 5). In immediate-response tasks, this analysis approach has previously revealed that the two competing motor alternatives race each other up to a stereotyped level at the time of response that is invariant across stimulus conditions or reaction time (RT), indicating a fixed threshold associated with triggering the contralateral action (Corbett et al., 2023; Devine et al., 2019; Feuerriegel et al., 2021; Kelly et al., 2021; O'Connell et al., 2012). When evidence onset is anticipated, motor preparation reflected in Mu/Beta amplitude reduction contralateral to both hands has been found to begin building towards threshold even before evidence appears (Kelly et al., 2021), consistent with a dynamic 'Urgency' signal, defined as an evidence-independent buildup component that adds to cumulative evidence, gradually reducing the amount of cumulative evidence required to reach threshold and thus effectively implementing a collapsing decision bound (O'Connell et al., 2018). Moreover, fluctuations in pre-evidence Mu/Beta activity have been shown to predict trial-to-trial variability in choice and RT, in the way accumulator starting point variability does in models (Steinemann et al., 2018).

Here, despite the fact that the response had to be withheld until the post-stimulus response cue, MB-indexed motor preparation exhibited both of these effects (Figure 5). It rose significantly from -0.2 sec to 0 sec relative to evidence onset (main effect of time, $F(1,14) = 6.43$, $p = 0.02$), reflecting anticipatory urgency, before exhibiting evidence-dependent buildup after evidence onset, with contralateral preparation steepening in high-contrast trials, peaking at around 700 ms, and maintaining almost this level until the response period (Figure 5A). By averaging motor preparation traces across all trials regardless of condition or outcome, and extrapolating from a linear fit in the same baseline period, we could estimate the temporal profile of the collapsing bound, normalised relative to the initial bound height (Figure 5B, gray line). This could then be compared against the profile of collapsing bounds estimated from the models. Although both models showed a collapse that was less extensive than the empirical estimate, the Integration model provided a closer match than the Extremum-flagging model. Further examining Mu/beta lateralisation broken out by correct/incorrect outcome (Figure 5C), it is clear that in low-contrast trials, baseline levels of motor preparation already favoured the response that was eventually chosen before any evidence was presented (correctness x side interaction $F(1,14) = 8.85$, $p = 0.01$), reflecting significant starting-point variability that impacts eventual choice. As expected, the more difficult the trial (shortest duration), the greater the degree to which choice is linked to the baseline level randomly favouring one or the other alternative, and this is recapitulated in the Integration model (Figure 5D) but not the Extremum-flagging model (Figure 5E), the latter of which did not estimate appreciable starting-point variability at any neural constraint weighting (Tables S3, S5-8).

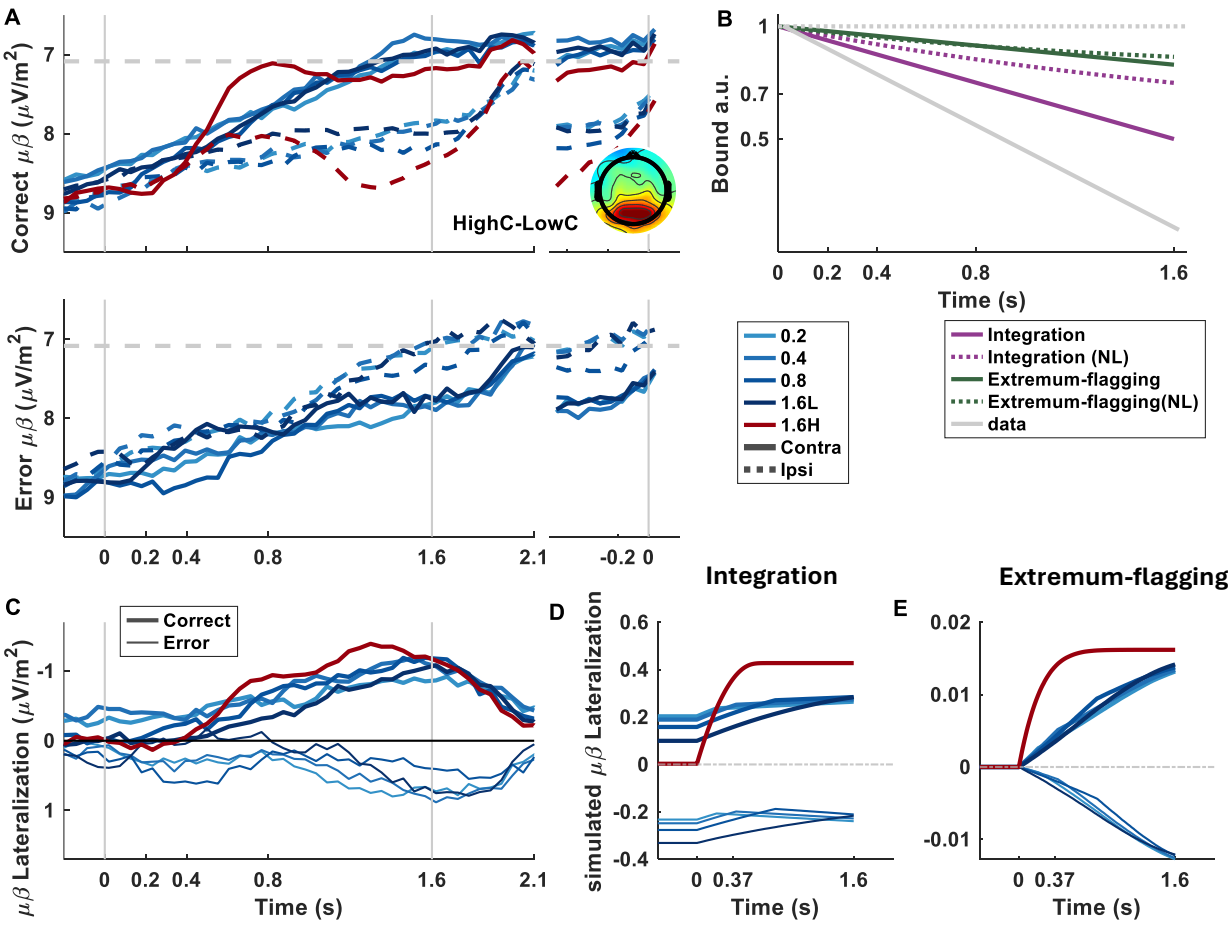

Figure 5. Observed and simulated motor preparation and bound dynamics. (**A**) Mu/beta (MB) decreases reflecting building motor preparation, contralateral and ipsilateral to the correct response, in correct (top) and error trials (bottom), shown aligned to evidence onset (left) and response (right). Contralateral traces, on correct trials, and ipsilateral traces, on error trials, reach a fixed threshold prior to response (marked by horizontal gray dashed line). Interestingly, on easy, high-contrast trials, this threshold (marked by the grey dashed horizontal line to aid comparison) is reached at approximately 800 ms and sustains at just below that level until the response. Note that on high-contrast trials, contralateral MB appears to reach a less decreased level prior to response compared to low-contrast trials. The topography inset shows the difference in MB between high-contrast and low-contrast conditions, revealing that this difference is unlikely motor-related, and more likely reflects greater attentional engagement in the harder low-contrast trials, known to be reflected in posterior alpha activity (Foxe & Snyder, 2011; Kelly & O'Connell, 2013). (**B**) collapsing bound estimated from the empirical data and the model fits. Linear (solid) and non-linear ('NL,' dashed) collapsing bounds were calculated for Integration and Extremum-flagging models. All traces were normalized to the initial bound height. (**C**) Observed Mu/beta lateralization for correct and error trials, demonstrating baseline choice-predictive effects that arise from starting-point variability. This can be compared to simulated average relative motor preparation at a neural constraint weighting of w = 10 for the Integration (D) and Extremum-flagging (E) models. Note again that models do not include non-decision delays, and so evidence-dependent dynamics begin at time zero in the simulated traces, whereas they are clearly delayed as usual in the empirically observed data. All simulations are from models containing both starting point variability and collapsing bounds.

Discussion

Sequential sampling models have long been known to provide excellent quantitative fits to decision behaviour on diverse tasks (Ratcliff et al., 2016), but the same ubiquity cannot be assumed of the role of protracted integration in practice, across all perceptual tasks (Cisek et al., 2009; Ludwig et al., 2005; Stine et al., 2020; Uchida et al., 2006). The present study examined a contrast-comparison task whose long, noiseless stimuli and delayed decision reports might seem to obviate protracted sampling, but whose subtle feature differences may be vulnerable to internal neural noise and hence benefit from

protracted integration. This emulates real-life situations where subtle features need to be judged in good visibility conditions in order to guide later actions, but is a situation that, due to the absence of information on decision timing and the dominance of one fixed evidence strength, offers very limited behavioural data for adjudicating between alternative mechanisms. Moreover, given the increasingly apparent task-dependence of neural decision mechanisms (Okazawa & Kiani, 2023; Parés-Pujolràs et al., 2025; Shushruth et al., 2022), one cannot simply extrapolate them from distinct task types such as immediate-response or expanded judgement tasks. Hence, to gather more information for model comparison while retaining the core task demands of interest, we manipulated the duration of the weak evidence without participants' knowledge, and traced the average neural dynamics of decision formation in the centro-parietal positivity (CPP) and spectral EEG measures of motor preparation. We found convincing evidence for protracted sampling: accuracy steadily improved with evidence duration, including a substantial increase from 800 to 1600-ms duration, and in difficult trials the CPP remained elevated throughout the stimulus period. CPP dynamics also established the operation of a bound, which was reached early and often for the rare strong-evidence trials, but less frequently and later for the weak trials. At the motor preparation level, pre-evidence, baseline biases toward the ultimately chosen alternative indicated the operation of starting point variability, and dynamic anticipatory preparation suggested a collapsing bound effect. However, quantitative modelling revealed that although the addition of neural constraints did offer many more insights than was possible from behaviour alone, it could not identify one uniquely viable account of the data; specifically, when allowing for starting point variability and collapsing bounds, similarly close fits to both behaviour and CPP dynamics could be attained by an integration model and a non-integration model based on extrema detection.

Extrema detection models were proposed as accounts of decision behaviour decades ago (Watson, 1979), and although integration-based models such as the DDM have dominated behavioural modelling in the intervening years due to their mechanistic appeal and versatility, it was recently shown that these two accounts are surprisingly competitive in their ability to explain a range of important behavioural effects (Stine et al., 2020). Here, we went a step further by considering plausible ways in which such an extrema detection model could produce a gradual buildup in its trial-averaged decision variable dynamics, and thereby predict neural decision signal dynamics. One mechanism, which continually tracked the most extreme sample encoded so far ('Extremum-tracking'), could fit behavioural data but could not simultaneously capture CPP dynamics because it predicts an initially steep and decelerating temporal profile that is not observed in the CPP (Figure 4L). A second mechanism, in which a stereotyped signal flagged single-trial bound crossings that could produce gradual buildup when averaged, could capture both behaviour and CPP dynamics to a degree comparable with Integration. Although this may be the first time that centroparietal signal buildup has been captured quantitatively through a non-integration mechanism, the general idea of centroparietal activity being triggered by decision completion was a part of many decades-old hypotheses regarding the function of the classic 'P3b' ERP component, centred on concepts of stimulus-response activation, context updating and memory storage (Polich 2007; Verleger et al 2020). However, while successful in this particular delayed-response paradigm, the degree to which such a process could account for decision signal dynamics previously observed on other tasks such as those with immediate-response requirements, remains unclear. In discrete-trial, forced-choice tasks, the response-locked CPP builds with a shallower slope and reaches a lower amplitude for weaker relative to stronger evidence, which in bounded accumulation frameworks is explained by evidence-dependent drift rates and a collapsing bound (Steinemann et al., 2018). The version of the Extremum-flagging model in the current study assumes a stereotyped signal triggered by each bound crossing and thus it should not show any such variation by evidence strength when aligned to the response; however, in principle, if one allowed for the possibility that the flag signal was smaller in amplitude for weaker stimuli (e.g. if it scaled with confidence) and more temporally jittered with respect to response execution, it could be possible to produce response-locked CPP waveforms that cross-over each other for high vs low evidence strength (because low-evidence builds more slowly

to a lower final amplitude) in the way that is observed. Whether this can be achieved with reasonable estimates of non-decision delays will have to be explicitly tested in future work.

While the Integration and Extremum-flagging models agreed on the operation of protracted sampling throughout the stimulus period unless a bound was reached, and the relative insensitivity of the simulated CPP to duration due to its always-positive nature, there were notable differences in some key quantitative predictions. The Integration model predicted a substantially higher ratio of drift rate for high versus low contrast (8.7) than the Extremum-flagging model (4.9), which was closer to the ratio of the contrast difference itself (a factor of 4). Plausible reasons for a drift rate ratio exceeding the contrast difference ratio are the well-known nonlinear contrast response function of early visual cortical neural activity and the fact that the high-contrast condition had a highly salient evidence onset to trigger attentive sampling, compared to the low-contrast condition. However, the ratio of buildup rates of relative motor preparation signals (3.5) is much lower. Although this could be partially ascribed to blunted buildup rate in the high-contrast conditions due to many early decision terminations and the blurring effect of time-frequency analysis, and blunted evidence-strength dependent buildup effects have been previously seen for delayed-response compared to immediate-response tasks (Twomey et al 2016; Roitman & Shadlen 2002), one could not be confident that this could account for the wide discrepancy observed here. Meanwhile, the quantitative degree of starting point variability and bound collapse was substantially greater in the Integration model, toward the levels exhibited empirically in motor preparation signals. Because the Extremum-flagging model relies on a precise bound setting on an evidence representation with high drift rate in order to explain the accuracy improvements across durations, it is highly sensitive to even a small bound collapse, and starting point variability effectively adds noise that cannot be absorbed by the precise bound setting, hence leading the Extremum-flagging model to favour a near-zero level of starting point variability. The inferred operation of collapsing bounds is intriguing from the point of view that responses were deferred in this task until after stimulus offset; however, given the prevalently weak evidence strength and tight 500-ms deadline after stimulus onset, it may be that subjects experienced urgency to commit to a choice in a timely fashion. Indeed, it has recently been shown that the dynamics of the Contingent Negative Variation in delayed response tasks is consistent with a cognitive-level urgency to terminate a decision even without the requirement of immediate response execution (McCone et al., 2026).

In the neurally-constrained modelling approach we adopted here, we conducted a sensitivity analysis where changes in the model fits were evaluated as a function of increasing emphasis on fitting the neural data. A potential shortcoming of this is that it is hard to announce a quantitative ‘winner’ among the models in cases such as this, as this depends on how much one favours behaviour versus neural data. We believe that this was the right approach nevertheless, because in principle, with such narrow margins it could be possible to make further tweaks to aspects of the models to produce slightly better quantitative fits and it would be hard to know when to stop in an unbiased way. One element, for example, that in principle could impact the relative quantitative fit of the models pertains to the fall-down dynamics of the CPP once a bound is reached. While we have access to this in immediate response paradigms, we do not in delayed-response paradigms and it is not clear that one can be readily extrapolated from the other. Aside from the uncertain shape of the CPP fall-down, which in the case of integration could be linear rather than the sine-shaped function used here for consistency with the Extremum-flagging model, fall-down dynamics would also be affected by how soon the accumulation quits or the flagging signal onsets following commitment (bound-crossing), and whether there is some amount of post-commitment sampling as has been observed in immediate-response tasks (Murphy et al 2015; Grogan et al 2023), all of which may have some impact on the morphology of the average waveforms. Relatedly, a larger number of subjects may have smoothened and refined the empirical average waveforms used to constrain the models, but the potential changes in quantitative fits are unlikely to have substantively altered the model comparison outcome, relative to the impact of tweaks related to such nuances as fall-down or post-commitment dynamics.

In summary, the present study presents a non-integrative sampling mechanism that can rival Integration in capturing not only behavioural accuracy of delayed perceptual decision reports about a protracted stimulus, but also decision variable dynamics reflected in neural decision signatures during that stimulus, which lends itself to further quantitative testing in a broader range of conditions.

Acknowledgments

H.H., K.S.M. and S.P.K. were supported by a Wellcome Trust Investigator Award (219572/Z/19/Z). R.G.O. was supported by Horizon 2020 European Research Council Consolidator Grant In Decision 865474.

Competing Interests

The authors declare no competing interests

Methods

Data Collection

Participants

Sixteen healthy human participants (9 males, 4 females (aged 18-40); the gender of three participants was not recorded) were invited to take part in our experiment, which was approved by the Human Research Ethics Committee for the Sciences, University College Dublin in accordance with the Declaration of Helsinki. Sample size was not based on formal power analysis, but was close to similar EEG studies. All participants had normal or corrected to normal vision with no history of psychiatric diagnosis, sensitivity to flickering light, or head injury. Participants gave us written and informed consent and were compensated at €10 / hour. All together the participants completed 188 blocks of 80 trials (15040 trials in total).

Contrast Discrimination Task

Participants were instructed to perform a delayed-response, two-alternative forced choice contrast discrimination task, programmed using PsychToolbox-3 (Brainard & Vision, 1997; Kleiner et al., 2007; Pelli, 1997) in MATLAB and presented on a gamma-corrected 51 cm Cathode Ray Tube monitor (Dell E771p) with a refresh rate of 75 Hz and resolution of 1024 768. Participants were seated at a viewing distance of 57 cm inside a dark, sound-attenuated booth with their heads stabilized in a chin rest with forehead support and their left and right thumbs resting on the left and right buttons of a symmetrically shaped computer mouse, which formed the response alternatives. Eye position was monitored continuously throughout the experiment using a remote eye tracker (EyeLink 1000, SR Research, 1,000Hz).

Stimulus and procedure. The stimulus contained two interleaved grating patterns with a spatial frequency of 1 cycle per degree tilted at −45 and +45 degrees relative to the vertical meridian (Figure 1A), which flickered on and off in alternation, with each presented for one screen refresh on every 4th refresh, thus generating two phase-opposed steady-state visual evoked potentials at 18.75 Hz that cancel each other when contrasts are equal (see below). A central fixation dot (0.2° diameter, grey) was presented, on which the participants fixated, followed 1200 ms later by the overlaid grating stimuli at 50% contrast (baseline period). After 600 ms of the baseline period, the fixation dot changed colour to yellow to warn the participants that evidence would be presented after another 600 ms. Evidence consisted of an equal-sized increment/decrement in the contrast of the two gratings such that one appeared brighter, which lasted for up to 1600 ms with the stimulus returning to baseline for shorter

durations until the stimulus was finally removed at the end of the 1600 ms period. Participants were not informed that the duration of evidence would vary, and the task structure encouraged a belief that it would always last 1600 ms. There were five equiprobable, randomly interleaved evidence conditions. In one easy condition, a high contrast difference (±40%, i.e., one grating 10% and the other 90%) was shown for the full 1600 ms, which served to reinforce the belief that evidence lasted for 1600 ms. The remaining four conditions consisted of a low contrast difference (±10%), which was presented for 200, 400, 800, or 1600 ms seconds. This contrast difference was difficult, with performance reaching 75% only on the longest duration trials (Figure 1B), so that participants were unlikely to detect with any confidence that the stimulus returned to baseline on the shorter duration trials (Figure 1A). This lack of explicit knowledge of variability in evidence duration is what differentiates our task from other variable-duration tasks (Kiani et al., 2008). Once the stimulus was removed, a go-cue instructed participants to indicate the higher contrast grating by clicking a mouse button with the thumb of the corresponding hand within 500 ms. The mouse was positioned flat and centrally in front of the participants, with fingers resting flat on the desk. Feedback was presented at the end of each trial in text and audio format. The text feedback informed the subjects whether they responded correctly ('Correct'), incorrectly ('Error'), prematurely (responded before the go-cue; 'Clicked too Early'), or missed the deadline of 500 ms to confirm their choice ('Too Late'). Although the short deadline was set on the response-cue to encourage the participants to indicate their decision immediately after the stimulus offset, since accuracy was the critical behavioural measure we analysed all responses up to 1 second after the go-cue. 0.4% of trials were missed (no button pressed up to 1 sec) and 2.6 % were premature responses (pressing before go-cue at stimulus offset), and these were excluded from calculating the accuracy. Participants received 10 points for correct, timely trials and 0 points for the other responses. At the end of each block, feedback was indicated on the mean accuracy and cumulative points gained during the block and participants were encouraged to attain as high a score as possible.

Training. Before performing the main experiment, participants conducted a practice version of the task to become familiar with the experiment and to be able to achieve the desired correct responses within the required deadlines. The training contained a series of short blocks of 20 trials with a fixed contrast difference shown for the full 1600 ms without any duration manipulation. The contrast difference started off with a high contrast difference of ±50%, dropping to ±20% with a 10% reduction in each block. They were required to achieve >90% in the easiest (±50%) and >70% in the hardest conditions (±20%), before progressing to the next difficulty level.

Statistical Analysis

Accuracies and response times were measured for each condition based on all trials with a valid response, which excluded 0.4% of trials that were missed (no response before 1s following the go-cue) and 2.6% of premature responses prior to the go-cue. Response times (RT) were recorded from the onset of evidence in milliseconds (ms).

One-way repeated measures analyses of variance (ANOVA) were carried out on the five task conditions, followed by pairwise comparisons (paired-sample t-tests), for RTs, accuracies, and neural amplitudes. For all statistical analyses, prior investigation was carried out to determine whether the data were suitable for parametric analysis by screening for any major violations of normality. Mauchley's test of sphericity was also conducted to assess the assumption of equal variances of the differences across all condition pairings in the rmANOVAs, and Greenhouse-Geisser corrections were applied when the sphericity assumption was violated.

EEG Acquisition and Preprocessing

Continuous EEG was acquired using a 128-channel ActiveTwo system (BioSemi) at a sample rate of 512 Hz. Vertical and horizontal eye movements were monitored using two vertical electrooculogram (EOG) electrodes placed above and below the left eye and two horizontal EOG electrodes placed at the outer canthus of each eye. Customized scripts in MATLAB (Mathworks, Natick, MA) and the EEGLAB toolbox (Delorme & Makeig, 2004) were used to analyse EEG data. EEG data were low-pass filtered by convolution with a 77-tap Hanning-windowed sinc function to give a 3-dB corner frequency of 38 Hz with strong attenuation at the mains frequency (50 Hz) (Widmann et al., 2015). Each block was detrended quadratically. Noisy and flat-lined channels were identified by eye by comparing

standard deviations to neighbouring channels and were interpolated based on surrounding channels using spherical splines. No high-pass filter was applied. Initially, EEG data were segmented into stimulus and response-aligned epochs. Stimulus-aligned epochs were extracted from stimulus onset to 960 ms after stimulus offset. All epochs were baseline-corrected relative to the interval of 106.6 ms pre-evidence onset (2 SSVEP cycles). We again applied spherical spline interpolation to any channel with excessively high variance with respect to neighbouring channels in each trial. Trials were rejected if an artifact occurred between evidence onset and stimulus offset (any channel with absolute magnitude >70 μV). One participant was excluded from EEG analyses due to excessive perspiration artifacts. Finally, to minimize the effects of volume conduction and avoid the possibility of overlapping signals confounding our interpretation of key signals such as the CPP (Kelly & O'Connell, 2013; Philiastides et al., 2014), a current source density (CSD) transformation was applied to each trial (Kayser & Tenke, 2006).

Electrophysiological data analysis

We extracted neural signatures of three processing levels of decision making: sensory evidence encoding, accumulation of evidence, and motor execution.

Neural signature of Sensory Evidence (Steady State Evoked Potential, SSVEP)

Flickering a stimulus at a certain frequency generates an oscillatory signal with the same frequency, known as the steady-state visual evoked potential (SSVEP), which provides a robust representation of physical contrast due to its narrow-band content (Regan, 1989). To obtain a direct measure of the sensory evidence upon which observers based their decisions (i.e. the relative contrast of the overlaid gratings) we generated a difference-SSVEP (d-SSVEP) by flickering the left- and right-tilted gratings in alternation at 18.75 Hz. This produced a signal that was zero, on average, for the equal-contrast baseline stimulus because the alternating gratings cancelled each other. When contrast differed, an aggregate d-SSVEP was produced with a phase dictated by the higher-contrast grating. The d-SSVEP signal was derived for each single trial by computing a Fourier Transform on untapered (boxcar) 320-ms time-windows (exactly 6 SSVEP cycles) in time-steps of 53.33 ms (1 cycle), and extracting the complex value representing oscillatory phase and magnitude at the 18.75-Hz frequency bin. We used the average complex oscillatory 18.75-Hz signal produced by the low-contrast 1600-ms condition to provide a common reference. This reference phase was calculated by trial-averaging the complex signal for the intermediate contrast level (S_{ref}). This was done independently for each electrode and time point. The angle of this resulting complex average was $_{ref}$ = arg(S_{ref}) then used as the reference phase to which all signal-trial data were aligned using phase rotation. Subsequently, these reference phases were subtracted from all trials, effectively rotating the phase data to correct for arbitrary offsets. One implication of the current approach is that we are unable to assess potential anticipatory modulations preceding the onset of the evidence. This means no reliable phase information is available for the pre-evidence period. However, this has minor concern in the current study, as our primary interest lies in modulations occurring during the evidence presentation. Oz electrode where the d-SSVEP was maximal on grand average was used to calculate this signal. The grand-average topography of the d-SSVEP was calculated by averaging from 300 ms to 1500 ms after stimulus onset (grey shading in Figure 3A) for the low-contrast, 1600-ms condition.

Neural signature of Evidence Accumulation (Centro Parietal Positivity)

Decision formation was examined by measuring the dynamics of the centro-parietal positivity (CPP), which has previously been shown to exhibit many properties predicted of evidence accumulation (O'Connell et al., 2012; Twomey et al., 2016). The evidence-locked CPP, baseline corrected to 106.6ms pre-evidence onset, was measured from a subset of electrodes chosen from a broad cluster of electrodes over centroparietal scalp (centred around Pz) at which the grand-average, high-contrast CPP was maximal at its peak latency (513 ms). 6 electrodes where the signal had the highest amplitude were selected from this cluster for each subject and the CPP was calculated by averaging these electrodes (see (Geuzebroek et al., 2023) for similar approach).

Motor Level decision threshold (Mu-Beta-Band spectral amplitude)

As in previous work (de Lange et al., 2013; Donner et al., 2009; Kelly & O'Connell, 2013; Steinemann et al., 2018), effector-selective motor preparation (Hanks & Summerfield, 2017; Kelly & O'Connell, 2015) was measured in the reduction of amplitude across the Mu/Beta bands (8-30 Hz) over the motor cortex (electrodes C3 and C4). Here, we calculated Mu/Beta amplitude using a short-time Fourier transform with a boxcar window size of 320 ms and time-steps of 53.33 ms, averaging across frequencies in the 8-30 Hz range but omitting 18.75 Hz to exclude the SSVEP. To produce a relative motor preparation signal that would reflect differential evidence for the decision, lateralized Mu/Beta was calculated by subtracting ipsilateral from contralateral amplitude, with respect to the hand corresponding to a correct answer on each trial. In the analysis of Figure 5, we computed average waveforms for correct and error trials separately, for the separate contralateral and ipsilateral signals without subtraction. The mean amplitude of contralateral signals across conditions at the time of response was taken to represent the 'motor threshold' (gray line in Figure 5A).

Model Fits to behaviour only

We explored five decision-making models, of which one was integration, three were extrema detection, and one was a snapshot strategy (Figure 2A). In all models, evidence samples were drawn from a normal distribution with a positive mean ("drift rate") during the evidence period and zero mean during the equal-contrast periods, with positive/negative values representing samples favouring the correct/incorrect alternative. What differs among these models is how the decision variable was formed from this evidence and used to determine the choice, as explained below. Evidence in all models had a standard deviation of 0.1 arbitrary units, which acted as the scale reference point. Core free parameters were common across models to enable direct comparisons of their fitting performance. These consisted of a separate "drift rate" for high and low contrast trials, dictating the strength of sensory evidence, and a "bound" parameter to set positive and negative criterion levels that the DV needed to reach to trigger a correct and incorrect decision, respectively.

All models were fit to the accuracies across the five experiment conditions (one 1600-msec high-contrast condition and four low-contrast conditions with durations of 200, 400, 800 and 1600 ms) by minimizing the χ^2 based statistic, G^2 (Ratcliff & Smith, 2004), described in detail below. The input sensory evidence signal for all models consisted of 1600 ms of Gaussian noise, plus a mean that stepped up and down according to a boxcar function during the periods of evidence to a height determined by the drift rate. To simulate accuracies across the five conditions, Monte-Carlo simulations of such trials were generated for a given set of parameter values and these parameters were optimized to match the empirical accuracies by minimizing the objective function (G^2) using a SIMPLEX algorithm (MATLAB's *fminsearch* function), generating a new Monte-Carlo simulation at each step of the optimization process. The G^2 statistic captured the difference in trial proportions between the real and simulated data for each of the five conditions, and was calculated using the equation:

$$G^2 = 2n\left(\sum_{r=1}^{2}\sum_{c=1}^{5} p_{r,c} log \frac{p_{r,c}}{\pi_{r,c}}\right)$$

Where n is the number of trials per condition, pr,c and r,care the observed and predicted proportions of responses r (correct/error) for each of the five experimental conditions, c. Models with lower G^2 values fit the data more accurately.

The Akaike Information Criterion (AIC) was used to evaluate the goodness of fit (Kelly et al., 2021) while penalising for model complexity, calculated from G^2 using the equation:

$$AIC = G^2 + 2f$$

where f is the number of free parameters.

To find the best-fitting parameters for a given model, we took three steps: 1) The G^2 was evaluated for a large grid of parameter vectors using 1000 simulated trials per parameter set, within the parameter limits listed in Table S4, with 10 points along each parameter dimension; 2) Next, we identified the grid point with the best G^2 value for each model and ran the SIMPLEX algorithm with 20,000 trials per iteration starting from that parameter vector to refine the parameter estimates. To facilitate the

convergence of the SIMPLEX algorithm, the random number generator was set to a fixed seed across all model simulations in a given fit. Then in order to illustrate the degree of uncertainty associated with this monte-carlo simulation approach (Figure 2C), we repeated the fits (step 1 and 2) with 10 different random seeds (i.e. 10 different instantiations of noise). 3) Starting from the parameter vectors attaining the lowest G^2 in the previous step, the parameter estimates were refined by repeating each instantiation of the SIMPLEX algorithm with 100,000 trials per iteration, with the final parameter estimate and associated G^2 computed from a new random seed taken as the optimum. This approach starting with a grid search has the advantage of ensuring a comprehensive exploration of the parameter space. We expanded the parameter value limits (Table S4) whenever the fits landed at the outer limit of the existing range. We also verified that the results did not change when we ran the step-2 fits from the 10 best grid points rather than the one best gridpoint with 10 random seeds.

Integration Models

We used a simple version of the well-known drift-diffusion model to implement temporal integration (Ratcliff & McKoon, 2008). The samples of momentary evidence were integrated perfectly over time into a decision variable, DV(t), given by:

$$DV(t) = DV(t-1) + d \times dt + N(0, s^2 dt)$$

where d(t) is the drift rate at time t that was non-zero only for the duration of evidence, dt is the discrete time increment (set to 2 ms), and the last term refers to Gaussian noise with mean 0 and standard deviation of s×dt. Decisions were terminated by reaching the upper/lower bound, or if no bound was reached, determined by the sign of the DV at stimulus offset (Kiani et al 2008).

Three variants of this model were fitted, the first with only one drift rate applied equally to both high and low contrast-differences to provide a worst-case benchmark; the second with two separate drift rates (UIntg) and no decision bound so that choices were always determined at stimulus offset by the sign of the DV; and the third with a bound parameter included alongside the two drift rates (BIntg).

Extrema Detection models

Rather than accumulate evidence samples over time, extrema detection models compared each independent sample of momentary evidence to a criterion or Bound level, terminating upon the first sample to exceed it. The decision variable was thus given by

$$DV(t) = d \times dt + N(0, s^2 dt)$$

Three variants of this Extrema Detection Model were tested, differing only in how choice was determined when no bound was exceeded within the evidence period: i) by a random (50-50) Guess (labelled 'ExtG' for extrema detection with guess default; see Stine et al., 2020), ii) based on the sign of the last evidence sample ('ExtL'), or iii) based on the sign of the most extreme (sub-bound) sample encountered throughout the evidence period ('ExtE'). This third variant, labelled 'Extremum-tracking,' entailed a memory component to update the extrema as and when they were encountered throughout the trial. Although memory is sometimes associated with integration in terms of potentially shared mechanisms (Wong & Wang, 2006), we class this variant as a non-integration model because the one remembered item (the most recent extremum) is updated sparsely and replaces rather than builds on the previous value, unlike integration which entails continuous, iterative updates. Like the Integration model, we tested three versions of this model: one with one drift rate and no bound, one with two drift rates and no bound, and one with two drift rates and a bound. Unbounded versions of the other two Extrema detection models were not implemented because, in the absence of a bound, they would trivially default to a random guess on all trials, or a decision based on just the last sample, which cannot provide accuracy improvements across the first 3 durations.

SnapShot with Guess

In this model, a single, randomly-selected sample of momentary evidence was used to determine the decision. In unbounded variants, the sign of this sample determined the choice regardless of its magnitude. In bounded variants, the sample determined the choice only if it exceeded the bound and a

random guess was produced instead when the sample fell short of the bound. Again, three versions of this model were run, respectively with one drift rate and no bound, two drift rates and no bound, and two drift rates with a bound.

Neurally-constrained Models

The models above were fitted to behaviour alone, and several of them captured accuracy equally well. To resolve this ambiguity, three of these successful models for which a means of producing buildup dynamics in the average decision variable could be conceived - the Integration model and two Extrema Detection models - were subsequently fitted jointly to behaviour and to the average CPP. To do this, a model-based simulation of the CPP was needed to compare to the empirical CPP. This simulation was derived in different ways in each model.

In the Integration model, the CPP was simulated as the absolute value of differential cumulative evidence as previously described (Afacan-Seref et al., 2018). The use of the absolute value operator is based on the assumption that the CPP reflects the sum of two neural populations (one for each choice alternative) that activate in a mutually exclusive way to represent positive and negative values of the single, differential decision variable (Kelly et al., 2021). This is consistent with observations that the CPP builds with positive polarity regardless of choice and that the CPP reaches a stereotyped threshold level at response during urgency-free, continuous monitoring decisions, even though it derives from two accumulator processes only one of which ultimately triggers the decision (Kelly & O'Connell, 2013). Based on post-decision CPP dynamics observed in immediate response tasks and easy delayed-response conditions, we assumed that the decision variable falls back to baseline in a stereotyped way upon reaching the bound. For consistency with the non-integration models below, we modelled this fall-back as a quarter-sine wave (from peak to zero-crossing).

In the Extremum-tracking model, the CPP was simulated again as the absolute value of a differential DV, with the DV in this case being an iteratively-updated representation of the most extreme evidence sample encountered so far. As in the integration model, a quarter-sine fall-down from the bound was assumed if and when it was crossed. In the final model based on Extrema detection with a random Guess default, the decision variable is again the momentary evidence itself, subjected to bounds, but here it is assumed that a centroparietal signal is only evoked upon such bound crossings. We assumed that the trial-averaged CPP was the result of averaging single-trial discrete event waveforms that were triggered by, and therefore time-locked to, decision termination. We modelled this decision termination-flagging event as a half (upper) sinusoidal function whose width is a free parameter and whose amplitude is fixed, starting at the bound crossing time. When averaged, these half-sines can resemble any arbitrary waveform depending on the distribution of bound-crossing times. In theory, this formulation allows the Extrema Guess model to produce a signal with ramp-like properties similar to the CPP without accumulating evidence. For trials where no bound was reached, the CPP was simulated as zero throughout the evidence period. For consistency with this last, 'Extremum-flagging' model, the fall-down dynamics of the simulated CPP after each bound-crossing for the Integration and Extremum-tracking models was assumed to take the shape of a quarter-sine function (peak to zero-crossing), whose width was similarly set as a free parameter in the neurally-constrained fits. Note that the 'FlagWidth' parameter referred to half of the sinusoidal period across all 3 models for consistency, and so the actual duration of the peak-to-baseline fall-down in the Integration and Extremum-tracking models was half of the estimated value of this parameter. As in the Integration mode, the absolute value of each single-trial DV waveform was taken before averaging in both Extrema detection models.

Based on the above simulations of the CPP, neurally constrained models were implemented by combining the same G^2 metric quantifying behavioural fit with a metric of CPP fit with variable weighting in a joint objective function to be optimized by the same SIMPLEX algorithm as used previously. Specifically, the joint objective function was given by

$$G_p^2 = G^2 + w((R_H)^2 + (R_L)^2)$$

$$R^2 = \frac{RSS}{TSS} = \frac{\sum_i^n (y_i - x_i)^2}{\sum_i^n (x_i - \bar{x})^2}$$

where R^2 is the residuals sum of squares (RSS) over the total sum of squares (TSS) of the differences between simulated, y_i, and empirical CPP, x_i, waveforms, computed for the high (subscript H) and low (L) contrast conditions. The 4 low-contrast conditions were averaged so that the models would not be informed about amplitude variations across them, allowing these variations to be used as an independent evaluation of model fit. The neural weighting was set to each of the values w=0, 0.1, 1, 10, 100, 1000 so that we could conduct a sensitivity analysis of how model fits varied with the degree of emphasis placed on the neural fit. To compute the neural R^2 term, we had to first apply appropriate time-windowing to exclude segments of the CPP waveform that did not clearly map onto the decision variable in the model, and normalize the real and simulated CPP so that they would have the same units. Specifically, we extracted windows of the observed CPP waveforms starting from 200 ms, the point at which the low and high-contrast CPP traces begin to diverge and presumably therefore the moment of initial impact of evidence, up to 900 ms in the high-contrast condition and up to 1600 ms in the low contrast conditions. The earlier end time for the high-contrast condition was chosen to avoid the bulk of a separate negative-going potential that appeared in the latter half of the stimulus-evaluation period in that condition, whose relationship to the decision process was not obvious and which would therefore risk biasing the measurement of the CPP during that time window. We then shifted the waveforms down so that they began at zero amplitude and then divided by the mean across all sample points to normalize them. The same normalization procedure was applied to the model-simulated CPPs, which gave the simulated and empirical waveforms comparable units.

In extracting the to-be-modelled CPP waveform starting at the onset of evidence-driven buildup, we remove the need to estimate transmission or other non-decision delays, which had also been unnecessary in behaviour-only modelling due to not fitting reaction time. Without loss of generality, we could therefore assume for simplicity that evidence accumulation and extremum-detection processes began without delay at evidence onset, and that the 'flag' signal began immediately upon bound crossing. It is important to note that this measure to remove unnecessary timing parameter estimation still allows for decision variable onset delays and flag-onset delays that may occur in reality. In all three models, we allowed the width of the sine-function describing the post-bound fall-down (quarter-sine) and flag (half-sine) to vary as a free parameter, on the basis that we cannot estimate these single-trial quantities from observed average waveforms without knowing bound-crossing times.

Although motor preparation signals were not used in the model fits, it was additionally of interest to simulate these signals from all models for qualitative comparison with Mu/Beta lateralisation (Figure 3C). To do this, we assumed that the relative motor preparation directly reflected the differential decision variable but with the distinction that it sustained at the level of the bound upon reaching it, rather than immediately falling like the CPP. From this perspective of motor preparation, the Extremum-flagging model can be considered highly similar in nature to the discrete stepping model proposed by Latimer et al., (2015), which features stochastically occurring, discrete steps in neural activity reflecting transitions from an undecided to a decided state.

Beyond the core parameters of drift rate and bound, sequential sampling models are typically endowed with additional parameters to enable them to achieve a comprehensive behavioural fit (Ratcliff et al., 2016). We therefore explored a range of additional parameters and their impact on the neurally-constrained fits. Specifically, for each of the three model classes - Integration, Extremum-tracking and Extremum-flagging - we included one extra model feature at a time, alongside the two drift rates and bound of the basic model. We describe each additional feature in turn:

Sampling onset time (SampT). When sensory evidence is weak, it is difficult to distinguish evidence from noise and this can lead to variability in the onset time of sequential sampling (Devine et al., 2019). Though subtle, the CPP in this task appeared to begin to rise before the onset of evidence, suggesting a potentially early sampling onset (Figure 3B). SampT was a free parameter that was allowed to take a value before or after time zero, which in the models marked the time when informative contrast-difference evidence first became available for sampling. Figure 4 - Figure Supplement 1 shows that this extra parameter did not substantively improve the fit of either of the competitive models.

Starting point variability (sz) Variability in the starting point of the decision variable allows for random trial-to-trial fluctuations in bias (Bogacz et al., 2006; Purcell et al., 2012). We implemented starting point

variability as a uniform distribution of starting levels centred at the neutral position between the correct/error bounds whose standard deviation was the free parameter sz. Figure 4 - Figure Supplement 1 illustrates that adding this extra parameter substantially improved the fit of the integration model. Interestingly, this nonzero starting-point variability is supported by the motor signal activity at the stimulus onset (Figure 5), although the model was not informed about this empirical effect.

Time-varying bound (,$t_{0.5}$) While constant bounds offer a parsimonious account of the decision criterion, bounds that change the criterion amount of evidence required for a decision as a function of elapsed time in the trial have been shown to better capture neural and behavioural data in certain cases, such as when urgency to respond is imposed through deadlines (Kelly et al., 2021), and although responses were deferred until after the long stimulus in the current task, a tight deadline of 0.5 sec was imposed at that point, which in combination with the predominant very weak evidence may have led to an experience of urgency. Such dynamically changing bounds have been shown to be optimal in certain circumstances involving unpredictable task difficulty (Malhotra et al., 2018) and are a critical feature of some extrema detection models (Thura et al., 2012). To determine whether models with time-varying bounds provide better fits to the data than models with fixed bounds, we implemented a time-varying bound using a linear function, given by

$$b(t) \;=\; \pm b_0(1-\kappa t)$$

where b_0 is the initial boundary separation, time (t), represents the collapse rate, allowed to be positive or negative (increasing bound). We also fit an alternative, nonlinear collapsing bound model with a hyperbolic ratio function, given by:

$$b(t) \;=\; \pm b_0(1-\kappa\frac{t}{t\;+\;t_{0.5}})$$

where $t_{0.5}$ the semi-saturation time. The positive and negative versions of this bound function correspond to the correct and incorrect alternative. This temporal function has been previously used to quantify the hypothetical urgency signal in physiological data (Churchland et al., 2008; Hanks et al., 2011; Voskuilen et al., 2016), except that the amount of collapse is expressed as a fraction of the total boundary separation rather than as an absolute value, which simplified fitting and led to more interpretable parameter estimates. Our results indicate that incorporating time-varying bounds improved the model fits for all the models (Figure 4 - Figure Supplement 1). The optimal bound change for the two extrema-detection models was always a decreasing one (positive values of ; Table S3&S4).

Between trial drift rate variability (eta) Although our stimulus does not include fluctuations in contrast, the internal representation of this contrast may fluctuate from trial to trial, with factors such as alertness and attention (Ratcliff et al., 2016). To allow for such between-trial variability in encoding sensory information, an extra parameter for trial-to-trial variability in drift rate was added, captured as Gaussian variability with standard deviation eta, centred on the drift rate parameters. This additional parameter improved the fit for integration model slightly (lower G^2+Penalty) while did not improve the fit for Extremum-flagging model (Tables S3, S5-8).

Accumulation duration (Dur) A central question in this study was whether judgements about prolonged static sensory evidence involve sequential sampling over the full evidence period, or only a portion of it. To explore this possibility in all three models, we implemented the models with one extra parameter ('Dur') that determined the time window starting from the onset of evidence encoding during which evidence was sequentially sampled for decision formation. The results revealed that the basic Integration model did improve with the addition of this duration parameter, but the estimate of this duration needed to fit the neural dynamics was less than 800 ms, therefore precluding reproduction of the significant increase in behavioural accuracy from the 800-ms to 1600-ms duration condition (Tables S3, S5-8).

Leaky accumulation. Another way that integration of evidence can be limited is through leak, whereby older samples are discounted in the cumulative total, and the decision variable gradually decays to zero in the absence of sufficient evidence. We implemented leaky accumulation through an additional leak

parameter which scales-down the weighting of previous evidence relative to the current sample, according to:

$$DV(t) \ = \ (1 - leak) \times DV(t-1) \ + \ d(t) \times dt \ + N(0, s^2 dt)$$

The results demonstrated minimal information loss over time (leak = 0.0002 at w=10), with only a small benefit to fit from including the parameter, in agreement with previous findings of low leakage (Kiani et al., 2008).

The fitting procedure for the neurally-constrained models followed the same general strategy as the model fits to behaviour alone, starting with evaluation of objective function at each point of a comprehensive grid, but now running the SIMPLEX algorithm from the 100 best grid points as starting parameter vectors. In the most complex model, which extends the basic model by including starting point variability and a non-linear collapsing decision bound, five grid points were used for each of three parameters (starting point variability, collapse rate and semi-saturation time) instead of 10 to limit computational expense (see Table S4). Running this procedure on all three model classes, we found that whereas the best objective function values for the Integration and Extremum-flagging models were 7.98 and 9.95, respectively (taking w=10 as representative), the best achieved by Extremum-tracking was 19.76. We therefore set Extremum-tracking aside and proceeded to refine the two competitive models, the best performing of which included both collapsing bound and starting point variability. To further ensure convergence on global optima despite the additional complexity, we ran additional fits with starting vectors generated from the parameter estimates of models that had fewer free parameters. Specifically, for the 'szC' models (with free parameters for both collapsing bound and starting point variability), we started a fit from the best fitting 'C' model and with sz set to zero, half of and equal to its estimated value in the 'sz' model. Analogously, we generated two more 'szC' starting vectors by taking the best fitting 'sz' model parameter values and adding an additional collapse rate equal to zero, equal to half and equal to the collapse rate found in the 'C' model.

References

Afacan-Seref, K., Steinemann, N. A., Blangero, A., & Kelly, S. P. (2018). Dynamic interplay of value and sensory information in high-speed decision making. *Current Biology*, *28*(5), 795–802.

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.

Brainard, D. H., & Vision, S. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433–436.

Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, *11*(6), 693–702.

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, *29*(37), 11560–11571.

Corbett, E. A., Martinez-Rodriguez, L. A., Judd, C., O'Connell, R. G., & Kelly, S. P. (2023). Multiphasic value biases in fast-paced decisions. *Elife*, *12*, e67711.

de Lange, F. P., Rahnev, D. A., Donner, T. H., & Lau, H. (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *Journal of Neuroscience*, *33*(4), 1400–1410.

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21.

Devine, C. A., Gaffney, C., Loughnane, G. M., Kelly, S. P., & O'Connell, R. G. (2019). The role of premature evidence accumulation in making difficult perceptual decisions under temporal uncertainty. *Elife*, *8*, e48526.

Donchin, E., & Coles, M. G. (1988). Is the P300 component a manifestation of context updating?. *Behavioral and brain sciences*, *11*(3), 357-374.

Donner, T. H., Siegel, M., Fries, P., & Engel, A. K. (2009). Buildup of choice-predictive activity in human motor cortex during perceptual decision making. *Current Biology*, *19*(18), 1581–1585.

Evans, N. J., & Wagenmakers, E.-J. (2020). Evidence accumulation models: Current limitations and future directions. *The Quantitative Methods for Psychology*, *16*(2), 73–90.

Feuerriegel, D., Jiwa, M., Turner, W. F., Andrejević, M., Hester, R., & Bode, S. (2021). Tracking dynamic adjustments to decision making and performance monitoring processes in conflict tasks. *Neuroimage*, *238*, 118265.

Foxe, J. J., & Snyder, A. C. (2011). The role of alpha-band brain oscillations as a sensory suppression mechanism during selective attention. *Frontiers in Psychology*, *2*, 154.

Geuzebroek, A. C., Craddock, H., O'Connell, R. G., & Kelly, S. P. (2023). Balancing true and false detection of intermittent sensory targets by adjusting the inputs to the evidence accumulation process. *ELife*, *12*, e83025.

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. (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience*, *30*(1), 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038
Grogan, J. P., Rys, W., Kelly, S. P., & O'Connell, R. G. (2023). Confidence is predicted by pre-and post-choice decision signal dynamics. *Imaging Neuroscience*,*1*, imag-1.
Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience*, *31*(17), 6339–6352.
Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife*, *3*, e02260.
Hanks, T. D., & Summerfield, C. (2017). Perceptual decision making in rodents, monkeys, and humans. *Neuron*, *93*(1), 15–31.
Hyafil, A., de la Rocha, J., Pericas, C., Katz, L. N., Huk, A. C., & Pillow, J. W. (2023). Temporal integration is a robust feature of perceptual decisions. *Elife*, *12*, e84045.
Kayser, J., & Tenke, C. E. (2006). Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: I. Evaluation with auditory oddball tasks. *Clinical Neurophysiology*, *117*(2), 348–368.
Kelly, S. P., Corbett, E. A., & O'Connell, R. G. (2021). Neurocomputational mechanisms of prior-informed perceptual decision-making in humans. *Nature Human Behaviour*, *5*(4), 467–481.
Kelly, S. P., & O'Connell, R. G. (2013). Internal and external influences on the rate of sensory evidence accumulation in the human brain. *Journal of Neuroscience*, *33*(50), 19434–19441.
Kelly, S. P., & O'Connell, R. G. (2015). The neural processes underlying perceptual decision making in humans: Recent progress and future directions. *Journal of Physiology-Paris*, *109*(1–3), 27–37.
Khodadadi, A., & Townsend, J. T. (2015). On mimicry among sequential sampling models. *Journal of Mathematical Psychology*, *68*, 37–48.
Kiani, R., Hanks, T. D., & Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *Journal of Neuroscience*, *28*(12), 3017–3029.
Kleiner, M., Brainard, D., & Pelli, D. (2007). *What's new in Psychtoolbox-3?* https://pure.mpg.de/rest/items/item_1790332/component/file_3136265/content
Latimer, K. W., Yates, J. L., Meister, M. L. R., 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. https://doi.org/10.1126/science.aaa4056
Ludwig, C. J., Gilchrist, I. D., McSorley, E., & Baddeley, R. J. (2005). The temporal impulse response underlying saccadic decisions. *Journal of Neuroscience*, *25*(43), 9907–9912.
Malhotra, G., Leslie, D. S., Ludwig, C. J. H., & Bogacz, R. (2018). Time-varying decision boundaries: Insights from optimality analysis. *Psychonomic Bulletin & Review*, *25*(3), 971–996. https://doi.org/10.3758/s13423-017-1340-6
McCone, H., Devine, C. A., McNickle, E., Dully, J., Geuzebroek, A. C., McGovern, D. P., ... & O'Connell, R. G. (2026). A Movement-Independent Signature of Urgency During Human Perceptual Decision Making. *Journal of Neuroscience*.
Murphy, P. R., Robertson, I. H., Harty, S., & O'Connell, R. G. (2015). Neural evidence accumulation persists after choice to inform metacognitive judgments. *elife*, *4*, e11946.
Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus–norepinephrine system. *Psychological Bulletin*, *131*(4), 510.
O'Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735.
O'Connell, R. G., & Kelly, S. P. (2021). Neurophysiology of Human Perceptual Decision-Making. *Annual Review of Neuroscience*, *44*(1), 495–516. https://doi.org/10.1146/annurev-neuro-092019-100200
O'Connell, R. G., Shadlen, M. N., Wong-Lin, K., & Kelly, S. P. (2018). Bridging neural and computational viewpoints on perceptual decision-making. *Trends in Neurosciences*, *41*(11), 838–852.
Okazawa, G., & Kiani, R. (2023). Neural Mechanisms That Make Perceptual Decisions Flexible. *Annual Review of Physiology*, *85*(1), 191–215. https://doi.org/10.1146/annurev-physiol-031722-024731
Parés-Pujolràs, E., Kelly, S. P., & Murphy, P. R. (2025). Dissociable encoding of evolving beliefs and momentary belief updates in distinct neural decision signals. *Nature Communications*, *16*(1), 3922.
Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.
Pfurtscheller, G., & Da Silva, F. L. (1999). Event-related EEG/MEG synchronization and desynchronization: Basic principles. *Clinical Neurophysiology*, *110*(11), 1842–1857.
Philiastides, M. G., Heekeren, H. R., & Sajda, P. (2014). Human scalp potentials reflect a mixture of decision-related signals during perceptual choices. *Journal of Neuroscience*, *34*(50), 16877–16889.
Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, *118*(10), 2128–2148.
Purcell, B. A., & Palmeri, T. J. (2017). Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*, *76*, 156–171.
Purcell, B. A., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2012). From salience to saccades: Multiple-alternative gated stochastic accumulator model of visual search. *Journal of Neuroscience*, *32*(10), 3433–3446.
Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*(2), 333.
Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, *20*(4), 260–281.
Regan, D. (1989). Human brain electrophysiology. *Evoked Potentials and Evoked Magnetic Fields in Science and Medicine*. https://cir.nii.ac.jp/crid/1572543025661076864
Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience*, *22*(21), 9475-9489.
Shushruth, S., Zylberberg, A., & Shadlen, M. N. (2022). Sequential sampling from memory underlies action selection during abstract decision-making. *Current Biology*, *32*(9), 1949–1960.
Steinemann, N. A., O'Connell, R. G., & Kelly, S. P. (2018). Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nature Communications*, *9*(1), 3627.
Steinemann, N., Stine, G. M., Trautmann, E., Zylberberg, A., Wolpert, D. M., & Shadlen, M. N. (2024). Direct observation of the neural computations underlying a single decision. *Elife*, *12*, RP90859.
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.
Thura, D., Beauregard-Racine, J., Fradet, C.-W., & Cisek, P. (2012). Decision making by urgency gating: Theory and experimental support. *Journal of Neurophysiology*, *108*(11), 2912–2930. https://doi.org/10.1152/jn.01071.2011
Trueblood, J. S., Heathcote, A., Evans, N. J., & Holmes, W. R. (2021). Urgency, leakage, and the relative nature of information processing in decision-making. *Psychological Review*, *128*(1), 160.
Twomey, D. M., Kelly, S. P., & O'Connell, R. G. (2016). Abstract and effector-selective decision signals exhibit qualitatively distinct dynamics before delayed perceptual reports. *Journal of Neuroscience*, *36*(28), 7346–7352.
Twomey, D. M., Murphy, P. R., Kelly, S. P., & O'Connell, R. G. (2015). The classic P300 encodes a build-to-threshold decision variable. *European Journal of Neuroscience*, *42*(1), 1636–1643. https://doi.org/10.1111/ejn.12936
Uchida, N., Kepecs, A., & Mainen, Z. F. (2006). Seeing at a glance, smelling in a whiff: Rapid forms of perceptual decision making. *Nature Reviews Neuroscience*, *7*(6), 485–491.
Verleger, R. (2020). Effects of relevance and response frequency on P3b amplitudes: Review of findings and comparison of hypotheses about the process reflected by P3b. *Psychophysiology*, *57*(7), e13542. https://doi.org/10.1111/psyp.13542
Voskuilen, C., Ratcliff, R., & Smith, P. L. (2016). Comparing fixed and collapsing boundary versions of the diffusion model. *Journal of Mathematical Psychology*, *73*, 59–79.
Waskom, M. L., & Kiani, R. (2018). Decision making through integration of sensory evidence at prolonged timescales. *Current Biology*, *28*(23), 3850–3856.
Watson, A. B. (1979). Probability summation over time. *Vision Research*, *19*(5), 515–522.
Weindel, G., Borst, J. P., & Maanen, L. van. (2025). Decision-making components and times revealed by the single-trial electro-encephalogram. *eLife*, *14*. https://doi.org/10.7554/eLife.108049.1
Widmann, A., Schröger, E., & Maess, B. (2015). Digital filter design for electrophysiological data–a practical approach. *Journal of Neuroscience Methods*, *250*, 34–46.
Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience*, *26*(4), 1314–1328.