nature human behaviour

ARTICLES

https://doi.org/10.1038/s41562-020-00967-9

Neurocomputational mechanisms of prior-informed perceptual decision-making in humans

Simon P. Kelly[1,2,3] ✉, Elaine A. Corbett[2,3] and Redmond G. O'Connell[3]

To interact successfully with diverse sensory environments, we must adapt our decision processes to account for time constraints and prior probabilities. The full set of decision-process parameters that undergo such flexible adaptation has proven to be difficult to establish using simplified models that are based on behaviour alone. Here, we utilize well-characterized human neurophysiological signatures of decision formation to construct and constrain a build-to-threshold decision model with multiple build-up (evidence accumulation and urgency) and delay components (pre- and post-decisional). The model indicates that all of these components were adapted in distinct ways and, in several instances, fundamentally differ from the conclusions of conventional diffusion modelling. The neurally informed model outcomes were corroborated by independent neural decision signal observations that were not used in the model's construction. These findings highlight the breadth of decision-process parameters that are amenable to strategic adjustment and the value in leveraging neurophysiological measurements to quantify these adjustments.

The ability to adapt our decision processes to varying task demands is a critical facet of higher cognitive function. Most famously, we can expedite our decisions when we are under time pressure, and we can prioritize choice alternatives that have a higher prior probability of being correct. Computational decision models, in which sensory evidence is accumulated over time until a bound is reached and triggers an action[1,2], have provided a powerful framework for examining the mechanisms that underlie such adaptations[3–5], yet our understanding remains very incomplete. As these models have traditionally been based solely on behaviour, they have been developed with judiciously reduced complexity to avoid overfitting and model mimicry. This emphasis on parsimony extended to favouring models that can account for any given task manipulation by adjusting the fewest possible parameters[6]. Whether this approach captures the full range of computational adjustments enacted by the brain remains uncertain. Furthermore, there now exist several alternative model variants that disagree regarding core elements of model structure, such as whether bounds are constant or time-dependent[7]. Resolving these disagreements has proven to be very difficult using behavioural modelling alone[8,9], yet is essential because different structures can lead to fundamentally different conclusions about the nature and extent of the key adaptations[10–13].

Neurophysiological signals that exhibit dynamics consistent with bounded evidence accumulation have now been observed in multiple species, especially in activity associated with planning decision-reporting actions[14–20]. These neural signatures can bolster model-based inferences in several complementary ways. They provide an empirical means to test key model predictions concerning decision signal dynamics. More fundamentally, they can be used to guide the construction of the model—that is, what process components and parameters to include—and also to quantitatively constrain certain model parameters that correspond directly to analogous signal parameters[21,22]. Such constraints, in principle, enable cognitive models to take on greater levels of complexity and to capture a wider range of effects without necessarily increasing the number of free parameters[10]. These advantages have not yet been fully utilized to examine human perceptual decisions. Here we used recently characterized human electroencephalographic (EEG) signatures of decision formation to construct, constrain and validate a bounded accumulation model that establishes the computational adjustments by which humans adapt to time constraints, stimulus strength and prior probability.

Guided by decision signal observations, we constructed a neurally informed (NI) model that differs from conventional models such as the drift diffusion model (DDM) in two fundamental ways. First, whereas the DDM assumes only one source of build-up—namely, sensory evidence that is accumulated at a mean 'drift rate' that scales with evidence strength—the NI model includes an early-onsetting, evidence-independent 'urgency' signal that adds to cumulative evidence[22–25] in driving motor preparation towards an action-triggering threshold. Second, whereas the DDM estimates only a single non-decision time that encapsulates all residual time that is not spent on bounded evidence accumulation, the NI model isolates three distinct temporal components: evidence-encoding onset, accumulation onset[26] and post-decision motor-execution time.

Our results revealed multiple distinct computational adaptations among these build-up and time-delay components in addition to the shifts in decision criterion (bound/starting level) that are typically observed[3,5,27]. Moreover, certain adaptation effects were fundamentally at odds with those indicated by the standard DDM. Under speed pressure, the DDM indicated a substantially shortened non-decision time and reduced drift rate similar to other recent studies[6,11,28–33]; by contrast, the NI model indicated only very small changes in non-decision time and an enhancement of drift rate that acts to partially recoup the accuracy losses incurred by lower bounds. Dynamic urgency, the role of which has been controversial in the behavioural modelling literature[7], was strongly observed in all of the regimes and was strategically adjusted to increase the

[1]Department of Biomedical Engineering, City College of the City University of New York, New York, NY, USA. [2]School of Electrical and Electronic Engineering and UCD Centre for Biomedical Engineering, University College Dublin, Dublin, Ireland. [3]Trinity College Institute of Neuroscience and School of Psychology, Trinity College Dublin, Dublin, Ireland. ✉e-mail: simon.kelly@ucd.ie

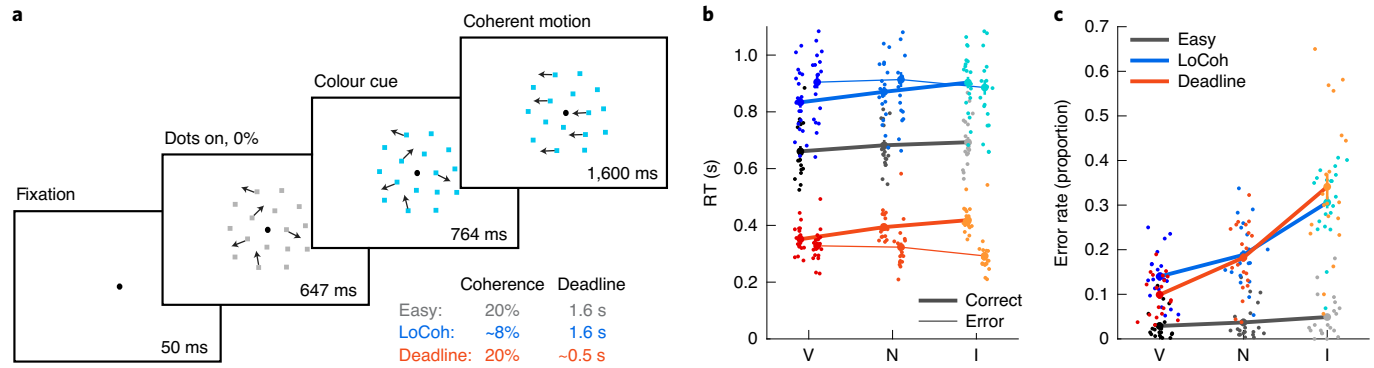

Fig. 1 | Prior-cued motion discrimination task and behavioural results. a, The trial structure, with fixed delay and cue durations across all trials. The coherence and deadline settings for the Easy, weak-evidence (LoCoh) and time-pressured (Deadline) regimes are listed underneath. In the LoCoh regime, coherence was individually set to achieve 80% accuracy (mean ± s.d., 8.4 ± 1.7%). In the Deadline regime, a shorter deadline was individually set (485 ± 61 ms) to be close to the median RT of the Easy regime to induce strong speed pressure. **b**, Mean RT significantly varied across valid (V), neutral (N) and invalid (I) trials ($F_{2,38} = 41.8$, $P < 0.001$, partial $\eta^2 = 0.69$) and across regime ($F_{2,38} = 218$, $P < 0.001$, partial $\eta^2 = 0.92$), with a significant interaction between the two ($F_{4,76} = 3.23$, $P = 0.040$, partial $\eta^2 = 0.14$; valid–invalid difference bigger for the LoCoh and Deadline regimes relative to the easy regime; all pairwise comparisons are provided in Supplementary Table 3). **c**, Error rate also significantly varied with validity ($F_{2,38} = 70.6$, $P < 0.001$, partial $\eta^2 = 0.79$) and across regime ($F_{2,38} = 95.5$, $P < 0.001$, partial $\eta^2 = 0.83$; error rates were higher in the LoCoh and Deadline regimes compared with the Easy regime; all pairwise comparisons are provided in Supplementary Table 3), with a significant interaction between the two ($F_{4,76} = 27.2$, $P < 0.001$, partial $\eta^2 = 0.59$; valid–invalid difference was bigger for the LoCoh and Deadline regimes compared with the Easy regime). For **b** and **c**, data are mean ± s.e.m. after between-participant variance was factored out. Individual participants are shown as background dots to convey inter-individual variability in performance.

likelihood of reaching a bound within the imposed deadlines. Finally, the NI model indicated a role for prior-probability-based biases in accumulation rate (drift rate bias), which the DDM failed to identify through standard model comparison. Critically, neural decision signal features that were not incorporated into the construction of the NI model exhibited modulations that corroborated the parameter differences and matched model-simulated decision formation dynamics.

Results

Scalp EEG signals were recorded while twenty human participants performed a prior-cued motion discrimination task under three different blocked regimes that involved low demands (high coherence and long deadline ('Easy')), weak evidence (low coherence and long deadline ('LoCoh')) and speed pressure (high coherence and short deadline ('Deadline')), respectively. In each trial, initially grey, incoherently moving dots changed colour to indicate whether the direction of upcoming motion would more likely be to the left, to the right or equally likely (Fig. 1a), and, after a fixed predictable delay, suddenly transitioned to coherent motion. The participants were asked to respond with their left hand for leftward motion or with their right hand for rightward motion, and they earned points for reporting correct decisions within the deadline.

Behaviour was strongly affected by both regime and prior probability in the expected ways (Fig. 1b,c). Relative to the Easy regime, error rates were higher in both the Deadline and LoCoh regimes, and response times (RTs) were much faster in the Deadline regime and slower in the LoCoh regime. Validly cued trials (in which motion direction was congruent with cue colour) had greater accuracy, faster correct responses and slower error responses compared with invalid trials, and these biases were stronger in both the LoCoh and Deadline regimes compared with the Easy regime. In the remainder of the Results section, we first lay out the NI model construction and the empirical neural effects on which it was based; we then describe the outcomes of the model fit to behaviour, contrasting with those of a standard DDM; and, finally, we describe follow-up neural and model-based analyses to confirm each of the effects indicated by the NI model.

Neurally informed model construction. To determine the decision-process adaptations that gave rise to these behavioural patterns, we leveraged EEG signatures of motor preparation and execution to inform the construction of an accumulation-to-bound model. The model had the usual basic ingredients of a bound, starting-point bias, drift rate (mean of the evidence being accumulated) and non-decision time, but with key, NI elaborations, including dynamic urgency and distinct pre- and post-decision delays. To enable adjustments of this wider range of parameters to be tested while keeping the number of free parameters as low as in standard models, we quantitatively constrained the parameters of bound, starting points and post-decision execution delays to match their corresponding neural signal characteristics (motor threshold, baseline motor preparation and motor potential onset latency) as described below.

Motor-level decision threshold. We first sought to confirm our previous findings[19,25] that motor preparation signals reach a constant level at response, consistent with a fixed action-triggering threshold and similar to observations in sensorimotor regions of the monkey brain[14–16]. We traced the motor preparation signals representing left and right responses through Mu–Beta-band spectral amplitude (MB signal, integrated over 8–30 Hz)[17,34] over the right- and left-hemisphere motor areas, respectively. Plotting the MB amplitude just before response across all of the conditions and for several RT bins revealed that the motor preparation signal contralateral to the button pressed reached a stereotyped threshold level regardless of the timing of the decision, the perceptual difficulty, the speed pressure, the correctness of the response and the prior information available (Fig. 2b; the ipsilateral signal amplitudes are shown in Supplementary Fig. 1). Accordingly, we set a constant bound on the NI model and, to enable a straightforward linkage to the empirical motor preparation signals, we cast the decision process as a race between two parallel build-up signals to reach that bound[16] (Fig. 2d).

Motor-level starting points. We next examined the neural dynamics of anticipatory motor preparation before evidence onset. Replicating

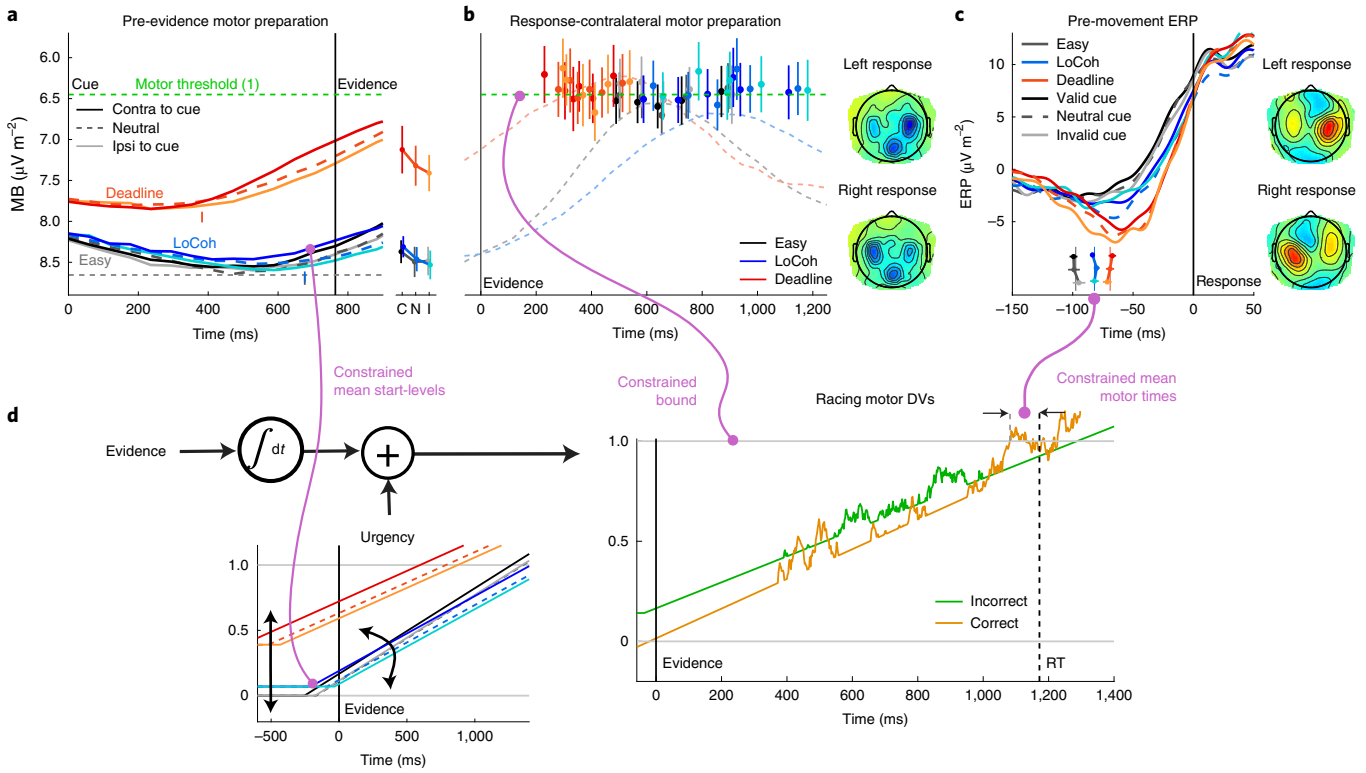

Fig. 2 | Neurally informed (NI) model constrained by motor-related neural signal amplitudes and latencies. a, Timecourse of motor preparation (Mu-beta (MB) amplitude reduction) after cue onset. The ticks beneath the waveforms indicate the first time point at which the temporal slope (collapsed across cue validity) reaches a significantly positive value (382 ms post-cue for the deadline regime and 676 ms for the Easy and LoCoh regimes). Given that MB is measured in windows of 294 ms centred at each timepoint, these build-up onsets are too early to be driven by evidence encoding. Note that the *y* axis is flipped such that increasing motor preparation (decreasing MB) is upwards. Amplitudes at the pre-evidence time point (−59 ms) used for model constraints are shown on the right. **b**, Mean pre-response MB amplitude contralateral to the response, for each prior condition and regime, with the correct responses divided into five equally sized RT bins. The mean amplitude across conditions represents the 'motor threshold', which is indicated by a green dashed line in **a** and **b**. As a reference, the average MB waveforms for neutrally cued, correct trials are plotted as dashed traces. **c**, Contralateral, response-locked motor potentials. The inflection point was taken to indicate the time at which motor preparation gives way to motor execution processes, and was estimated as the timepoint at which the first-derivative signal (Supplementary Fig. 2) passes through zero. These times are indicated below the waveforms for all nine conditions. Topographies show the amplitude at 0 ms relative to −20 ms to isolate the sharp, response-locked positive-going deflection from any concomitant slower potentials. **d**, In the NI model, two racing decision variables (DV) corresponding to left/right motor preparation (MB) are driven by dynamic urgency summed with cumulative differential evidence. Mean starting levels were quantitatively constrained by pre-evidence MB amplitudes in **a**, and post-decision motor execution times by the pre-response execution potential latencies in **c**. A single example trial from the LoCoh regime is shown, in which, due to random trial-to-trial variations in starting level and urgency build-up rate, the correct alternative happens to start lower yet ultimately wins the race due to cumulative evidence. Vertical axes in both graphs are in arbitrary units. Supplementary Figure 3 explicitly lays out all of the assumed mathematical relationships and offers a plausible systems-level implementation. For **a**–**c**, data are mean ± s.e.m. after factoring out between-participant variance.

previous studies[4,27,34], pre-evidence motor preparation levels were relatively increased under speed pressure (main effect of regime in 3×3 repeated-measures analysis of variance (RMANOVA) on MB amplitude, $F_{2,38}=8.08$, $P=0.0064$, partial $\eta^2=0.26$) and for the more probable alternative (main effect of prior cue, $F_{2,38}=6.53$, $P=0.0037$, partial $\eta^2=0.30$), reflecting the dominant decision criterion adjustments that are long-established in standard models (Fig. 2a). Interestingly, in all of the regimes, motor preparation starts to build before sensory evidence could possibly yet have an impact on it, especially in the Deadline regime. Such evidence-independent build-up is the hallmark of dynamic urgency-signal components, which drive the decision process toward its threshold even in the absence of evidence[22,23]. Mean pre-evidence motor preparation amplitudes were used to directly constrain the corresponding mean starting-point parameter values of the NI model, and an unconstrained mean urgency rate parameter was included for each regime to allow for a linear increase in urgency from those pre-evidence levels (Fig. 2d).

Post-decision motor time. We next estimated post-decision delays—the 'motor time' component of non-decision time intervening between bound-crossing and response completion—by measuring the onset of motor execution potentials relative to RT[19,25,35] (Fig. 2c; Methods). We found that this onset significantly differed as a function of regime ($F_{2,38}=20.56$, $P<0.001$, partial $\eta^2=0.52$; pairwise comparisons are provided in Supplementary Table 4) but not cue validity ($F_{2,38}=0.04$, $P=0.96$, partial $\eta^2=0.002$, $BF_{excl}=24.24$, where BF_{excl} is the Bayes factor comparing a model excluding versus including the factor of validity). We therefore constrained the motor time in our model to these onset latencies per regime (−97.5, −82 and −69 ms for Easy, LoCoh and Deadline, respectively), with no variation across cue validities. The remainder of non-decision time is classically attributed to sensory encoding delays beginning at stimulus onset[3]. We reasoned that sensory evidence representations are themselves unlikely to differ appreciably in their timing across regimes, whereas the time at which accumulation begins is

Table 1 | Goodness-of-fit metrics

	NI model			DDM		
Model version	G^2	**AIC**	**BIC**	G^2	**AIC**	**BIC**
All regime effects, drift rate bias	46.9	80.9	182.1	168.8	204.8	311.9
No deadline effect on drift rate, drift rate bias	54.5	86.5	181.8	174.9	206.9	302.2
No regime effect on tnd, drift rate bias	59.5	89.5	178.8	579.0	611.0	706.2
No regime effect on urgency rate, drift rate bias	58.1	88.1	177.4			
All regime effects, no drift rate bias	52.7	80.7	164.1	169.2	199.2	288.5

Goodness-of-fit metrics are provided for the full NI model and the full drift diffusion model (DDM), as well as versions of the models with each non-dominant regime or priors effect omitted in turn. We provide G^2 as well as the AIC and Bayesian information criterion (BIC) metrics, which penalize less and more strictly for complexity, respectively, for reference. Given the purpose of the study to establish the full range of effects with convergent support from neural data, we used AIC for model selection throughout. Lower values of G^2, AIC and BIC indicate a better fit. For the NI model, tnd indicates accumulation onset only, because post-decision motor times were constrained by the neural data; for the DDM, tnd refers to the total non-decision time.

probably adjustable, as suggested by recent neurophysiological evidence of pre-evidence accumulation under some circumstances[36], and behavioural evidence for its strategic adaptation across speed/accuracy regimes[26]. We therefore included a single evidence onset parameter and three independently adjustable accumulation onset parameters for the three regimes.

Behavioural modelling results. We fit the NI model developed above to behaviour and compared its outcomes to a standard DDM[37] with a similar number of free parameters (the full parameter listings are provided in Supplementary Tables 1 and 2). The first question is whether the NI model, having had key parameters constrained by EEG signal measurements, can quantitatively capture behaviour as effectively as the versatile DDM. In fact, the NI model provided a superior fit (Table 1), which was mainly attributable to its ability to recapitulate the lack of rightward skew in the RT distributions, the fast, purely prior-cue-based responses in the Deadline regime and the full extent of the variations in accuracy as a function of RT (Fig. 3).

The more important question for our present purposes is what these models indicate regarding the key parameter adjustments that are made to adapt to speed/perceptual demands and prior probability. As models were fit to grand average data (point estimates of all of the parameters are provided in Supplementary Tables 1 and 2), we assessed the reliability of parameter effects across participants by fitting the DDM and NI models to 500 bootstrap samples of the 20 participants and computing the 95% confidence intervals (CI) of the differences with respect to Easy and neutral trials (Methods). As expected, the DDM indicated criterion shifts that agreed with those reflected in neural motor preparation starting-level modulations; the bound separation was approximately halved in the Deadline regime (0.070 versus 0.037; difference CI = [0.030,0.036]) and biases towards the more probable alternative in all of the regimes (starting-point bias for Easy, CI = [0.0010,0.0036]; LoCoh, CI = [0.0041,0.0080]; Deadline, CI = [0.006,0.011]). The models also generally agreed on the effects of weak evidence; both of the models indicated the expected drift rate reduction in the LoCoh regime (LoCoh drift rate as a proportion of Easy for the DDM, 0.45, CI = [0.41,0.48]; and for the NI model, 0.39, CI = [0.35,0.44]), and a lengthening of non-decision time under weaker evidence in the DDM (0.537 s versus 0.424 s, difference CI = [0.077,0.148]) mirrored by a similar lengthening of the preaccumulation delay in the NI model (0.269 s versus 0.162 s, difference CI = [−0.009,0.186]), countering the slighter shortening of motor execution delay measured in pre-response potentials (Fig. 2c).

However, the DDM and NI models drew fundamentally different conclusions regarding build-up and delay adaptations to speed pressure (Fig. 4). The DDM fit replicated several recent modelling studies[6,11] in finding a reduced drift rate in the Deadline regime relative to the Easy regime (as a proportion, 0.88, CI = [0.80,0.98]) and shortened non-decision time (by 0.162s, CI = [0.134,0.197]). By contrast, the NI model indicated an enhanced drift rate (factor of 1.39, CI = [1.23,1.75]) alongside a reduced urgency rate (0.47 versus 0.66, difference CI = [0.00,0.30]) and only a small (29 ms) but reliable shortening of the motor execution component of non-decision time (Fig. 2c), with no reliable change to accumulation onset (difference CI = [−0.041,0.077]).

The regime effects reported by each model were significant in the sense that Akaike information criterion (AIC) values increased substantially for alternative model versions that were constrained to have one non-decision time across regimes, a single drift rate per coherence or—in the case of the NI model—a single urgency rate (Table 1). These effects were reliably reproduced across many randomly seeded model fit iterations (Extended Data Fig. 1) as well as across bootstrap samples (Extended Data Fig. 2).

As a function of prior probability, both models reliably indicated positive drift rate bias parameters in all of the regimes (Supplementary Tables 1 and 2 and Extended Data Fig. 2d). In fact, the behavioural data contained a qualitative signature of drift rate bias in the form of significant prior probability effects on accuracy for even the very longest RTs (paired *t*-tests on the seventh RT bin of the conditional accuracy function of Fig. 3a; Easy, $t_{19} = 2.28$, $P = 0.034$, CI = [0.004,0.088]; LoCoh, $t_{19} = 5.23$, $P < 0.001$, CI = [0.070,0.164]; Deadline, $t_{19} = 2.50$, $P = 0.022$, CI = [0.009,0.104]), a pattern that tends not to be predicted by starting-point biases alone (Supplementary Fig. 4). However, whereas the NI model produced an almost identical AIC value with and without a drift rate bias parameter (Table 1), the improvement in fit of the DDM with the inclusion of drift rate bias was very small and did not overcome the complexity penalty of the AIC. This is presumably because the abovementioned conditional-accuracy-based signature of drift rate bias represents a relatively small section of the full behavioural data in this experiment.

To further validate the model fitting procedure and results, we fitted the models to simulated data with manipulated build-up and delay effects. This confirmed that both the superior fit quality of the NI model and the DDM's misestimation of drift rate and non-decision time under speed pressure are systematically related to the presence of pre-evidence accumulation and dynamic urgency (Extended Data Fig. 3). Indeed, the DDM and NI model performed similarly when data were simulated without these two features. We then validated the procedure for selecting among NI models (Table 1) by verifying that, when a given parameter adjustment effect was absent from simulated data, the AIC values correctly favoured the model that included no such effect. Furthermore, when the parameter effect was introduced at increasing magnitudes, the AICs were observed to increasingly favour the model that allowed for the effect (Extended Data Fig. 4).

Model validation through independent neural data. We next tested the validity of the NI model using neural decision signal data that were not used to construct or constrain the model. For this, we leveraged a distinct neurophysiological signature of decision formation that is known as the centro-parietal positivity (CPP)[19,38]. Similar to MB, the CPP undergoes an evidence-dependent, RT-predictive build-up during decision formation[19,39], but also exhibits important functional differences. First, the CPP has been shown to be motor-independent, tracing cumulative evidence even when no

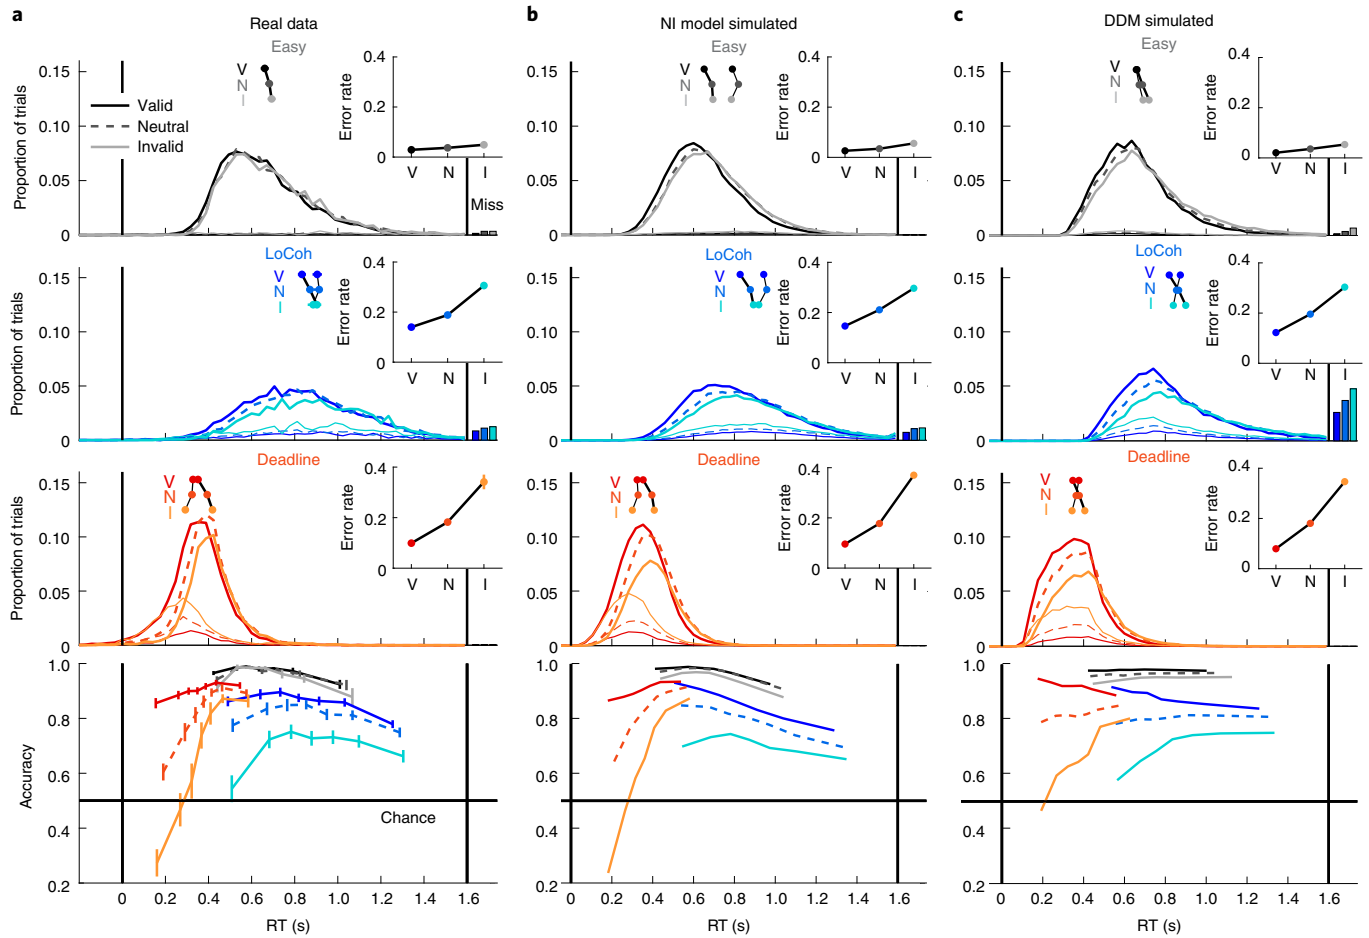

Fig. 3 | Comparison of real versus model-simulated behaviour. a, Empirical RT distributions for correct (thick) and incorrect (thin) responses under the Easy, LoCoh and Deadline regimes, averaged across participants. The proportion of misses (no response within 1,600 ms of evidence onset) for each condition is shown as a bar on the right. The mean RT reproduced from Fig. 1b is marked above the RT distributions, and the mean error rates reproduced from Fig. 1c are plotted in the insets to assist with comparison with the model-simulated data. Data are mean ± s.e.m.; error bars are, in some cases, smaller than the marker size. At the bottom, conditional accuracy functions are shown for all of the conditions, quantifying the accuracy in each of seven equal-sized RT bins. To follow-up on a regime × validity × RT-bin interaction ($F_{24,456} = 14.3$, $P < 0.001$, partial $\eta^2 = 0.43$), a validity × RT-bin RMANOVA was performed in each regime, revealing significant main effects of both factors in all three regimes, and a significant validity × RT-bin interaction in the Deadline and LoCoh regimes (Supplementary Table 5). Follow-up tests focusing on the inverted-U nature of the functions showed that the slowest RT bin had significantly lower accuracy than the second-slowest bin for LoCoh and Easy regimes (but not the Deadline regime, which lacked slow RTs in general), and the fastest RT bin had significantly lower accuracy than the second-fastest in all three regimes (Supplementary Table 6). **b**, Simulated behavioural data from the NI model, plotted in an identical manner to the real behaviour. **c**, Simulated behavioural data from the constant-bound DDM. A version of this figure for the models without drift rate biases is provided in Supplementary Fig. 4. Supplementary Figure 5 shows conditional accuracy functions for each individual participant, demonstrating that grand average behavioural patterns are representative of those exhibited by individual decision-makers.

overt action is required or when stimulus-to-response mappings are unknown[19,40]. Second, the CPP and MB have been shown to undergo distinct modulations under speed pressure. In contrast to MB, which adjusts its starting levels and reaches a fixed threshold at response, it is the CPP's amplitude at response that varies inversely with urgency and not its starting level[25]. These characteristics are consistent with the CPP purely reflecting the cumulative evidence component of the decision process, unadulterated by the urgency components that manifest at the motor level. Thus, the CPP offers a means to independently test the predicted impact of the adaptations indicated by the NI model on evidence accumulation dynamics.

To generate dynamic decision variable simulations with which to compare the empirical MB and CPP waveforms, the urgency and evidence accumulation components of the NI model's decision process were encoded and recorded separately for each trial. Their summed contributions were assumed to comprise the decision variable at the effector-selective, action-triggering level (MB, just as assumed for LIP[22,23]), whereas the CPP was assumed to reflect only the cumulative evidence, in rectified form[41] (a detailed model implementation is provided in Supplementary Fig. 3). Qualitatively, the real (Fig. 5a,b,d,e) and simulated (Fig. 5h,i,k,l) signals appear to be very well matched both in their initial post-stimulus build-up dynamics and their final dynamics and amplitudes just before the response.

Two aspects in particular serve to affirm important structural elements of the model. First, the combination of starting-point variability and dynamic urgency predicts that the level of cumulative evidence reached at response varies as an overall inverted-U function of RT (Fig. 5n). Specifically, the very fastest responses in the deadline regime tend to occur when the process randomly starts close to the correct bound and, therefore, needs little cumulative evidence to reach it; beyond those earliest RTs, the cumulative

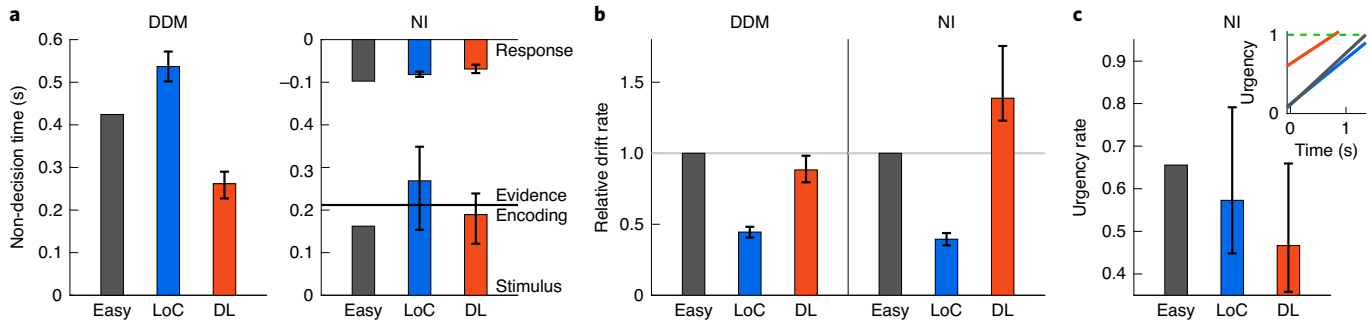

Fig. 4 | Parameter estimates showing regime effects of the DDM and NI models. a, The unitary non-decision time estimates of the DDM (left), contrasted with the NI model's (right) separate estimates of evidence-encoding onset (horizontal line), accumulation onset (bottom bars) and post-decision motor times constrained by the neural data (top bars). **b**, Estimated drift rate, normalized with respect to the value in the Easy regime to facilitate cross-model comparison. **c**, Urgency rate estimated by the NI model. Inset: the linear time functions representing the full urgency signals, including both offset (starting point) and slope, demonstrating that, despite the shallower slope under the Deadline regime, it is aimed at crossing the threshold much earlier due to starting point elevation. The error bars show the 95% CIs of the parameter effects relative to the easy regime estimated from 500 bootstrap samples (Extended Data Fig. 2). Note that the variability indicated by these CIs reflects not only the across-participant variability in the adoption of the adjustment strategies, but also the reliability with which parameter effects can be estimated using our model and fitting procedure, which in turn is influenced by the cross-participant variations in the quality of the EEG signals used in the neural constraints.

evidence required to trigger a response decreases with RT owing to the increasing contribution that urgency makes to the decision variable[25]. This very pattern was borne out in the empirical CPP amplitudes (Fig. 5g; RT-bin × regime interaction in RMANOVA, $F_{8,152}=3.61$, $P=0.048$, partial $\eta^2=0.16$; correlation with model-simulated cumulative evidence, Pearson $r=0.75$, $P<0.001$ by permutation test). Furthermore, as we described previously[25], there was no effect of prior cue or regime on the CPP measured at evidence onset (Fig. 6; 3×2 RMANOVA; Supplementary Table 7), providing additional confirmation that the CPP does not directly receive an urgency input.

Second, the model assumes variable and often earlier accumulation onset with respect to evidence encoding. A direct consequence is that the trial-averaged accumulator signal begins to climb distinctly earlier than the time at which sensory evidence causes it to accelerate, a feature that is apparent in both the simulated and real CPP waveforms (one-sample *t*-test on signal slope in first 118 ms post-evidence, $t_{19}=2.46$, $P=0.024$, CI = [0.0012,0.0144]; no significant effects of regime or cue in 2×3 RMANOVA; Supplementary Table 8).

In addition to these affirmations of the model's structure, the neural decision signal dynamics provided further empirical signatures of the various adaptations across regimes and prior probability conditions indicated by the NI model:

Regime effects. A salient speed pressure adaptation indicated by the DDM was a ~150 ms shortening of non-decision time in the Deadline relative to the Easy regime, whereas the NI model did not indicate a significant difference in accumulation onset between these regimes. Indeed, both the simulated and real stimulus-locked CPP appear to accelerate at approximately the same time. To assess this quantitatively, we compared the time at which the stimulus-locked CPP reached 20% of its maximum value in the Easy and Deadline regime, and this differed by no more than 1 ms on average (jackknife *t*-test against 0, $t_{19}=0.06$, $P=0.95$, $BF_{01}=4.3$, where BF_{01} is the Bayes factor for the null versus alternative hypothesis).

The NI model indicated an enhanced drift rate in the Deadline regime, opposite to the reduction effect indicated by the DDM. Build-up rates of the CPP signal agreed with the NI model (stimulus-locked slope, $F_{2,38}=21.5$, $P<0.001$, partial $\eta^2=0.53$; response-locked, $F_{2,38}=3.97$, $P=0.027$, partial $\eta^2=0.17$; slope enhancement under the deadline regime reached significance for only the response-locked waveforms, both when measured in time windows aligned to the response, and aligned to the regime-specific motor times; pairwise comparisons are provided in Supplementary Table 9; Fig. 5d,e). Note that more points were awarded to participants when correct and on-time in the Deadline regime to offset the risk of demotivation due to the regime's increased difficulty. This led to a slightly (12%) higher total number of points obtained in the Deadline regime compared with the Easy regime, but this did not significantly relate to the drift rate enhancement (Supplementary Fig. 7 and Discussion).

The NI model also indicated a reduced rate of urgency build-up in the difficult regimes, particularly the Deadline regime. The neural signatures of motor preparation (Fig. 5a,b) exhibited build-up rates that ranked in the same way as the model-estimated urgency rates; the deadline condition exhibited the shallowest slope before the response (stimulus-locked, $F_{2,38}=3.73$, $P=0.033$, partial $\eta^2=0.16$; response-locked, $F_{2,38}=7.25$, $P=0.005$, partial $\eta^2=0.28$; both regimes were significantly different from easy; Supplementary Table 10), as in the model-simulated waveforms (Fig. 5i).

Prior probability effects. Technically, the AIC comparisons indicate that the inclusion of drift rate bias in the DDM is not justified against the increased complexity. By contrast, the NI model indicated a similar AIC with and without drift rate bias. The neurophysiological data provided crucial additional supporting evidence for drift rate biases. Specifically, pre-response CPP amplitude was modulated significantly by prior probability in the deadline regime but not the others (one-way RMANOVAs testing validity effects in each regime, following a 3×3 RMANOVA regime × validity interaction, $F_{4,76}=3.27$, $P=0.046$, partial $\eta^2=0.15$; Supplementary Table 11). Indeed, with drift rate bias included, the model-simulated amplitudes mirrored this tendency for prior biases to disproportionately affect amplitude in the deadline regime (Fig. 5m). Validity effects on CPP amplitude arise due to starting-point shifts, because the closer to the bound the motor signals start, the less scope there is for evidence accumulation. Why would these amplitude modulations manifest in the deadline regime but not LoCoh? As shown in Fig. 7a, due to the growing dynamic urgency component, the negative influence of a starting-point increase on the peri-threshold accumulator amplitude is cancelled out by the positive influence of drift bias. To examine this further, we repeated the model fits but this time we included the degree of mismatch between simulated

and real CPP amplitude modulations in the objective function as a penalty term. We found that the model that included drift rate bias increasingly outperformed the model with no drift rate bias as the emphasis on capturing the CPP modulations was increased (Fig. 7b).

Discussion

Adaptations to speed pressure and prior probability are perhaps the most well-studied manifestations of cognitive flexibility in sensorimotor decision-making, yet a comprehensive understanding of all of the decision-process adjustments that underlie this ability has eluded the field, largely due to the limited constraints afforded by behavioural data. Here we leveraged the additional constraints offered by well-characterized human neurophysiological decision signals to construct a model that enables a more detailed parsing of computational adaptations without unduly increasing the number of free parameters. Moreover, our approach enabled the model to be evaluated not only on the basis of quantitative fits to behaviour but also on the basis of its ability to explain dynamic neural signal features that were not used in constructing the model.

While behavioural modelling studies have tended to converge on boundary adjustments as the principal mediator of speed–accuracy trade-offs, some recent studies have suggested that speed pressure may also shorten non-decision time and reduce drift rate[6,11,28–33]. Although our DDM fit replicated these effects, our NI model and neural decision signals provided strong convergent evidence for an enhancement of drift rate, not a reduction. This is consistent with recent findings that representations of the sensory evidence itself are enhanced under speed pressure along with corresponding increases in decision signal build-up rate[25,42]. Our participants earned 12% more points overall in the Deadline regime compared with the Easy regime and, although the much greater regularity of negative feedback makes it improbable that participants perceived the Deadline regime as generally more valuable, it was important to consider the possibility that the drift rate boost might reflect an automatic enhancement of processing due to the participants experiencing higher overall rewards[43–45]. Several analyses speak against this (Supplementary Fig. 7): (1) bootstrap samples with stronger drift rate boosts were not those obtaining greater rewards; (2) model simulations show that, without drift rate enhancement, rewards would have been no greater than in the easy regime; and (3) other elements of processing were not enhanced (for example, reduced urgency rate). Thus, rather than an automatic reward-related enhancement, the drift rate boost probably reflects a motivated, strategic adaptation engaged to recoup the accuracy losses induced by lower decision bounds.

Our neurally constrained model made it possible to parse what is usually a unitary non-decision time parameter into its key pre- and post-decision constituents. In doing so, the model revealed that these constituents are each independently adjustable. Consistent with recent electromyography findings[25,46], we found that motor execution times, estimated from response-locked potentials, were shortened under speed pressure. Motor time was also shortened in the weak-evidence (LoCoh) regime, whereas accumulation onset was lengthened, presumably because the longer deadline allows for postponement of accumulation until evidence is strongly encoded, in line with recent proposals of strategic accumulation onset adjustments[26,36,47].

The question of whether decision bounds tend to be constant or collapsing has been contentious. Several neurophysiological studies in monkeys[22,23] and humans[24,25] have reported that motor preparation has a dynamic urgency component that implements a collapsing bound. Meanwhile, most behavioural model comparison studies have concluded in favour of constant bounds, although it has often been difficult to account for the increased complexity and, therefore, penalization of collapsing bound models[8,10]. Here we observed a striking empirical signature of dynamic, evidence-independent urgency in the early launch of motor preparation before evidence onset, particularly in the deadline regime. Such pre-evidence motor preparation was in fact proposed in early EEG research[48] and was demonstrated in neuronal activity in monkey frontal eye fields during compelled response tasks in which the 'go' command preceded the target stimulus[49]. Our NI model revealed that dynamic urgency had an important role in all three regimes. Interestingly, its build-up rate was slowest in the deadline regime, but this was coupled with a much earlier onset from an already increased level, so that it was nevertheless aimed at crossing the bound much earlier than in the longer-deadline regimes (Fig. 4c). The reason for adopting this particular urgency strategy is unclear, but it may relate to the narrow time window for rewarded responses in this regime. Given that the variability of temporal estimation scales with the duration that is estimated[50,51], and assuming that the variability of urgency build-up rate may similarly scale with its mean, it may be that earlier and shallower launches of urgency can be aimed more precisely. Although further research is needed to fully establish the principles that govern such adaptation patterns, these results serve to demonstrate the extent of strategic flexibility of urgency signals.

Drift rate biases induced by prior stimulus probability have rarely been reported in previous behavioural modelling research[3,28,52], and the present study affirms that they are difficult to detect using standard models and model-selection procedures. Although the accuracy biases for long RTs provide a qualitative signature of their presence, the AIC metrics indicated that their inclusion is justified against the increased complexity only for the NI model. We were able to confidently conclude in favour of the biases on the basis of their ability to account for CPP amplitude modulations. This finding complements another recent study that showed strong drift rate biases due to relative reward value in severely time-constrained

Fig. 5 | Real and simulated average decision signal dynamics and amplitudes. All data are for correct trials only. **a**,**b**, Stimulus-locked (**a**) and response-locked (**b**) MB amplitude reflecting motor preparation contralateral and ipsilateral to the (correct) response executed. Waveforms were averaged across cue validities to maximize signal robustness (traces broken out by validity are provided in Supplementary Fig. 6). **c**, Amplitudes of contralateral and ipsilateral MB just before the response, broken out by prior cue validity. **d**,**e**, Stimulus-locked (**d**) and response-locked (**e**) CPP for each regime. Note that, although the CPP continues to rise and peak beyond the point at which we measure its pre-response amplitude (a 59 ms window centred −97.5 ms before the response, that is, the earliest motor execution onset), this probably represents post-decision accumulation[76,77], effectively giving rise to an overshoot in the CPP's trajectory[25]. **f**, CPP amplitude broken out by prior cue validity. **g**, Pre-response CPP amplitude measured in five RT bins for each prior cue type and regime highlights an overall inverted U-shaped relationship with RT. For **a**–**g** (that is, all of the plots showing the real data), data are mean ± s.e.m. after between-participant variability was factored out. **h**–**n**, The same plots as in **a**–**g** were produced for simulated decision signal trajectories from the best fitting NI model. Note that, whereas the real signals fall back or rebound at or shortly after the time of the response, our simulations continue to evolve beyond this point; because these post-decision dynamics are immaterial to the accuracy and timing of the action choices, there was no imperative to mimic these fall-down dynamics. Note also that all simulated CPP waveforms have been shifted on the time axis by a fixed amount to approximately line up with the real waveforms, in line with an estimated 134 ms delay between cumulative evidence encoding at the CPP level and its impact at the motor preparation level (Methods). Note that the same *y*-axis is used across all graphs within each row of the figure; graphs **a**–**c** use the *y* axis from **a**, **d**–**g** use the *y* axis from **d**, **h**–**j** use the *y* axis from **h** and **k**–**n** use the *y* axis from **k**.

sensorimotor decisions[41], and adds to evidence of drift rate biases arising from other sources, such as previous exposure in recognition memory tasks[53], biases towards 'no' responses[54] and choice history[55].

In addition to providing more accurate inferences regarding parameter adjustments, the goodness of fit of the NI model to behavioural data was far superior to that of the DDM in this dataset. Model recovery simulations verified that the degree of this fit superiority was due to the presence of early accumulation onsets (that is, preceding evidence encoding) and dynamic, anticipatory urgency (Extended Data Fig. 3). These features are probably induced by the imposition of response deadlines, and bring about a reduced rightward skew in RT distributions and many purely prior-based

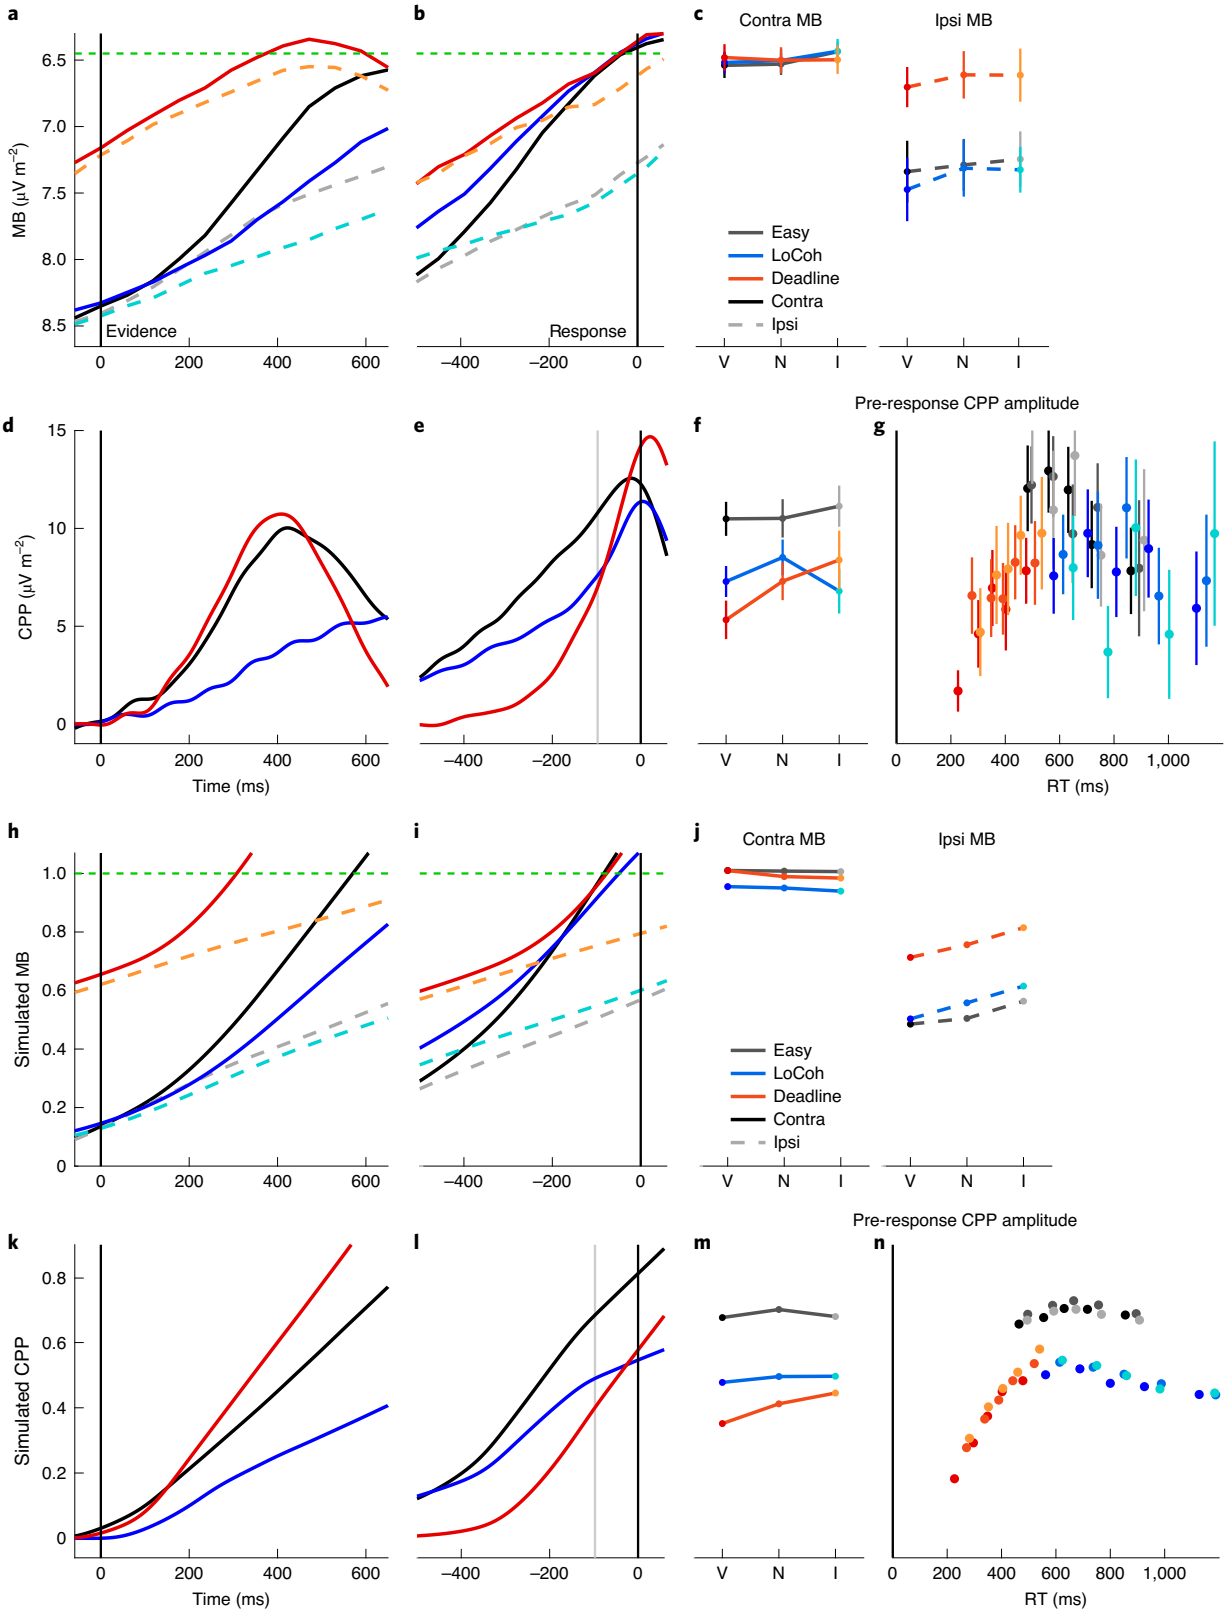

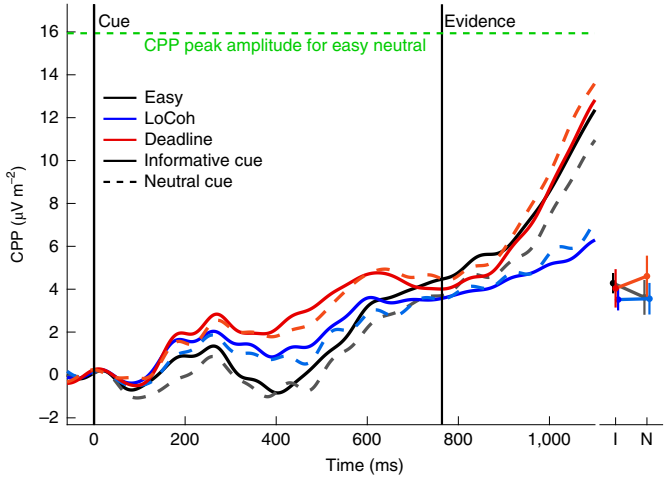

Fig. 6 | Cue-evoked ERP waveforms. Event-related potential waveforms for neutral and informative cues in the cue-evidence period, at electrodes in the focus of the CPP. Note that, as the colour cue itself must be decoded and processed, there is a positive-going shift in the waveform during the cue-evidence interval, which is greater in the regimes that strongly call for incorporation of the prior probability information (main effect of regime at 450 ms, $F_{2,38} = 5.86$, $P = 0.021$, partial $\eta^2 = 0.24$), but this difference dissipates by the time of evidence onset (Supplementary Table 7). Amplitudes in the pre-evidence time interval −59–0 ms are shown on the right; the error bars show the s.e.m.

fast responses in the deadline regime that the DDM cannot capture by construction. In the DDM, fast, purely prior-based responses are precluded by the fact that the process is suspended during the non-decision time and choices can be driven only by evidence that favours the correct alternative. Our NI model enables the timing of multiple stages of the decision process to be independently estimated, therefore allowing for the possibility that both urgency and the accumulation process can be initiated in advance of stimulus-discriminating evidence becoming available. This in turn allows for more phases of biased versus accurate responding, as expressed in these data. Another noteworthy feature of our task is that the coherent motion did not suddenly appear against a blank field but, rather, transitioned from incoherent motion to avoid contaminating evoked potential components[19,39]. This in itself is unlikely to have been a critical factor in our findings because the evidence nevertheless did appear suddenly through a step change from random to coherent motion and this transition occurred after a fixed delay that participants were extensively trained to estimate using high-coherence trials. Although the degree of fit superiority of the NI model is unlikely to be so large for tasks that are more DDM-compatible (for example, with no deadlines[56] and lower RT limits[3]), the present findings affirm the general principle that, if the structure of the model ignores important elements of the neural decision process, inferences can be substantially misattributed.

Our findings demonstrate the value not only in the use of neural data to inform cognitive models, but also, conversely, in the application of cognitive modelling to correctly interpret neural signatures

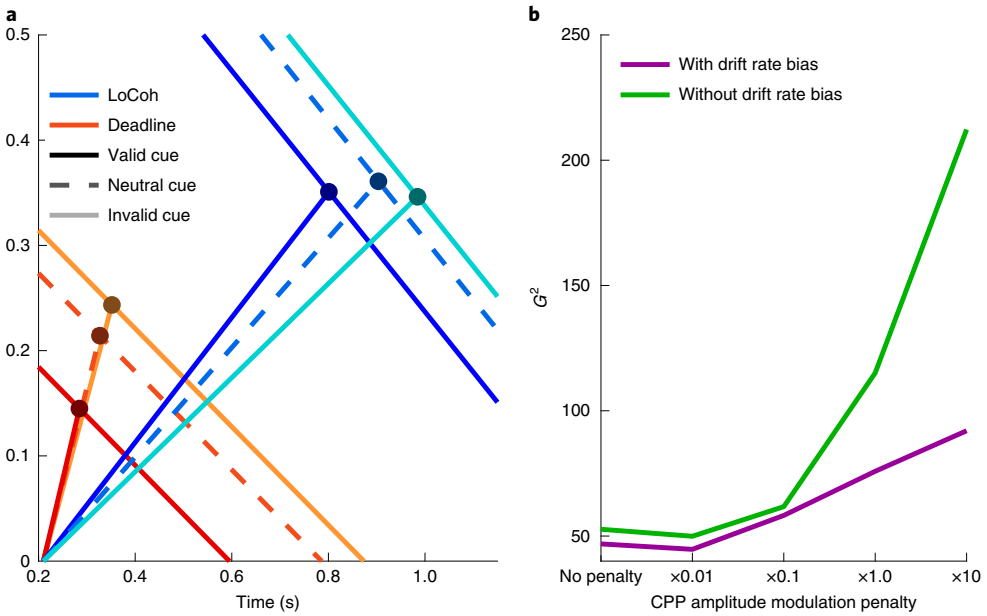

Fig. 7 | Relationship between drift rate biases and CPP amplitude modulations. a, Simulation of average evidence accumulation trajectories (upward-sloping lines) and collapsing bounds (downward-sloping lines) using the NI-model-estimated parameter values in the LoCoh and Deadline regimes, demonstrating how the accumulator's height at bound-crossing, reflected in pre-response CPP amplitude, is impacted by drift rate biases. A growing urgency component contributing to motor preparation effectively sets a collapsing bound on evidence accumulation, and prior probability biases in starting motor preparation shift this bound downwards and upwards for trials with valid and invalid cues, respectively. For a given drift rate, the latter bound shifts cause the accumulator to reach higher levels for the invalid trials compared with the valid trials. However, a bias in drift rate tends to counteract this effect because the slower build-up on invalid trials makes the accumulator cross the collapsing bound later and, therefore, lower down. In the LoCoh regime, the bounds are high and the drift rate is low, such that a relatively small bias in the drift rate is enough to fully cancel out the effect of starting-point bias on the terminal accumulator amplitude. **b**, G^2 values quantifying goodness of fit to the behavioural data for model fits, using an objective function with an increasing penalty term that quantifies divergences of the priors effect on CPP amplitude simulated from the model from that of the real data. The larger this penalty term, the more the parameter values of the fit are dictated by the CPP amplitudes instead of the behavioural data; therefore, increasing this penalty term will necessarily decrease the goodness of fit to behaviour. However, a greater emphasis on capturing the CPP amplitude modulations results in the model with drift rate biases becoming distinctly superior to the model without drift rate biases. This reflects the fact that, without drift rate biases, it is not possible to produce prior probability modulations of terminal amplitude as small as those shown in **a** or observed in both the real and simulated CPP (Fig. 5f,m). Vertical axes in both graphs are in arbitrary units.

of decision formation. Although certain neural signal features translate relatively directly to corresponding model parameters, such as pre-evidence starting levels and their adjustment across conditions, other signal features, such as build-up rate, bear a less direct correspondence and are better analysed with reference to dynamic model simulations[57]. Such simulations were instrumental here. For example, although we replicated our previous finding[25] of a CPP amplitude reduction over RT in the longer-deadline conditions—which is diagnostic of a collapsing bound—the deadline regime showed the opposite trend despite the clear indications of dynamic urgency in the model and MB traces. The model simulation confirmed that this increase in CPP amplitude over RT arises even with a collapsing bound due to the large degree of starting-point variability relative to the low bound.

At the same time, some cautionary notes are warranted regarding this approach. Although informed by neural signals, our model is still very much a process model abstracted from neural implementation, especially from the level of spiking neural networks for which there are now a growing suite of powerful models[58–60]. As applied here, our NI model relies solely on well-established functional characteristics of the neural decision signals, but much research is still needed to expose the neural architecture that gives rise to these signals and their properties for the EEG-informed models to further contribute to our understanding of the finer details of the neural systems implementing decision computations. Furthermore, although neural constraints enable the exploration of a wider range of processes and parameters, not all aspects of model structure will have strong guidance from neural data, and researcher degrees of freedom must be carefully considered just as in conventional behavioural modelling[11]. Finally, the NI modelling approach that we applied here dealt with aggregate grand-average data to ensure the reliability of neural constraints, but future work will have to explore the degree to which single-participant EEG can be used to model individual differences, which would be particularly valuable in clinical research.

In summary, our findings add to the growing demonstrations that decision-process adaptations to speed pressure and prior probability are multifaceted. They further show how establishing the degree and nature of those adaptations can critically depend on model structure, and that neural signatures of the decision process can offer critical guidance in this regard.

Methods

Participants. Twenty-two human participants (8 female; aged 20–32 years) each participated in three experimental sessions, the first for psychophysical training and the remaining two for EEG recordings during task performance. No statistical methods were used to predetermine the sample size, but it was larger than that of previous studies that reported speed/accuracy manipulation effects[25]. Two of the participants were excluded from all analyses due to excessive perspiration artifacts in at least one EEG session. All of the participants had normal or corrected-to-normal vision. Informed written consent was obtained from all of the participants, and all of the procedures were approved by the Institutional Review Board at The City College of New York.

Set-up. The participants were seated in a dark, electromagnetically shielded booth, with their heads stabilized in a chin rest with forehead support placed 57 cm away from a cathode ray tube monitor (frame rate 85 Hz; resolution, 1,024 × 768) with a black background. The participants rested their hands on the table in front of them, with the forefingers of their left and right hands resting on the left and right buttons of a symmetrically shaped computer mouse, which formed the response alternatives. Eye position was monitored continuously throughout task performance using a remote eye tracker (EyeLink 1000, SR Research, 1,000 Hz).

Task. A trial was initiated by the participant by simultaneously pressing both mouse buttons, provided their measured eye position was within 3° of a central, white, 5 px × 5 px fixation point. After a short delay of 50 ms, grey incoherently moving dots appeared, changed colour after 647 ms (11 flicker frames) and finally became coherent after a further 764 ms (13 flicker frames; Fig. 1a). The fixed timing meant that the participants could predict evidence onset even under weak evidence. The seamless transition from zero to non-zero coherence at evidence onset avoids non-task-specific evoked potentials caused by sudden luminance increments, which would otherwise overlap with and obscure the relevant decision signals[39,40,61], while nevertheless producing a sudden-onset of evidence as in other paradigms. Cue colours were yellow, green and cyan, assigned to indicate a 75% prior probability of leftward, neutral (50/50) and 75% rightward motion; the colour meanings were fixed for a given participant but counterbalanced across participants. The three prior-probability colours were each equiluminant with the initial light grey colour of the dots, as verified using a photometer. The dots remained at a constant coherence level for 1,600 ms before the display was extinguished for all of the conditions. If the participant indicated the correct decision before a deadline, by a left- or right-hand button press for left or right motion, respectively, he/she earned points that translated to monetary value. Trials ended with feedback text indicating the number of points won if correct and on time, or indicating zero points and the reason if a response was 'wrong' or 'too slow'. The participants were instructed to maintain fixation for the duration of each trial, until feedback was presented. Participants performed the task under three regimes, which were run in separate blocks of trials: (1) Easy, which had high coherence (20%) and a long deadline (1,600 ms); (2) LoCoh, which had low coherence (ranging from 5% to 12%, individually titrated; see below) and a long deadline (1,600 ms); and (3) Deadline, which had high coherence (20%) and a short deadline (388–640 ms, individually titrated).

The dot-motion stimulus comprised an average of 118, 6 px × 6 px dots presented within an aperture of 8° diameter (density, 2.3 dots per degree squared) centred on fixation, flickering on and off at a frame rate of 17 Hz (two refreshes on, three off)[40]. The on–off flicker was incorporated originally to provide a steady-state visual evoked potential (SSVEP) but, as in our other recent studies, we did not analyse this because it does not represent a sensory evidence signal specifically[39]. During incoherent motion, dots were placed randomly and independently within the aperture, whereas, during coherent motion, a proportion of the dots were randomly selected on each frame to be displaced by a fixed distance of 0.353° in either the leftward or rightward direction on the following frame, resulting in a motion speed of 6° s^{-1}. The task was programmed using PsychToolbox[62] for MATLAB (MathWorks).

Training. To practice the task and to determine individually titrated levels of difficulty in the Deadline and LoCoh regimes, the participants first underwent a training session. In this session, a series of short blocks of 24 trials were first run with a 1,600 ms deadline, starting at 70% coherence and stepping down by 10–20% after achieving >90% correct responses at each level until the participant reached 20% coherence and was fully accustomed to the fixed timing of events and the motion speed. A 60-trial adaptive staircase procedure (one up, three down)[63] was then run to estimate the coherence level at which the participant performed at 80% accuracy with the same 1,600 ms deadline, and this was repeated until a stable level was indicated. Next, the deadline was changed from 1,600 ms to the median RT from the last 20%-coherence block and 48 trials were run to estimate the participant's typical RT distribution under time pressure. The deadline was then further shortened towards the median RT of that fast block to further increase time pressure and, provided that participants reported experiencing intense time pressure, this value was set as the deadline in the Deadline regime. As a final step in the training session, the colour cues were introduced. The participant performed a practice block of 96 trials of each of the three finalized task regimes to end the training session.

Procedure. Over two EEG recording sessions following the training, each participant performed 10 blocks of 96 trials of each of the three regimes (two participants spread the 30 blocks over 3 sessions due to fatigue). Five blocks of a given regime were run consecutively at a time to ensure that the participants would settle into a strategy and to minimize cross-over. The orders of the regimes were randomized across sessions and participants. Small adjustments to the LoCoh coherence level or to the tight deadline in the Deadline regime were made by the experimenter as necessary from block to block when the difficulty seemed to be too easy or difficult. To minimize the degree to which blocks of the different regimes differed in overall value, participants were rewarded 40 points per correct trial in the Easy regime, 50 points in the LoCoh regime (recall titration for 80% accuracy) and 70 points for the Deadline regime (less than double the easy regime to allow for further reduction in median correct RT with practice). Using this scheme, across the 20 participants, the average reward per trial was 40.6, 40.4 and 45.5 points for the Easy, LoCoh and Deadline regimes, respectively—a significant, albeit small, difference (one-way RMANOVA, $F_{2,38} = 8.43$, $P = 0.007$, partial $\eta^2 = 0.31$; post hoc differences significant for Deadline compared with both Easy ($t_{19} = 2.69$, $P = 0.015$, CI = [1.08,8.73]) and LoCoh ($t_{19} = 3.48$, $P = 0.0025$, CI = [2.02,8.11]). The reason that the average reward per trial exceeded 40 points in the Easy regime was because the points per correct trial was mistakenly set to 70 for 5 of the 200 easy blocks (across all of the participants) and to 50 for 27 of them. One point corresponded to US$0.0031 and, at the end of the experiment, two blocks were randomly selected within each regime, and the total points of those six blocks translated into a dollar amount in compensation for the participants' time, in addition to US$12 per hour for training and EEG setup time. Data collection and analysis were not performed blind to the conditions of the experiments.

Behavioural data analysis. RTs were measured relative to the onset of coherent motion evidence. The first button click following the prior cue was always logged as the response for a given trial, even if made before evidence onset. Such negative RTs were regarded as a legitimate manifestation of the mechanisms employed to perform the task, rather than as contaminant responses[64]. Similarly, responses that occurred after the deadline in the Deadline regime were included in all analyses and modelling. Trials in which no response was made by 1,600 ms were counted as misses. These trials are shown in the behavioural data (Fig. 3a), but were excluded from electrophysiological analysis because they cannot be response-locked. In RMANOVAs conducted on mean RT, we tested only correct RT because the Easy regime contained insufficient error trials. To compute RT distributions for plotting (Fig. 3a), we divided the trials into non-overlapping RT bins of 35.3 ms width starting from -176 ms and ending at 1,600 ms, and separately counted correct trials and error responses that fell into each RT bin. This was performed for each individual participant, expressed as a proportion of the total trials for that participant, and then averaged across participants. To compute conditional accuracy functions, we first divided each participant's pooled (correct and error) RT distribution into seven equally sized bins, then computed the proportion of correct trials within each of these RT bins and finally plotted response accuracy over the mean RTs of the bins for each condition, averaged over participants (Fig. 3a). A 3 regime × 3 validity × 7 RT-bin RMANOVA was run on conditional accuracy functions and the three-way interaction was followed up with pairwise tests of bin 1 versus bin 2 and bin 7 versus bin 6 to verify the inverted-U trend, and between valid and invalid for bin 7 to establish the behavioural signature of drift rate bias.

Electrophysiological data analysis. Continuous electrophysiological data were recorded in d.c. mode from 97 scalp electrodes with a sampling rate of 500 Hz and an online reference at site FCz (ActiCap, Brain Products). All offline analysis was performed using custom-generated MATLAB scripts (MathWorks)[65]. Data were first low-pass filtered by convolution with a 58-tap Hanning-windowed sinc function that was designed to provide a 3 dB corner frequency of 41.5 Hz and a local extremum of attenuation coinciding with the mains frequency (60 Hz), while also avoiding phase distortion and ringing artifacts[66]. The data were further high-pass filtered with a cut-off of 0.1 Hz to remove slow drift (third-order Butterworth). Data were epoched from −176 ms to 2,600 ms relative to prior cue onset (colour change) and baseline-corrected with respect to an interval of 118 ms centred on $t = 0$ ms (two cycles of the SSVEP). Channels with excessively high variance with respect to neighbouring channels and channels that saturated or flat-lined during a given block were identified and interpolated (spherical splines). The single-trial EEG data were then transformed to current source density (CSD)[67] to reduce the spatial blurring effects of volume conduction, in particular reducing the overlap of frontocentral and occipital negativities with centroparietal electrodes at which the CPP is measured[39,68]. Trials were rejected from the analysis if a blink was identified in the eye tracker trace, or if the maximum absolute CSD-transformed amplitude exceeded 600 μV m^{-2} for any electrode, at any time up to 150 ms after the response. This resulted in 3–29% (mean 11%) of trials being rejected.

Cue-locked (−176–1,100 ms with respect to the colour-change prior cue), evidence-locked (−59–650 ms with respect to coherence onset) and response-locked (−500–59 ms with respect to the button press) ERPs were extracted from the longer single-trial epochs. Evidence-locked and response-locked waveforms were baseline-corrected with respect to the 59 ms interval just before evidence onset. An additional offline low-pass filter (103-tap windowed sinc with 3 dB attenuation at 6 Hz and 90 dB at 17 Hz) was applied to the ERP traces to remove SSVEP oscillations before plotting. MB activity was measured using a short-time Fourier transform (STFT) applied to windows of 294 ms (5 SSVEP cycles) stepped by 59 ms (one SSVEP cycle) at a time, and by taking the mean amplitude in the range of 8–30 Hz but omitting the 17 Hz component to avoid contamination by the SSVEP. MB motor preparation indices were measured at standard 10–20 sites C3 and C4 (ref. [69]), whereas the CPP was measured as the mean amplitude in a cluster of four electrodes between standard sites Pz and CPz. The topographic foci were identified on the basis of response-locked signal topographies that were collapsed across all conditions and were therefore impartial to the effects being tested for in the study.

For the analysis of motor threshold level (Fig. 2b and Supplementary Fig. 1), MB amplitude was measured at the first STFT timepoint after −97.5 ms (the earliest motor execution potential onset in the easy regime; see below) with respect to the button click time. To generate the scatter plot, we took each of the nine task conditions (regimes × prior cues) and, after shaving off the outlying 5% on each side of the distribution, divided the distribution of MB amplitude values into five equally sized bins and plotted the mean amplitude against the mean RT for each bin. We also overlaid the averages for error trials without dividing into quantiles due to low trial count. We measured motor-level starting points for statistical testing and for model constraints using the last STFT point before evidence onset, 706 ms post-cue. To verify the pre-evidence timing of urgency onset in Fig. 2a, we performed running t-tests against zero for the temporal slope quantified as simply the consecutive two-point difference in the MB motor preparation signals and marked the timepoint at which this became significantly positive and remained so thereafter.

As MB signals represent spectral amplitudes estimated in 294 ms intervals, the temporal resolution is far too low to examine what may be subtle differences in the post-decision motor time component of non-decision time. We therefore traced the higher-resolution timecourse of broadband event-related potentials measured at the same standard motor sites (C3/C4), where, just before response completion, the slow negativity reflecting preparation gives way to a sudden positive-going deflection associated with generating the motor response[70] (Fig. 2c). Here we examined the response-contralateral signal only (standard site C3 for right button presses averaged with C4 for left button presses) for correct trials. To ensure that the suddenness of this potential was conveyed clearly, we did not apply the additional low-pass filter for this plot. In keeping with previous studies[19,25,35], we took the inflection point of this signal as the point of commitment to a decision alternative and onset of motor execution. To accurately estimate these inflection points, we computed the first timepoint at which the first derivative (slope) of the signal shifts from negative to positive, after applying the low-pass filter to reduce the impact of noise (Supplementary Fig. 2). To test for significant differences in this marker of the onset of motor execution (Fig. 2c), we performed a jackknifing procedure in which the zero-crossing time was computed for the average signals of 19 participants at a time, with each participant systematically excluded in turn, and scaled-up the s.e. of each condition by 19 before conducting an RMANOVA and follow-up comparisons[71].

Neural analyses for model validation. Pre-evidence CPP amplitude with respect to a pre-cue baseline, which was used to test for starting point effects of regime and prior cue validity, was measured by integrating across a 59 ms window ending at evidence onset (Fig. 6). Pre-response CPP amplitude (Fig. 5) was measured in a 59 ms time window centred on -97.5 ms, the earliest inflection point of the contralateral motor cortical ERP (C3/C4) estimated as described above. CPP signal slope was tested for early accumulation by fitting a line to the 118 ms period beginning at evidence onset. CPP temporal slope differences across regimes were also tested by again fitting a line in 118 ms windows, starting from 200 ms post-evidence and ending at −97.5 ms pre-response. The 118 ms (2 SSVEP cycles) window length was chosen as a compromise between the increased robustness of measurement bought by longer time windows and the need to ensure that the window does not extend beyond the period of evidence accumulation for most trials, which is very short for the deadline regime. The equivalent was performed for MB by taking the temporal difference between two consecutive STFT data points together covering a 118 ms period. As the pre-response MB signal slope would be partially driven by the pre-RT period, often overlapping with the pre-evidence period in the Deadline regime due to very fast responses, we also tested for differences in stimulus-locked slopes between the Easy and Deadline regimes and found that the Deadline motor preparation became significantly shallower at 206 ms, even when excluding any trials with a RT of faster than 300 ms to avoid potential contamination by the post-response beta rebound effect[72] ($t_{19} = -2.7798$, $P = 0.012$, CI = [−0.121,0.017]).

To test the correlation in the pattern of CPP amplitude variation over RT between the real and simulated data, we computed the correlation coefficient for correct trials across all regimes, cue validities and RT bins, and compared it against a null distribution of 10,000 correlation coefficients computed after shuffling the values, finally computing the probability of a value as extreme as the real value using the normal cumulative probability function. To test the relative timing of the stimulus-locked CPP point of acceleration, we used a jackknife procedure to compute the time at which it reached 20% of its maximum value in the easy and deadline regimes for the 20 subgroups of 19 participants and scaled-up the s.d. of estimates by $n - 1$ before conducting a t-test. Note that it is not possible to precisely estimate the true time of accumulation onset from the CPP waveforms owing to the gradual nature of initial noise accumulation, the timing jitter across trials and participant and the low-pass filtering of the CPP that is necessary to remove the SSVEP. Nevertheless, this test of the time to reach the 20% mark serves as a robust proxy and could certainly be expected to detect any large differences on the order of those estimated by the DDM. Although the differences are small, the neural data in fact trends in the direction of evidence accumulation starting sooner in the easy regime compared with the deadline regime, consistent with the NI model parameter estimates and directly opposite to the later onset suggested by the DDM.

General statistical approach. We used repeated-measures two-tailed t-tests and ANOVAs as appropriate to test for differences in behavioural and neural amplitude and slope measures across regimes, prior probability conditions and outcomes. For all error bars, between-participant variability was factored out so that only variance relevant to a repeated-measures experimental design remains. All CIs are the 95% CIs. Sphericity was verified using the sphericity function of the Python package pingouin. Bayes factors are provided for any null results (JASP v.0.11.1).

Model fitting. Models were fit to behavioural data by minimizing the χ^2-based statistic G^2 (refs. [53,73]), quantifying the divergence of model-simulated from real behavioural data grand-averaged across participants. To summarize the real data, we divided each of 18 RT distributions (3 regimes × correct/error × 3 validities) into 6 RT bins separated by the quantiles [0.1 0.3 0.5 0.7 0.9] and counted trials in each of these bins as well as the misses for each condition, expressed these counts

as proportions of overall trial counts for each participant, and then averaged both the proportions and the quantile values across participants. Those same quantiles were then used to obtain the corresponding proportions of trials in the simulated data for any given set of model parameters. The G^2 statistic was computed using a Monte-Carlo simulation of a large number of trials (see below) and summing the divergences in trial proportions between the real and simulated data using the equation

$$G^2 = 2\left(\sum_{r=1}^{3}\sum_{v=1}^{3} n_{r,v}\left[\sum_{o=1}^{2}\sum_{q=1}^{6} p_{r,v,o,q}\log\frac{p_{r,v,o,q}}{\pi_{r,v,o,q}} + p^{m}_{r,v}\log\frac{p^{m}_{r,v}}{\pi^{m}_{r,v}}\right]\right)$$

where $p_{r,v,o,q}$ and $\pi_{r,v,o,q}$ are the observed and predicted proportions of responses in bin q, bounded by the quantiles, of outcome o (correct/error) of regime r (easy/LoCoh/deadline) and validity v (valid/neutral/invalid). $n_{r,v}$ is the number of valid trials per regime and validity. G^2 characterizes goodness of fit to accuracy and RT distributions simultaneously, as the fit to the joint distributions depends on the relative choice probabilities. Models with a lower G^2 value fit the data more accurately.

To facilitate comparison of the explanatory power of the NI model with more conventional modelling procedures, we fit a standard DDM to the behavioural data. In its fullest form, the model included 18 free parameters: a separate bound b, drift rate d, drift rate bias Δd, non-decision time tnd and starting-point bias z for each of the three regimes, and a single starting-point variability (uniform distribution, range, sz), non-decision time variability (uniform, range, st) and drift rate variability (Gaussian; standard deviation, sd) parameter fixed across all of the regimes. This model corresponds to the 'full' diffusion model in common use[11], and allows for additional effects of regime aside from the dominant bound adjustment, such as the modulation of drift rate and/or non-decision time[6,28], and allows for drift rate biases in addition to starting-point biases[52]. To assess the degree to which allowing for these non-standard parameter differences was warranted to optimize the fit, we also fit three other versions of the diffusion model: one that differed from the fullest model only in disallowing any drift rate biases (15 free parameters); one that differed only in setting one drift rate per coherence (resulting in 2 rather than 3 drift rates across regimes; 17 free parameters); and one that differed only in allowing one non-decision time to cover all of the regimes (16 free parameters). As is typical, we set the within-trial Gaussian noise s.d. of $s = 0.1$ as the scaling parameter.

To guide the construction of the NI model, we first focused our analyses on the decision signal dynamics at the outset of the decision process (just before evidence onset) with respect to the levels reached at its termination. In motor preparation signals reflected in MB, we observed a stereotyped threshold level just before the response (Fig. 2b) and therefore constrained the motor preparation values across the nine conditions for the starting level at −59 ms relative to evidence onset. For convenience and intuitive interpretation without loss of generality, we scaled all MB amplitude values linearly such that the average level reached just before (first STFT point after −97.5 ms) the response for the signal contralateral to the response corresponded to a 'bound' level of 1, and the highest MB amplitude (lowest motor preparation level) during the cue-evidence interval in the neutral easy condition corresponded to zero motor preparation (Fig. 2a). A starting level in another regime of, for example, 0.5 can therefore be readily interpreted as motor preparation half-way towards execution threshold. To further constrain the motor time delay between post-threshold initiation of the cortical response execution process and the registration of the button click, we computed the time at which the first derivative (temporal slope) of the 'M1' ERP signal over motor cortex contralateral to the button pressed passed through zero, indicating the inflection point at which motor execution is assumed to begin (Fig. 2c). The motor time in the model was fixed to this time point calculated separately for each of the three regimes. There were therefore 12 fully constrained parameters in the NI model comprising nine motor preparation starting levels and three motor times. It should be noted that other signal features, such as build-up onset and build-up rate, are unsuitable for directly constraining parameters because, when measured in trial- or participant-averaged data, they do not bear a unique, one-to-one correspondence to model parameter values[57]. Instead, we used these features to validate the model by testing for its predicted evidence accumulation and motor preparation dynamics (Fig. 5). With the scaling parameter set by the bound equalling 1, we then added 17 free parameters to make the fullest NI model, including an urgency rate U (temporal slope of motor level urgency), drift rate D (mean of the Gaussian evidence distribution), accumulation onset time Tac and drift rate bias ΔD for each regime, and, in common across regimes, a single parameter for each of the factors of within trial evidence noise s, urgency rate variability Su, motor starting point variability Sz, non-decision time variability St and evidence onset time Te. Evidence onset time is the time at which the mean of the noisy evidence steps from zero to a positive value determined by the drift rate. Accumulation start time was assumed to be set independently from evidence onset, on the basis of the idea that sensory areas represent and, therefore, make available an evidence representation in a highly stereotyped, invariant manner from trial to trial, but the time at which that representation begins to be accumulated can vary both randomly and as a function of task demands[26]. Non-decision time variability was equally apportioned to the accumulation start time and the motor non-decision time as it was deemed to be probable that having separate timing variability parameters would lead to a lack of constraints. Thus, similar to the fullest DDM, the fullest NI model allowed for regime effects on drift rate and non-decision time (here specifically accumulation onset) and biases in drift rate, in addition to the usual dominant adjustments. However, a strong distinction is that the dominant adjustments are not free parameters—they are fixed by the motor preparation signal data. As for the DDM, we tested for the significant contributions of each non-dominant regime and priors effect by comparing the fit quality of versions of the model that had all of the same parameters except (1) with no drift rate biases allowed, (2) with only one urgency rate fixed across regimes, (3) with only one drift rate per coherence and (4) with only one accumulation onset time.

The fitting procedure for both the DDM and NI models followed the same general approach consisting of the following steps, each using a bounded SIMPLEX algorithm (fminsearchbnd in MATLAB[74]) for searching the parameter space. (1) We first fit a version of the model with all regime effects allowed but no drift bias from a large number of starting parameter vectors that were initialized with no regime differences, using 10,000 simulated trials per condition. (2) Starting from the parameter vectors estimated from the 25 best fits (lowest G^2) of the first step, we refined the fit of the same model by lowering the tolerance criteria for termination/convergence and simulating more trials per condition (20,000). This step resulted in an improvement in G^2 of 3–6% (reduction) with respect to the best fit of the first search. (3) Starting from the estimated parameter vectors from those 25 fits, we fit the fullest version of the model with drift rate biases free to vary from zero. For both models, we chose an initial level of drift rate bias (the same for all of the regimes) by increasing it and simulating the conditional accuracy function until a similar sized accuracy difference between valid and invalid trials was produced for the slowest RT bin of the simulated LoCoh condition. (4) Starting from those fitted parameter vectors, we tested for significant contributions of the non-dominant effects as described above, by fitting versions of the model with each individual effect nullified by reducing the number of parameters. For example, to test for the significance of non-decision time differences across regimes, we used a model version with only one non-decision time fixed across all of the regimes, and the starting parameter vector used an initial non-decision time computed from averaging the best fitting parameter values across the three regimes of the previous fit. In step 1, we selected 50 initial parameter vectors at random for the NI model but—for the DDM, to go to extra lengths to ensure the wider parameter space was comprehensively explored—we first evaluated G^2 for simulations of 1,000 trials per condition for a coarse grid of 668,250 parameter vectors widely varying bound, drift rate, non-decision time and all three intertrial variability parameters, but keeping all of the parameters fixed across regimes and starting point and drift rate biases set to zero. For each combination of the three variability parameters in this set, we took the [b, d, tnd] combination giving the lowest G^2 value and started the first coarse search from these 405 initial parameter vectors. Thus, the best 25 fits in the first step were taken from 405 preliminary fits in the case of DDM, and from 50 in the NI model. The use of a comprehensive grid search to initially identify viable starting parameter vectors, which was performed for the DDM, was not necessary to carry out for the NI model fit, because the neural data provided a natural guide for parameter ranges in which to randomly choose values. In the case of the NI model fit, this second refinement step also involved expanding from 1 accumulation onset time to 3, because it was not certain whether additional regime differences in this would be needed given that motor times were constrained. To facilitate convergence of the SIMPLEX algorithm, the same random seed was used for all runs of the simulation function for each of the steps. In a very final step for both the DDM and NI models, we computed G^2 values from the parameter estimates of each of the 25 model fits using 100,000 simulated trials per condition and a new random seed, and identified the parameter vector giving the lowest G^2 for each of the models. To guide model comparison, we computed AIC and BIC, which penalize for complexity as follows:

$$\mathrm{AIC} = G^2 + 2f$$
$$\mathrm{BIC} = G^2 + f\log(n)$$

where f is the number of free parameters and n is the mean total number of trials across participants. Simulated RT distributions in Fig. 3b,c were generated from these same simulations from the best fitting models. Although we list both AIC and BIC for reference in all fits (Table 1), we based our model selections on AIC, because BIC penalizes more severely for complexity and, as a result, tends to favour simple parameterizations at the expense of missing less dominant but reliable effects[75], which would run counter to our purposes of exploring such non-dominant effects.

We further assessed the reliability of the model fits and parameter effects in several ways. First, as the fitting procedure involved 25 separate fits starting from independent initial parameter vectors, we verified the consistency of the direction of regime effects across these fits (Extended Data Fig. 1). Second, to assess the reliability across participants, we generated 500 bootstrap samples of the 20 participants and fit the full DDM and the full NI model to each. In bootstrap sampling, samples of the same number as the original sample size (20 here) were randomly selected with replacement, meaning that a given participant can be included multiple times or not at all in any given bootstrap sample. In each case, we computed the aggregate behavioural data and the neural measurements to constrain the starting levels relative to pre-response threshold values and motor

times in the exact same manner as for the original 20 participants. For each bootstrap sample and each model, we ran two fits of the full model, starting with a parameter vector matching the final fit of the original 20 participants with and without collapsing the regime differences, and selected the fit with lower G^2. We plotted the 95% CIs of parameter effects relative to the easy regime estimated from these bootstrap samples in Fig. 4 and display the effects for all of the samples in Extended Data Fig. 2. To further ensure that there could be no one individual driving the parameter effects that we observed, we also conducted 20 additional fits of the DDM and NI model with 19 participants at a time, systematically leaving each participant out. There was no instance in which the direction of the parameter adjustments opposed that of the full 20.

To further validate the model fitting procedure itself and the results of Table 1, we generated behavioural datasets with manipulated build-up and delay effects, and verified that we obtained the expected model selection and parameter effect results. The two main findings of Table 1 were the superior fit quality of the NI model over the DDM, and the superior fit of the full NI model over model versions with any one of the four key adjustment effects disallowed. Two analyses therefore focussed on these findings. First, we established the role of the two main DDM-incompatible features of the NI model—pre-evidence accumulation and dynamic urgency—in the NI model's superior fit quality and accuracy. To this end, we generated 6 (2 × 3) datasets that had either no pre-evidence accumulation (that is, evidence-encoding onset always preceded accumulation onset) or the full amount estimated in the real data, and that had a rate of dynamic urgency build-up equal to 0, 0.5 or 1 times the amounts estimated in the real data. These datasets were simulated by manipulating certain aspects of the parameter vector estimated from the best NI model of the real data, always maintaining the same starting levels relative to bound and post-decision motor time as measured in the neural data. To make the mean RTs and error rates comparable to the real data, we scaled-up drift rates as we decreased urgency rates, so that, when there was zero urgency, the drift rates relative to the bound were the same as estimated by the DDM fits to the real data, and within-trial noise was increased until the mean squared error between the simulated and real error rates was minimized. We then ran the very same fitting procedure as used for the real data on these six simulated datasets.

Second, we validated the model selection procedure, establishing the presence of the key parameter effects of regime and prior probability by comparing models allowing each effect versus models disallowing each effect. The four principal effects were as follows: (1) non-decision time differences across regimes, (2) drift rate enhancement under deadline, (3) urgency rate differences across regimes and (4) the presence of prior probability biases on drift rate. For each of these, we took the parameter estimates of the full NI model and generated behavioural data from versions with the effect multiplied by 0, 0.5 and 1.5. Again, we fitted these 12 additional simulated datasets using the exact same fitting procedure.

Dynamic NI model simulations. To simulate the dynamics of decision formation at both the CPP and motor preparation level, we simply recorded the single-trial decision variable simulation timecourses at the CPP and MB levels for each trial and averaged them in the same way as for the real ERP data. To ensure consistency, we filtered the noisy simulated signals using the same low-pass filter in the case of the CPP and by convolving the MB signal with a boxcar function of the same duration as the FFT window used in the real data. The CPP was simulated as the absolute value of the differential cumulative evidence, as we have described previously[41]. To briefly recapitulate, this is based on the assumption that the CPP reflects the activity of two neural populations that together encode a single, signed quantity of cumulative differential evidence (Supplementary Fig. 3), one population representing positive values and the other negative values, although both do so with positive polarity. This is consistent with the observations that the CPP builds with positive polarity regardless of choice and the fact that, despite this assumed summation of two accumulator processes, the CPP reaches a stereotyped threshold level at response for urgency-free, continuous monitoring decisions[39].

Note that, without loss of generality, the model was implemented for behaviour-fitting with no assumed delay between the encoding of evidence increments at the CPP level and their impact on the motor preparation level, because it is not possible to estimate this delay from behaviour alone. The timing estimates related to evidence onset time and accumulation onset time should therefore be interpreted as the times at which these events register at the level of motor preparation, with the understanding that their upstream representation had been encoded some time before that. For the purposes of qualitative comparison between the real and simulated CPP data, we made a rough estimate of the CPP-MB transmission delay by simply observing differences in the timing of cross-over landmarks in the real and simulated data. There were four such discernible landmarks: where the stimulus-locked deadline regime trace crosses over the LoCoh and the easy regimes; and where the response-locked easy and deadline traces respectively cross over the LoCoh trace. The average temporal shift between real and simulated data was 134 ms, and we therefore shifted the CPP waveforms to the left by this amount of time in Fig. 6k,l. We emphasize that our purpose was not to pinpoint this timing delay precisely—which is probably not achievable in this manner—but, rather, to shift the traces by roughly the right amount to help with comparisons. The critical differences in timing across conditions are unaffected by this.

Joint model fit to behaviour and CPP. To use CPP amplitude modulations as additional constraints in the model fits to compare the explanatory power of NI models with and without drift rate bias (Fig. 7b), we simply re-ran the behavioural fitting procedures but with a penalty term added to the G^2 value to comprise the objective function. This penalty term was computed as the sum of squared differences between the percentage validity modulations (valid pre-response CPP amplitude with respect to the invalid condition) in the real and simulated CPP waveforms, summed across the three regimes, and weighted with a value logarithmically ranging between 0.01 and 10. Note that this particular method of effectively constraining the CPP amplitude falling out of Monte-Carlo simulations will necessarily result in worse fits to behaviour considered alone because the fitting algorithm itself is increasingly de-emphasizing the behaviour in its objective function and, as it so strongly emphasizes quantitative differences in ERP amplitudes, it will be highly sensitive to any noise at those amplitudes. The purpose of this analysis was to examine how the relative performance of the model with drift rate bias versus without drift rate bias changed with emphasis on capturing CPP amplitude modulations.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Consent was not obtained to publicly share individual participant data. However, grand average behavioural data used for modelling are available at https://osf.io/53gmw/ and other data are available from the corresponding author on request.

Code availability

All code is available publicly at https://osf.io/53gmw/.

Received: 7 September 2019; Accepted: 17 September 2020;
Published online: 14 December 2020

References

1. Smith, P. L. & Ratcliff, R. Psychology and neurobiology of simple decisions. *Trends Neurosci.* **27**, 161–168 (2004).
2. Link, S. W. & Heath, R. A. A sequential theory of psychological discrimination. *Psychometrika* **40**, 77–105 (1975).
3. Ratcliff, R. & McKoon, G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922 (2008).
4. Heitz, R. P. The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Front. Neurosci.* **8**, 150 (2014).
5. Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W. & Forstmann, B. U. Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *J. Neurosci.* **32**, 2335–2343 (2012).
6. Rae, B., Heathcote, A., Donkin, C., Averell, L. & Brown, S. The hare and the tortoise: emphasizing speed can change the evidence used to make decisions. *J. Exp. Psychol. Learn. Mem. Cogn.* **40**, 1226–1243 (2014).
7. Ratcliff, R., Smith, P. L., Brown, S. D. & McKoon, G. Diffusion decision model: current issues and history. *Trends Cogn. Sci.* **20**, 260–281 (2016).
8. Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R. & Brown, S. D. Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *J. Neurosci.* **35**, 2476–2484 (2015).
9. Evans, N. J., Hawkins, G. E. & Brown, S. D. The role of passing time in decision-making. *J. Exp. Psychol. Learn. Mem. Cogn.* https://doi.org/10.1037/xlm0000725 (2019).
10. O'Connell, R. G., Shadlen, M. N., Wong-Lin, K. & Kelly, S. P. Bridging neural and computational viewpoints on perceptual decision-making. *Trends Neurosci.* **41**, 838–852 (2018).
11. Dutilh, G. et al. The quality of response time data inference: A blinded, collaborative assessment of the validity of cognitive models. *Psychon. Bull. Rev.* https://doi.org/10.3758/s13423-017-1417-2 (2018).
12. Smith, P. L. & Lilburn, S. D. Vision for the blind: reconsidering blinded inference for decision models in light of visual psychophysics. *Psychon. Bull. Rev.* **27**, 882–910 (2020).
13. Heathcote, A. & Hayes, B. Diffusion versus linear ballistic accumulation: different models for response time with different conclusions about psychological mechanisms? *Can. J. Exp. Psychol.* **66**, 125–136 (2012).
14. Hanes, D. P. & Schall, J. D. Neural control of voluntary movement initiation. *Science* **274**, 427–430 (1996).
15. Roitman, J. D. & Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489 (2002).
16. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
17. Donner, T. H., Siegel, M., Fries, P. & Engel, A. K. Buildup of choice-predictive activity in human motor cortex during perceptual decision making. *Curr. Biol.* **19**, 1581–1585 (2009).

18. van Vugt, M. K., Simen, P., Nystrom, L., Holmes, P. & Cohen, J. D. Lateralized readiness potentials reveal properties of a neural mechanism for implementing a decision threshold. *PLoS ONE* **9**, e90943 (2014).
19. O'Connell, R. G., Dockree, P. M. & Kelly, S. P. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat. Neurosci.* **15**, 1729–1735 (2012).
20. Hanks, T. D. & Summerfield, C. Perceptual decision making in rodents, monkeys, and humans. *Neuron* **93**, 15–31 (2017).
21. Purcell, B. A. et al. Neurally constrained modeling of perceptual decision making. *Psychol. Rev.* **117**, 1113–1143 (2010).
22. Hanks, T., Kiani, R. & Shadlen, M. N. A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *eLife* **3**, e02260 (2014).
23. Churchland, A. K., Kiani, R. & Shadlen, M. N. Decision-making with multiple alternatives. *Nat. Neurosci.* **11**, 693–702 (2008).
24. Murphy, P. R., Boonstra, E. & Nieuwenhuis, S. Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nat. Commun.* **7**, 13526 (2016).
25. Steinemann, N. A., O'Connell, R. G. & Kelly, S. P. Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nat. Commun.* **9**, 3627 (2018).
26. Teichert, T., Grinband, J. & Ferrera, V. The importance of decision onset. *J. Neurophysiol.* **115**, 643–661 (2016).
27. Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U. & Nieuwenhuis, S. The neural basis of the speed-accuracy tradeoff. *Trends Neurosci.* **33**, 10–16 (2010).
28. Voss, A., Rothermund, K. & Voss, J. Interpreting the parameters of the diffusion model: an empirical validation. *Mem. Cogn.* **32**, 1206–1220 (2004).
29. Rinkenauer, G., Osman, A., Ulrich, R., Muller-Gethmann, H. & Mattes, S. On the locus of speed-accuracy trade-off in reaction time: inferences from the lateralized readiness potential. *J. Exp. Psychol. Gen.* **133**, 261–282 (2004).
30. Arnold, N. R., Bröder, A. & Bayen, U. J. Empirical validation of the diffusion model for recognition memory and a comparison of parameter-estimation methods. *Psychol. Res.* **79**, 882–898 (2015).
31. Donkin, C., Brown, S., Heathcote, A. & Wagenmakers, E.-J. Diffusion versus linear ballistic accumulation: different models but the same conclusions about psychological processes? *Psychon. Bull. Rev.* **18**, 61–69 (2011).
32. Ho, T. et al. The optimality of sensory processing during the speed-accuracy tradeoff. *J. Neurosci.* **32**, 7992–8003 (2012).
33. Heathcote, A. & Love, J. Linear deterministic accumulator models of simple choice. *Front. Psychol.* **3**, 292 (2012).
34. de Lange, F. P., Rahnev, D. A., Donner, T. H. & Lau, H. Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *J. Neurosci.* **33**, 1400–1410 (2013).
35. Dmochowski, J. P. & Norcia, A. M. Cortical components of reaction-time during perceptual decisions in humans. *PLoS ONE* **10**, e0143339 (2015).
36. Devine, C. A., Gaffney, C., Loughnane, G. M., Kelly, S. P. & O'Connell, R. G. The role of premature evidence accumulation in making difficult perceptual decisions under temporal uncertainty. *eLife* **8**, e48526 (2019).
37. Vandekerckhove, J. & Tuerlinckx, F. Diffusion model analysis with MATLAB: a DMAT primer. *Behav. Res. Methods* **40**, 61–72 (2008).
38. Kelly, S. P. & O'Connell, R. G. The neural processes underlying perceptual decision making in humans: recent progress and future directions. *J. Physiol. Paris* **109**, 27–37 (2015).
39. Kelly, S. P. & O'Connell, R. G. Internal and external influences on the rate of sensory evidence accumulation in the human brain. *J. Neurosci.* **33**, 19434–19441 (2013).
40. Twomey, D. M., Kelly, S. P. & O'Connell, R. G. Abstract and effector-selective decision signals exhibit qualitatively distinct dynamics before delayed perceptual reports. *J. Neurosci.* **36**, 7346–7352 (2016).
41. Afacan-Seref, K., Steinemann, N. A., Blangero, A. & Kelly, S. P. Dynamic interplay of value and sensory information in high-speed decision making. *Curr. Biol.* **28**, 795–802 (2018).
42. Heitz, R. P. & Schall, J. D. Neural mechanisms of speed-accuracy tradeoff. *Neuron* **76**, 616–628 (2012).
43. Serences, J. T. Value-based modulations in human visual cortex. *Neuron* **60**, 1169–1181 (2008).
44. Anderson, B. A., Laurent, P. A. & Yantis, S. Value-driven attentional capture. *Proc. Natl Acad. Sci. USA* **108**, 10367–10371 (2011).
45. Ramkumar, P., Dekleva, B., Cooler, S., Miller, L. & Kording, K. Premotor and motor cortices encode reward. *PLoS ONE* **11**, e0160851 (2016).
46. Spieser, L., Servant, M., Hasbroucq, T. & Burle, B. Beyond decision! Motor contribution to speed-accuracy trade-off in decision-making. *Psychon. Bull. Rev.* **24**, 950–956 (2017).
47. Purcell, B. A., Schall, J. D., Logan, G. D. & Palmeri, T. J. From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *J. Neurosci.* **32**, 3433–3446 (2012).
48. Coles, M. G., Gratton, G., Bashore, T. R., Eriksen, C. W. & Donchin, E. A psychophysiological investigation of the continuous flow model of human information processing. *J. Exp. Psychol. Hum. Percept. Perform.* **11**, 529–553 (1985).
49. Stanford, T. R., Shankar, S., Massoglia, D. P., Costello, M. G. & Salinas, E. Perceptual decision making in less than 30 milliseconds. *Nat. Neurosci.* **13**, 379–385 (2010).
50. Gibbon, J. Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* **84**, 279–325 (1977).
51. Jazayeri, M. & Shadlen, M. N. Temporal context calibrates interval timing. *Nat. Neurosci.* **13**, 1020–1026 (2010).
52. Dunovan, K. E., Tremel, J. J. & Wheeler, M. E. Prior probability and feature predictability interactively bias perceptual decisions. *Neuropsychologia* **61**, 210–221 (2014).
53. Ratcliff, R. & Smith, P. L. A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* **111**, 333–367 (2004).
54. de Gee, J. W. et al. Dynamic modulation of decision biases by brainstem arousal systems. *eLife* **6**, e23232 (2017).
55. Urai, A. E., de Gee, J. W., Tsetsos, K. & Donner, T. H. Choice history biases subsequent evidence accumulation. *eLife* **8**, e46331 (2019).
56. Katsimpokis, D., Hawkins, G. E. & van Maanen, L. Not all speed-accuracy trade-off manipulations have the same psychological effect. *Comput. Brain Behav.* https://doi.org/10.1007/s42113-020-00074-y (2020).
57. Purcell, B. A. & Palmeri, T. J. Relating accumulator model parameters and neural dynamics. *J. Math. Psychol.* **76**, 156–171 (2017).
58. Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
59. Niyogi, R. K. & Wong-Lin, K. Dynamic excitatory and inhibitory gain modulation can produce flexible, robust and optimal decision-making. *PLoS Comput. Biol.* **9**, e1003099 (2013).
60. Atiya, N. A. A., Rañó, I., Prasad, G. & Wong-Lin, K. A neural circuit model of decision uncertainty and change-of-mind. *Nat. Commun.* **10**, 2287 (2019).
61. Loughnane, G. M. et al. Target selection signals influence perceptual decisions by modulating the onset and rate of evidence accumulation. *Curr. Biol.* **26**, 496–502 (2016).
62. Brainard, D. H. The psychophysics toolbox. *Spat. Vis.* **10**, 433–436 (1997).
63. Wetherill, G. B. & Levitt, H. Sequential estimation of points on a psychometric function. *Br. J. Math. Stat. Psychol.* **18**, 1–10 (1965).
64. Ratcliff, R. & Tuerlinckx, F. Estimating parameters of the diffusion model: approaches to dealing with contaminant reaction times and parameter variability. *Psychon. Bull. Rev.* **9**, 438–481 (2002).
65. Delorme, A. & Makeig, S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J. Neurosci. Methods* **134**, 9–21 (2004).
66. Widmann, A. & Schröger, E. Filter effects and filter artifacts in the analysis of electrophysiological data. *Front. Psychol.* **3**, 233 (2012).
67. Kayser, J. & Tenke, C. E. Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: I. Evaluation with auditory oddball tasks. *Clin. Neurophysiol.* **117**, 348–368 (2006).
68. Twomey, D. M., Murphy, P. R., Kelly, S. P. & O'Connell, R. G. The classic P300 encodes a build-to-threshold decision variable. *Eur. J. Neurosci.* **42**, 1636–1643 (2015).
69. Gratton, G., Coles, M. G., Sirevaag, E. J., Eriksen, C. W. & Donchin, E. Pre- and poststimulus activation of response channels: a psychophysiological analysis. *J. Exp. Psychol. Hum. Percept. Perform.* **14**, 331–344 (1988).
70. Vidal, F., Grapperon, J., Bonnet, M. & Hasbroucq, T. The nature of unilateral motor commands in between-hand choice tasks as revealed by surface Laplacian estimation. *Psychophysiology* **40**, 796–805 (2003).
71. Miller, J., Ulrich, R. & Schwarz, W. Why jackknifing yields good latency estimates. *Psychophysiology* **46**, 300–312 (2009).
72. Pfurtscheller, G. & Lopes da Silva, F. H. Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clin. Neurophysiol.* **110**, 1842–1857 (1999).
73. Busemeyer, J. R. & Diederich, A. *Cognitive Modeling* (SAGE, 2010).
74. Nelder, J. A. & Mead, R. A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965).
75. Burnham, K. P. & Anderson, D. R. Multimodel inference: understanding AIC and BIC in model selection. *Sociol. Methods Res.* **33**, 261–304 (2004).
76. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).
77. Murphy, P. R., Robertson, I. H., Harty, S. & O'Connell, R. G. Neural evidence accumulation persists after choice to inform metacognitive judgments. *eLife* **4**, e11946 (2015).

Acknowledgements

We thank S. Tamang and G. Price for data collection. This work was supported by a research grant to S.P.K. and R.G.O. from the U.S. National Science Foundation (grant no. BCS-1358955), a research grant from Science Foundation Ireland to S.P.K. (grant no. 15/CDA/3591) and a European Research Council (ERC) Starting Grant to R.G.O. under the European Union's Horizon 2020 research and innovation programme (grant no. 638289). E.A.C. was supported by a Government of Ireland Postdoctoral Fellowship from the Irish Research Council (grant no. GOIPD/2017/1261) and European Union's

Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 842143. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

S.P.K. and R.G.O. conceived the research. S.P.K. collected and analysed the data. S.P.K. and E.A.C. developed and fit the computational models. S.P.K., E.A.C. and R.G.O. wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41562-020-00967-9.

Supplementary information is available for this paper at https://doi.org/10.1038/s41562-020-00967-9.

Correspondence and requests for materials should be addressed to S.P.K.

Peer review information: Primary Handling Editor: Marike Schiffer.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

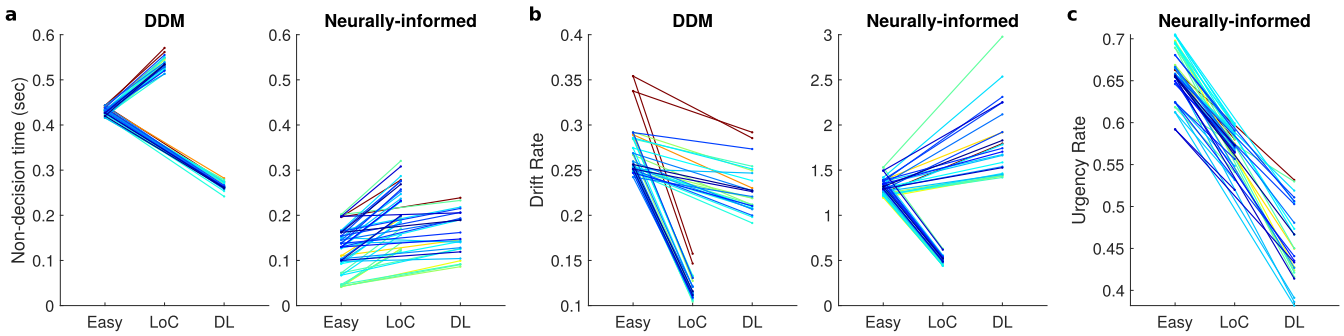

Extended Data Fig. 1 | Consistency of regime effects across independent fits. All 25 fits, even though each was started from an independent starting parameter vector originally and are of varying quality, agree on the qualitative trends across regimes for each parameter effect within each model. Bluer traces indicate lower G^2 values (better fits), whereas redder indicate higher G^2 (worse fits).

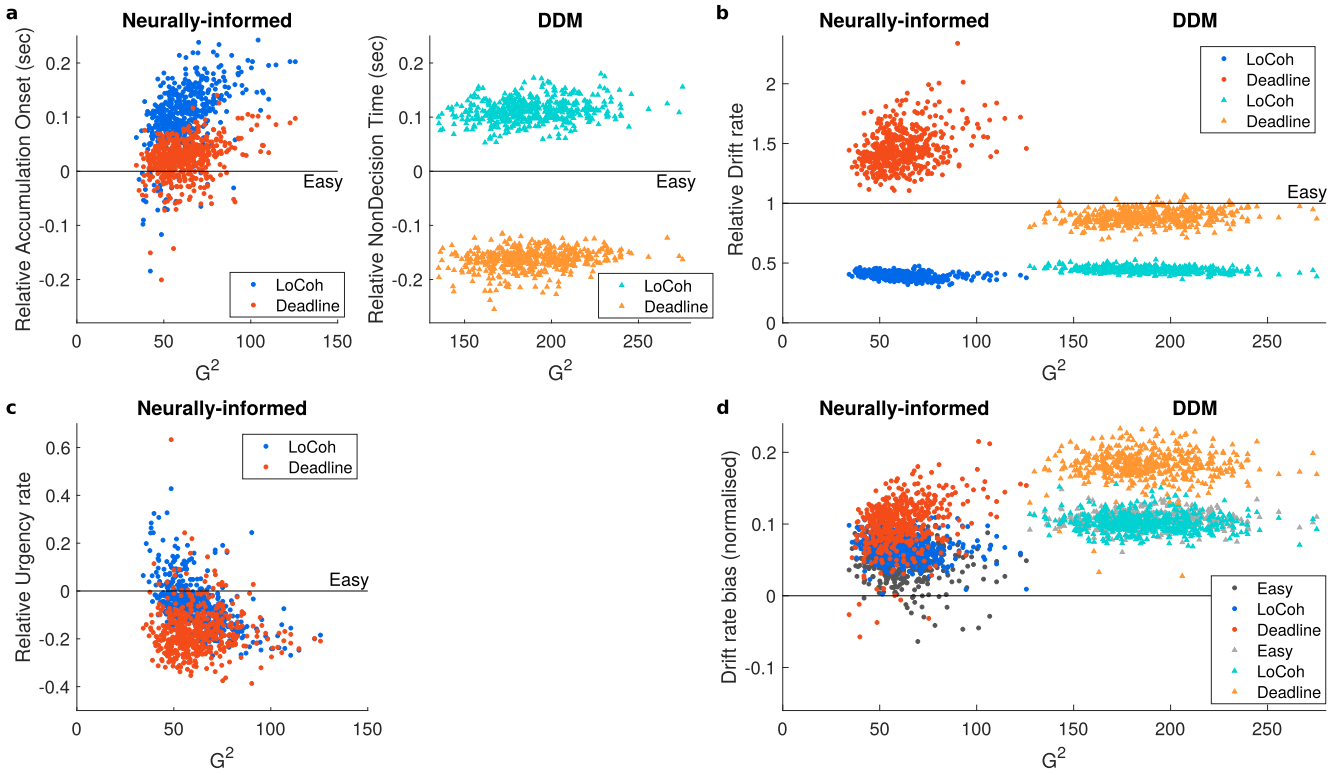

Extended Data Fig. 2 | Parameter adjustment effects across Bootstrap samples. To assess the reliability of adaptation effects indicated by the neurally-informed (NI) and drift diffusion (DDM) models, we fit the models to 500 bootstrap samples generated by sampling subjects with replacement. Scatterplots are shown of parameter effects (relative to the Easy regime or to zero as appropriate) against fit quality (G^2) along the x-axis so that reliability of the NI model's fit superiority can also be assessed. (**a**) Estimated accumulation onset parameters relative to the Easy regime in the NI model (left), and relative non-decision time for the DDM (right). (**b**) Estimated drift rate normalised to its value for the Easy regime (represented by a value of 1) to enable comparison between the DDM and NI model. (**c**) Estimated urgency rate expressed relative to the Easy regime for the NI model. (**d**) Estimated drift rate bias (positive means towards more probable alternative), normalised relative to the drift rate estimated for the Easy regime to facilitate comparison.

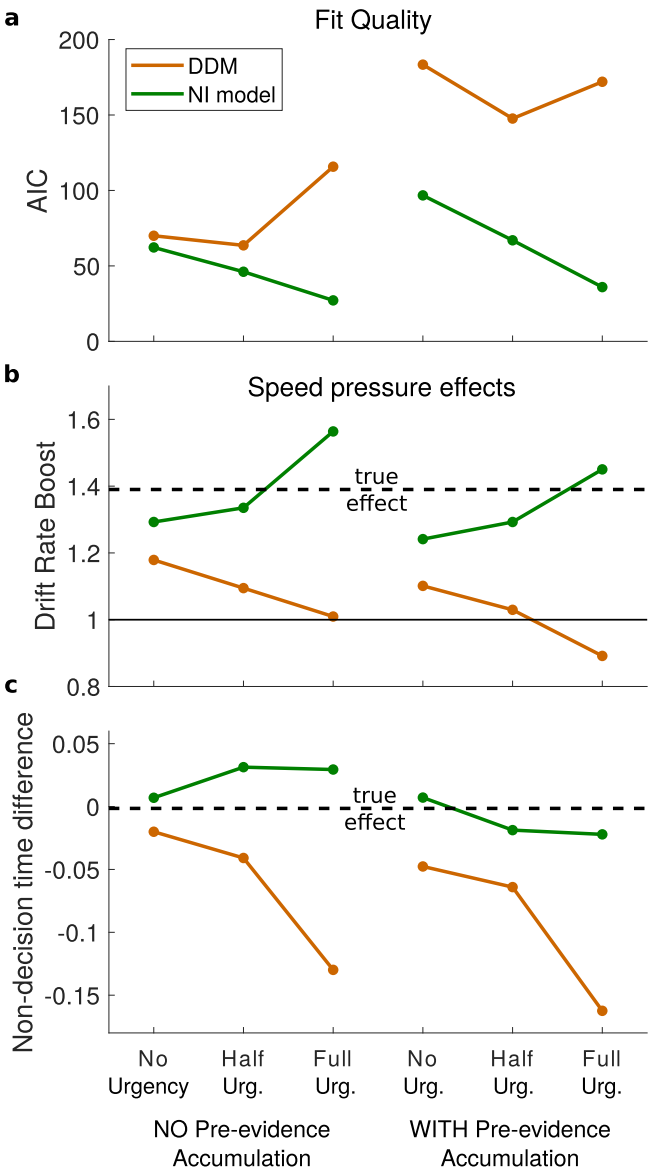

Extended Data Fig. 3 | Fit quality and key Speed pressure effects estimated by the DDM and neurally-informed (NI) models for simulated data. Six simulated datasets ranged from being fully DDM-compatible (no pre-evidence accumulation and zero dynamic urgency) to having the same degrees of pre-evidence accumulation and dynamic urgency as indicated in the fits to the real data (see methods). The DDM fits just as well as the Neurally-informed (NI) model for the most DDM-compatible data, but the superiority of the NI model - in terms of overall AIC (**a**) and in the accuracy of the relative drift rate (**b**) and non-decision time (**c**) estimated across regimes - systematically increases with the introduction of pre-evidence accumulation and dynamic urgency into the generated data. True effects computed from the parameters from which the data were generated are shown as horizontal dashed lines in (**b**) and (**c**). Drift rate (**b**) was normalised by the value in the Easy regime so that a value of 1 represents no effect. Total non-decision time for the NI model (**c**) was computed by summing the post-decision motor time and the post-stimulus accumulation onset, for comparison with the estimates of the DDM.

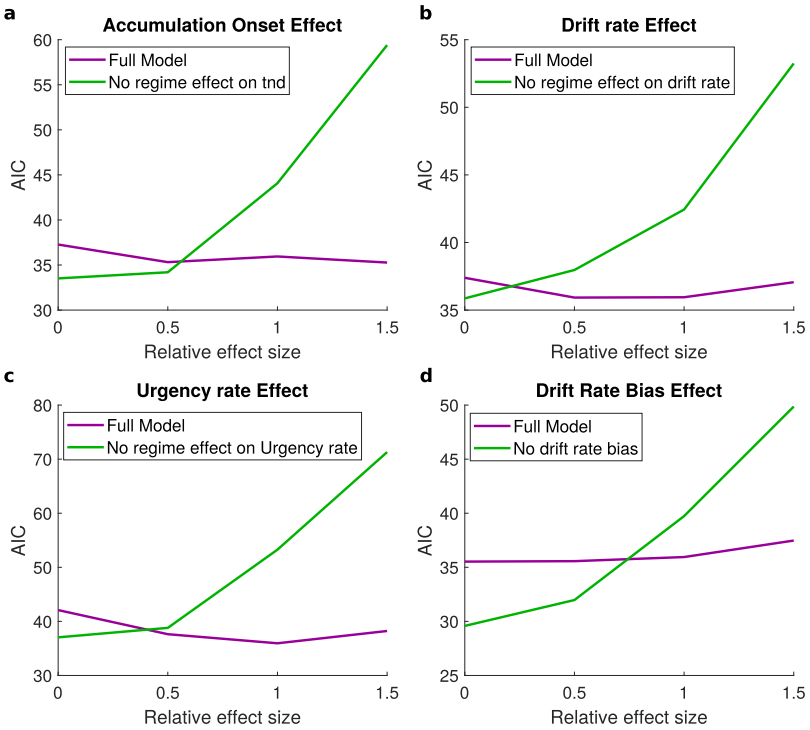

Extended Data Fig. 4 | Fit quality for simulated datasets manipulating the presence and size of each adjustment effect. For each of the 4 principal adjustment effects, **a**) non-decision time differences across regimes, **b**) drift rate enhancement under Deadline, **c**) Urgency rate differences across regimes, and **d**) the presence of Prior Probability biases on drift rate, AIC values increasingly favour the model allowing the effect as the size of the effect is increased. Importantly, in each case, the model that disallows the effect in question is favoured when the effect is fully eradicated (x 0) in the simulated data. The model selection correctly concludes in all 4 cases when the size of the effect was commensurate with the effect sizes arising from the real data (x 1).

nature research

Corresponding author(s): Simon Kelly

Last updated by author(s): Apr 16, 2020

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	We have made our task code (stimulus presentation and response recording) available along with analysis scripts and data.
Data analysis	EEGLAB version 9.0.4.5

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

Data availability

Consent was unfortunately not obtained to share individual subject data publicly. However, grand average behavioural data used for modelling is available at https://osf.io/53gmw/ and other data is available upon reasonable request.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	A sample-size of twenty-two subjects was arrived at based on having measured an effect size of more than 0.6 in oscillatory signatures of bias in a prior study (Kelly et al 2010, European J Neurosci). However, those were biases due to spatial attention, not motor biases as in the current study, and furthermore, there were many more dependent measures here than just motor-level biases. This therefore served as an approximate guide, but is a sample size that is comfortably high compared to similar repeated-measures EEG studies in the literature.
Data exclusions	Two subjects were excluded from all analyses due to excessive perspiration artifacts in at least one EEG session.
Replication	This experiment itself replicates many of the key findings of a prior study on an independent dataset and different task (Steinemann et al 2018): motor preparation was initially elevated under speed pressure but reached a constant level at the time of response, centro-parietal amplitude at response varied systematically with response time and buildup rate was higher under speed pressure. The findings have not yet been replicated for this specific task beyond the current dataset, however, nor have the modelling results here yet been replicated.
Randomization	In this repeated-measures design, we pseudorandomised the order of the conditions.
Blinding	In our study it was important for the subjects to know what condition they were performing, because we were examining their strategic adaptation to those conditions. Therefore, blinding would have been inappropriate in this case.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☐	☒ Human research participants
☒	☐ Clinical data
☒	☐ Dual use research of concern

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Human research participants

Policy information about studies involving human research participants

Population characteristics	There are no covariates as this is a repeated-measures design
Recruitment	Participants were recruited through advertising around the campus of City College New York. Therefore it was random selection but the sample contains mostly students in their 20s.
Ethics oversight	City College of New York IRB

Note that full information on the approval of the study protocol must also be provided in the manuscript.