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]

[1] Department of Biomedical Engineering, City College of the City University of New York, New York, New York 10031.

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

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

* Corresponding author: Simon P. Kelly (simon.kelly@ucd.ie); ORCID 0000-0003-4981-3099

Abstract

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 difficult to establish using simplified models based on behaviour alone. Here we exploited well-characterised human neurophysiological signatures of decision formation to construct and constrain a build-to-threshold decision model with multiple buildup (evidence accumulation, urgency) and delay components (pre- and post-decisional). The model indicated that all of these components were adapted in distinct ways, in several instances fundamentally differing from the conclusions of conventional diffusion modelling. Critically, 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.

Introduction

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 under time pressure, and we can prioritise 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 underlying such adaptations[3–5], yet our understanding remains very incomplete. Since traditionally these models have 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 which 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 very difficult through 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 exhibiting dynamics consistent with bounded evidence accumulation have now been observed in multiple species, especially in activity associated with the planning of decision-reporting actions [14–20]. These neural signatures can bolster model-based inferences in several complementary ways. Most obviously, 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 - 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, allow 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 exploited to examine human perceptual decisions. Here, we use recently characterised human electroencephalographic (EEG) signatures of decision formation to construct, constrain and validate a bounded accumulation model, which 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 buildup - 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 encapsulating all residual time not spent on bounded evidence accumulation, the NI model isolates three distinct temporal components of evidence encoding onset, accumulation onset [26], and post-decision motor-execution time.

Our results revealed multiple distinct computational adaptations among these buildup 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 dramatically shortened non-decision time and reduced drift rate similar to other recent studies [6,11,28–33]; in contrast, the NI model indicated only very slight changes to non-decision time and an enhancement of drift rate that acts to partially recoup the accuracy losses incurred by lower bounds. Dynamic urgency, whose role has been controversial in the behavioural modeling literature [7], was strongly observed in all regimes and strategically adjusted to increase the 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 bias’), which the DDM failed to identify through standard model comparison. Critically, neural decision signal features that were not incorporated into the NI model's construction 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, which respectively involved low demands (‘Easy’ - high coherence and long deadline), weak evidence (‘LoCoh’ - low

coherence and long deadline), and speed pressure ('Deadline' - high coherence and short deadline). In each trial, initially grey, incoherently moving dots changed color to indicate whether the direction of upcoming motion would be more likely leftward, rightward or equally likely (Fig. 1a), and after a fixed, predictable delay suddenly transitioned to coherent motion. Subjects were asked to respond with their left hand for leftward motion or with their right hand for rightward motion, in order to earn points for correct decisions reported within the deadline.

Figure 1 here

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 (where motion direction was congruent with cue colour) had greater accuracy, faster correct responses and slower error responses than invalid trials, and these biases were stronger in both the LoCoh and Deadline regimes than Easy. In the remainder of the Results section, we first lay out the neurally-informed (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, neurally-informed elaborations including dynamic urgency and distinct pre/post decision delays. To allow 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 via Mu/Beta-band spectral amplitude ("MB" signal, integrated over 8-30 Hz [17,34]) over right- and left-hemisphere motor areas, respectively. Plotting the MB amplitude just prior to response across all 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; see Supplementary Fig. 1 for ipsilateral signal amplitudes). Accordingly, we set a constant bound on the NI model, and to enable a straightforward linkage to the empirical motor preparation signals, cast the decision process as a race between two parallel buildup signals to reach that bound [16] (Fig. 2d).

Motor-level Starting Points: We next examined the neural dynamics of anticipatory motor preparation prior to evidence onset. Replicating previous studies [4,27,34], pre-evidence motor preparation levels were relatively elevated under speed pressure (main effect of Regime in 3x3 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 long-established in standard models (Fig. 2a). More strikingly, in all 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 buildup 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; see Methods). We found that this onset significantly differed as a function of Regime ($F(2,38)=20.56$, $p<0.001$, partial $\eta^2=0.52$; see Supplementary Table 4 for pairwise comparisons) but not Cue Validity ($F(2,38)=0.04$, $p=0.96$, partial $\eta^2=0.002$, BFexcl=24.24). We thus constrained the motor time in our model to these onset latencies per regime (-97.5, -82, -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 likely 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 thus included a single evidence onset parameter and three independently adjustable accumulation onset parameters for the three regimes.

Figure 2 here

Behavioural Modelling Results

We fit the neurally-informed (NI) model developed above to behaviour and compared its outcomes to a standard DDM [37] with a similar number of free parameters (See Supplementary Tables 1 and 2 for full parameter listings). An obvious 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).

Figure 3 here

The more important question for our current purposes, is what these models indicate regarding the key parameter adjustments that are made to adapt to speed/perceptual demands and prior probability. Since models were fit to grand average data (see Supplementary Tables 1 and 2 for point estimates of all parameters), we assessed reliability of parameter effects across subjects by fitting the DDM and NI models to 500 bootstrap samples of the 20 subjects and computing 95% confidence intervals (CI) of differences with respect to Easy/Neutral trials (see methods). As expected, the DDM indicated criterion shifts that agreed with those reflected in neural motor preparation starting-level modulations, with the bound separation being approximately halved in the Deadline regime (0.070 Vs 0.037; difference CI = [0.030 0.036]) and biases toward the more probable alternative in all 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, with both indicating the expected drift rate reduction in the LoCoh regime (LoCoh drift rate as a proportion of Easy for DDM: 0.45, CI = [0.41 0.48]; for NI model: 0.39, CI = [0.35 0.44]), and a lengthening of non-decision time under weaker evidence in the DDM (0.537s vs 0.424s, difference CI = [0.077 0.148]) mirrored by a similar lengthening of the pre-accumulation delay in the NI model (0.269 s vs 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 conclusions regarding buildup 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 relative to the Easy regime (as a proportion: 0.88, CI = [0.80 0.98]) and shortened non-decision time (by 0.162 s, CI = [0.134 0.197]). In stark 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 vs 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's Information Criterion (AIC) values increased substantially for alternate 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 regimes (Supplementary Tables 1 and 2; Extended 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 7th RT bin of 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 (See 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 DDM's improvement in fit with the inclusion of drift rate bias was very slight and did not overcome the AIC's complexity penalty. 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 fit the models to simulated data with manipulated buildup 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, left), 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. Further, 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).

Figure 4 here

Model validation through independent neural data

We next tested the NI model's validity via 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 known as the centro-parietal positivity ("CPP") [19,38]. Like 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 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 for comparison with 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]), while the CPP was assumed to reflect only the cumulative evidence, in rectified form [41] (see Supplementary Fig. 3 for a detailed model implementation). Qualitatively, the real (Fig. 5a,b,d,e) and simulated (Fig. 5h,i,k,l) signals appear very well-matched both in their initial post-stimulus buildup dynamics and their final dynamics and amplitudes just prior to the response.

Figure 5 here

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 thus needs little cumulative evidence to reach it; beyond those earliest RTs, the cumulative evidence required to trigger a response declines with RT because of 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 x 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). Further,

as in our previous work [25], there was no effect of Prior Cue or Regime on the CPP measured at evidence onset (Fig. 6; 3x2 rmANOVA; see Supplementary Table 7), providing additional confirmation that the CPP does not directly receive an urgency input.

Figure 6 here

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 2x3 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, whereas the NI model did not indicate a significant difference in accumulation onset between the Easy and Deadline 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, BF01=4.3).

The NI model indicated an enhanced drift rate in the Deadline regime, opposite to the reduction effect indicated by the DDM. Buildup rates of the CPP signal accorded with the NI model (Stimulus-locked slope: F(2,38)=21.5, p<0.001, partial η^2=0.53; Response-locked: F(2,38)=3.97, p=0.027, partial η^2=0.17; slope enhancement under the Deadline regime reached significance only for the response-locked waveforms, both when measured in time windows aligned to the response, and aligned to the regime-specific motor times; pairwise comparisons

in Supplementary Table 9; Fig. 5d,e). Note that more points were awarded 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 to Easy, but this did not significantly relate to the drift rate enhancement (see Supplementary **Fig. 7** and discussion).

The NI model also indicated a reduced rate of urgency buildup in the difficult regimes, particularly the Deadline regime. The neural signatures of motor preparation (Fig 5a,b) exhibited buildup rates that ranked in the same way as the model-estimated urgency rates, with the Deadline condition exhibiting the shallowest slope prior to 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 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. The NI model, in contrast, 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 3x3 rmANOVA Regime x 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 for evidence accumulation there is. Why would these amplitude modulations manifest in the Deadline regime but not LoCoh? As illustrated in Fig. 7a, due to the growing dynamic urgency component, the negative influence of a starting point increase on 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 including 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).

Figure 7 here

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 underlying 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-characterised 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 allowed the model to be evaluated not only based on quantitative fits to behaviour but also based on its ability to explain dynamic neural signal features not used in the model's construction.

While behavioural modelling studies have tended to converge on boundary adjustments as the principal mediator of speed-accuracy tradeoffs, some recent studies have suggested that speed pressure may also shorten non-decision time and reduce drift rate[6,11,28–33]. While our DDM fit replicated these effects, our neurally-informed model and neural decision signals provided strong convergent evidence for an enhancement of drift rate, not a reduction. This, in fact, accords 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 [25,44]. Our subjects earned 12% more points overall in the Deadline compared to the Easy regime, and although the much greater regularity of negative feedback makes it unlikely that subjects 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 their experiencing higher overall rewards [45–47]. 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, 3) other elements of processing were not enhanced (e.g. reduced urgency rate). Thus, rather than an automatic reward-related

enhancement, the drift rate boost likely 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 so doing, the model revealed that they are each independently adjustable. Consistent with recent electromyography findings [25,48], 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, while accumulation onset was lengthened, presumably because the longer deadline allows postponement of accumulation until evidence is strongly encoded, in line with recent proposals of strategic accumulation onset adjustments[26,36,49].

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 a dynamic urgency component to motor preparation, which implements a collapsing bound. Meanwhile, most behavioural model comparison studies have concluded in favour of constant bounds, though it has often been difficult to account for the increased complexity, and hence penalisation, 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 work [50] and demonstrated in neuronal activity in monkey frontal eye fields during compelled response tasks where the 'go' command precedes the target stimulus [51]. Our NI model revealed that dynamic urgency played an important role in all three regimes. Interestingly, its buildup rate was slowest in the Deadline regime, but this was coupled with a much earlier onset from an already elevated 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 may relate to the narrow time window for rewarded responses in this regime. Given that the variability of temporal estimation scales with the duration being estimated [52,53], and assuming the variability of urgency buildup rate may similarly scale with its mean, it may be that earlier and shallower launches of urgency can be aimed more precisely. While further work is needed to fully establish the principles governing 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 modeling work [3,28,54], 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 their inclusion is justified against the increased complexity only for the NI model. We were able to confidently conclude in favour of the bias based on their ability to account for CPP amplitude modulations. This finding complements another recent study showing strong drift rate biases due to relative reward value in severely time-constrained sensorimotor decisions [41], and adds to evidence of drift rate biases arising from other sources such as previous exposure in recognition memory tasks [55], biases towards 'no' responses [56] and choice history [57].

In addition to providing more accurate inferences regarding parameter adjustments, the NI model's goodness of fit 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 is owed to the presence of early accumulation onsets (i.e., preceding evidence encoding) and dynamic, anticipatory urgency (Extended Data Fig. 3). These features are likely induced by the imposition of response deadlines, and bring about a reduced rightward skew in RT distributions and many purely prior-based fast responses in the Deadline regime, which 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 only be driven by evidence that favours the correct alternative. Our NI model allows for the timing of multiple stages of the decision process to be independently estimated, thus 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 peculiar feature of our task is that the coherent motion did not suddenly appear against a blank field, but rather transitioned from incoherent motion in order to avoid contaminating evoked potential components [19,39]. This in itself is unlikely to have been a significant 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 (e.g. with no deadlines [58] 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 significantly 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 of decision formation. While 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 buildup rate bear a less direct correspondence and are better analysed with reference to dynamic model simulations [59]. Such simulations were instrumental here. For example, while 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. Though neurally-informed, 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 [60–62]. As applied here, our NI model solely relies on well-established functional characteristics of the neural decision signals, but much research is still needed to expose the neural architecture giving rise to these signals and their properties in order for the EEG-informed models to further contribute to our understanding of the finer details of the neural systems implementing decision computations. Furthermore, while neural constraints enable 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]. Lastly, the neurally-informed modeling approach we applied here dealt with aggregate grand-average data to ensure reliability of neural constraints, but future work will have to explore the degree to which single-subject EEG can be used to model individual differences, which would be particularly valuable in clinical research.

In sum, 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 subjects (8 female) aged between 20 and 32 years each participated in 3 experimental sessions, the first for psychophysical training and the remaining two for EEG recordings during task performance. No statistical methods were used to pre-determine sample size but it is larger than that of previous studies reporting Speed/Accuracy manipulation effects [25]. Two subjects were excluded from all analyses due to excessive perspiration artifacts in at least one EEG session. All subjects had normal or corrected-to-normal vision. Informed written consent was obtained from all subjects, and all procedures were approved by the Institutional Review Board at The City College of New York.

Setup. Subjects 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 1024 × 768) with a black background. They rested their hands on the table in front of them, with the forefingers of their left and right hand 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 with a remote eye tracker (EyeLink 1000, SR Research, 1000Hz).

Task. A trial was initiated by the subject by simultaneously pressing both mouse buttons, provided their measured eye position was within 3 degrees of a central, white, 5 x 5-pixel fixation point. After a short delay of 50 ms, gray, incoherently moving dots appeared, changed color after 647 ms (11 flicker frames), and finally became coherent after a further 764 ms (13 flicker frames, see Fig. 1a). The fixed timing meant that subjects 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 (see also [39,40,63]), while nevertheless producing a sudden-onset of evidence as in other paradigms. Cue colors were yellow, green and cyan, assigned to indicate a 75% prior probability of leftward motion, neutral (50/50) and 75% rightward with the color-meanings fixed for a given subject but counterbalanced across subjects. The three prior-probability colors were each equiluminant with

the initial light gray color of the dots, as verified using a photometer. The dots remained at a constant coherence level for 1600 ms before the display was extinguished for all conditions. If the subject 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.” Subjects were instructed to maintain fixation for the duration of each trial, until feedback was presented. Subjects 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 (1600 ms); 2) ‘LoCoh,’ which had low coherence (ranging from 5 to 12 %, individually titrated; see below), and a long deadline (1600 ms), and 3) ‘Deadline,’ which had high coherence (20%) and a short deadline (ranging from 388 to 640 ms, individually titrated).

The dot-motion stimulus was composed of an average of 118, 6 x 6 pixel dots presented within an aperture of 8 degrees diameter (density 2.3 dots/degree squared) centered on fixation, flickering on and off at a frame rate of 17 Hz (two refreshes on, three off; see also [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 degrees in either the leftward or rightward direction on the following frame, resulting in a motion speed of 6 degrees per second. The task was programmed using PsychToolbox [64] for Matlab (Mathworks, MA).

Training. In order to practice the task and to determine individually-titrated levels of difficulty in the Deadline and LoCoh regimes, subjects first underwent a training session. In this session, a series of short blocks of 24 trials were first run with a 1600 ms deadline, starting at 70% coherence, and stepping down by 10-20% upon achieving >90% correct responses at each level, until the subject 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, [65] was then run to estimate the coherence level at which the subject performed at 80% accuracy with the same 1600-ms deadline, and this was repeated until a stable level was indicated. Next, the deadline was changed from 1600 ms to the median RT from the last 20%-coherence block

and 48 trials were run to estimate the subject's typical RT distribution under time pressure. The deadline was then further shortened toward 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 color cues were introduced. The subject 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 subject performed 10 blocks of 96 trials of each of the three regimes (two subjects 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 subjects would settle into a strategy and to minimize cross-over. The orders of the regimes were randomized across sessions and subjects. 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 appeared too easy or difficult. To minimize the degree to which blocks of the different regimes differed in overall value, subjects 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 twenty subjects, the average reward per trial turned out to be 40.6, 40.4 and 45.5 points for the Easy, LoCoh and Deadline regimes, respectively, which, though small, was a significant difference (1-way rmANOVA: $F(2, 38) = 8.43$, $p = 0.007$, partial $\eta^2=0.31$; post-hoc differences significant for Deadline compared to 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 subjects) and to 50 for 27 of them. One point corresponded to $0.0031, and at the end of the experiment two blocks were randomly selected within each regime, and the total points of those 6 blocks translated into a dollar amount in compensation for the subjects' time, in addition to $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 prior to evidence onset. Such negative RTs were regarded as a legitimate manifestation of the mechanisms employed to perform the task, rather than as contaminant responses [66]. Similarly, responses occurring after the deadline in the Deadline regime were included in all analyses and modeling. Trials on which no response was made by 1600 ms were counted as misses. These trials are shown in the behavioural data (Fig. 3a) but excluded from electrophysiological analysis because they cannot be response-locked. In repeated measures ANOVAs 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 trials into non-overlapping RT bins of width 35.3 ms starting from -176 ms and ending at 1600 ms) and separately counted correct trials and error responses falling into each RT bin. This was done for each individual subject, expressed as a proportion of the total trials for that subject, then averaged across subjects. To compute conditional accuracy functions (CAF), we first divided each subject's pooled (correct and error) RT distribution into 7 equal sized bins, then computed the proportion of correct trials within each of these RT bins and finally plotted response accuracy over the bins' mean RTs for each condition, averaged over subjects (Fig. 3a). A 3 regime x 3 validity x 7 RT-bin repeated-measures ANOVA (rmANOVA) was run on conditional accuracy functions and the 3-way interaction was followed up with pairwise tests of bin 1 vs 2 and bin 7 vs 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 DC 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 in-house Matlab scripts (MathWorks, Natick, MA) with raw data-reading, channel interpolation and topographic plot functions from the EEGLAB toolbox [67]. In offline analysis, continuous data were first low-pass filtered by convolution with a 58-tap hanning-windowed sinc function 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 [68]. The data were further high-pass filtered with a cutoff of 0.1 Hz to remove slow drift (3rd-order Butterworth). Data were epoched from -176 ms to 2600 ms relative to prior cue onset (color change) and baseline-corrected with respect to a 118-ms interval centered on t = 0 ms (2 cycles of the SSVEP). Channels with excessively high variance with respect to neighboring 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 [69] to reduce the spatial blurring effects of volume conduction, in particular reducing the overlap of frontocentral and occipital negativities with centroparietal electrodes where the CPP is measured [39,70]. Trials were rejected from analysis if a blink was identified in the eye tracker trace, or if the maximum absolute CSD-transformed amplitude exceeded 600 μV/m2 for any electrode, at any time up to 150 ms following the response. This resulted in 3-29% (mean 11%) of trials being rejected.

Cue-locked (-176 to 1100 ms with respect to the color-change prior cue), evidence-locked (-59 to 650 ms with respect to coherence onset) and response-locked (-500 to 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 prior to evidence onset. An additional offline low-pass filter (103-tap windowed sinc with 3dB attenuation at 6 Hz and 90 dB at 17 Hz) was applied to the ERP traces to remove SSVEP oscillations prior to plotting. Mu/Beta-band (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 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 [71], while the CPP was measured as the mean amplitude in a cluster of 4 electrodes between standard sites Pz and CPz. The topographic foci were identified based on response-locked signal topographies that were collapsed across all conditions and thus impartial to the effects being tested for in the study.

For the analysis of motor threshold level (Fig. 2b, Supplementary Fig. 1), MB amplitude was measured at the first STFT timepoint following -97.5 ms (the earliest motor execution potential onset in the Easy regime, see below) with respect to the button click time. To make the scatter plot we took each of the 9 task conditions (regimes x prior cues) and after shaving off the outlying 5% on each side of the distribution, divided the distribution of MB amplitude values into 5 equal 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 prior to evidence onset, 706 ms post-cue. To verify the pre-evidence timing of urgency onset in Fig. 2a, we carried out running t-tests against zero for the temporal slope quantified as simply the consecutive two-point difference of the MB motor preparation signals and marked the timepoint at which this became significantly positive and remained so thereafter.

Since 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 (Fig. 2c; [72]). 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 the suddenness of this potential is conveyed clearly we did not apply the additional low-pass filter for this plot. In keeping with previous studies [19,25,35], we take 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 carried out a jackknifing procedure where the zero-crossing time was computed for the average signals of 19 subjects at a time, with each subject systematically excluded in turn, and scaled-up the standard errors of each condition by 19 before conducting an rmANOVA and follow-up comparisons (see [73]).

Neural analyses for model validation. Pre-evidence CPP amplitude with respect to a pre-cue baseline, 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 centered on -97.5 ms, the earliest inflection point of the contralateral motor cortical ERP (C3/C4) estimated as 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 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 carried out for MB by taking the temporal difference between two consecutive STFT data points together covering a 118-ms period. Since the pre-response MB signal slope would be partially driven by the pre-response

time 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 RT faster than 300 ms to avoid potential contamination by the post-response beta rebound effect[74] (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 regime for the 20 subgroups of 19 subjects and scaled-up the standard deviation of estimates by N-1 before conducting a t-test. Note that the true time of accumulation onset is not possible to precisely estimate from the CPP waveforms due to the gradual nature of initial noise accumulation, the timing jitter across trials and subjects, and the low-pass filtering of the CPP necessary to remove the SSVEP (see methods). 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. Though the differences are slight, the neural data in fact trends in the direction of evidence accumulation starting sooner in the Easy regime than the Deadline regime, consistent with the NI model parameter estimates and directly opposite to the later onset suggested by the DDM.

General statistical approach. Repeated-measures, two-tailed t-tests and ANOVAs were employed 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-subject variability has been factored out so that only variance relevant to a repeated-measures experimental design remains. All confidence intervals (CI) are 95% confidence intervals. Sphericity was verified using the 'sphericity' function of Python package 'pingouin.' Bayes factors are provided for any null results (JASP 0.11.1).

Model fitting. Models were fit to behavioural data by minimising the chi-square-based statistic G^2 [55,75] quantifying the divergence of model-simulated from real behavioural data grand-averaged across subjects. To summarise the real data, we divided each of 18 RT distributions (3 regimes x correct/error x 3 validities) into 6 RT bins separated by the quantiles [.1 .3 .5 .7 .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 subject, and then averaged both the proportions and the quantile values across subjects. 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 by 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), respectively. $n_{r,v}$ is the number of valid trials per regime and validity. G^2 characterises 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 neurally-informed model with more conventional modelling procedures, we fit a standard drift diffusion model (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 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 [54]. To assess the degree to which allowing for these non-standard parameter differences was warranted to optimise 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 regimes (16 free parameters). As is typical, we set the within-trial Gaussian noise standard deviation of s=0.1 as the scaling parameter.

To guide the construction of the neurally-informed model, we first focused our analyses on the decision signal dynamics at the outset of the decision process (just prior to 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 prior to 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 prior (first STFT point following -97.5 ms) to 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, say, 0.5 can thus be readily interpreted as motor preparation half-way towards execution threshold. In order 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. Thus, there were 12 fully constrained parameters in the neurally-informed model comprising nine motor preparation starting levels and three motor times. It should be noted that other signal features such as buildup onset and buildup rate are unsuitable for directly constraining parameters because, when measured in trial- or subject-averaged data, they do not bear a unique, one-to-one correspondence to model parameter values [59]. Instead, we use 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 equaling 1, we then added 17 free parameters to make the fullest neurally-informed 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 Δ**D** 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, based on the idea that sensory areas represent and thus 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 likely that having separate timing variability parameters would lead to a lack of constraints. Thus, like 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 in 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, respectively, 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 [76]) 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 initialised 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 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 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 inter-trial variability parameters but keeping all 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 done for the DDM, was not necessary to carry out for the neurally-informed 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 additionally 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 Akaike's (AIC) and Bayes Information Criterion (BIC), which penalise for complexity as follows,

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

$$BIC = G^2 + f \log(n)$$

where f is the number of free parameters and n is the mean total number of trials across subjects. Simulated RT distributions in Fig. 3B,C are 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 penalises more severely for complexity and as a result tends to favour simple parameterizations at the expense of missing less

dominant but reliable effects [77], 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, since 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 reliability across subjects, we generated 500 bootstrap samples of the 20 subjects 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) are randomly selected with replacement, meaning that a given subject 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 way as for the original 20 subjects. 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 subjects with and without collapsing the regime differences, and selected the fit with lower G^2. We plot the 95% confidence intervals of parameter effects relative to the Easy regime estimated from these bootstrap samples in Fig. 4 and display the effects for all 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 subjects at a time, systematically leaving each subject out. There was no instance where 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 buildup 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 4 key adjustment effects disallowed. Two analyses thus 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 x 3) datasets that had either no pre-evidence accumulation (i.e. evidence encoding onset always preceded accumulation onset) or the full amount estimated in the real data, and that had a rate of dynamic urgency buildup 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 minimised. We then ran the very same fitting procedure as used for the real data on these 6 simulated datasets.

Second, we validated the model selection procedure establishing the presence of the key parameter effects of Regime and Prior Probability through comparison of models allowing versus disallowing each effect. The 4 principal effects were 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. 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 x 0, x 0.5 and x 1.5. Again, we fit 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 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 done in previous work [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 (see Supplementary Fig. 3), one population representing positive values and the other negative values, though both do so with positive polarity. This accords 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 thus 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 4 such discernible landmarks: where the stimulus-locked Deadline regime trace crosses over the LoCoh and the Easy regimes, respectively, 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 therefore we shifted the CPP waveforms leftwards by this amount of time in Fig. 6k,l. We emphasise that our purpose was not to pinpoint this timing delay precisely, which is unlikely achievable in this way, but rather to shift the traces by roughly the right amount to aid comparisons. The critical differences in timing across conditions are unaffected by this.

Joint model fit to behaviour and CPP. In order to use CPP amplitude modulations as additional constraints in the model fits to compare the explanatory power of neurally-informed 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-emphasising the behaviour in its objective function, and since it so strongly emphasises quantitative differences in ERP amplitudes it will be highly sensitive to any noise in those amplitudes. The purpose of this analysis was to examine how the relative performance of the model with versus without drift rate bias changed with emphasis on capturing CPP amplitude modulations.

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 from the corresponding author upon request.

Code availability

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

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 (Regul Ed)* **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.* (2019). doi:10.1037/xlm0000725
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.* (2018). doi:10.3758/s13423-017-1417-2
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* (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,** (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. Cognit.* **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,** (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.e6 (2018).
42. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461,** 263–266 (2009).

43. Murphy, P. R., Robertson, I. H., Harty, S. & O'Connell, R. G. Neural evidence accumulation persists after choice to inform metacognitive judgments. *elife* **4,** (2015).
44. Heitz, R. P. & Schall, J. D. Neural mechanisms of speed-accuracy tradeoff. *Neuron* **76,** 616–628 (2012).
45. Serences, J. T. Value-based modulations in human visual cortex. *Neuron* **60,** 1169–1181 (2008).
46. Anderson, B. A., Laurent, P. A. & Yantis, S. Value-driven attentional capture. *Proc Natl Acad Sci USA* **108,** 10367–10371 (2011).
47. Ramkumar, P., Dekleva, B., Cooler, S., Miller, L. & Kording, K. Premotor and motor cortices encode reward. *PLoS ONE* **11,** e0160851 (2016).
48. 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).
49. 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).
50. 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).
51. 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).
52. Gibbon, J. Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* **84,** 279–325 (1977).
53. Jazayeri, M. & Shadlen, M. N. Temporal context calibrates interval timing. *Nat. Neurosci.* **13,** 1020–1026 (2010).
54. Dunovan, K. E., Tremel, J. J. & Wheeler, M. E. Prior probability and feature predictability interactively bias perceptual decisions. *Neuropsychologia* **61,** 210–221 (2014).
55. Ratcliff, R. & Smith, P. L. A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* **111,** 333–367 (2004).
56. de Gee, J. W. *et al.* Dynamic modulation of decision biases by brainstem arousal systems. *elife* **6,** (2017).
57. Urai, A. E., de Gee, J. W., Tsetsos, K. & Donner, T. H. Choice history biases subsequent evidence accumulation. *elife* **8,** (2019).
58. Katsimpokis, D., Hawkins, G. E. & van Maanen, L. Not all Speed-Accuracy Trade-Off Manipulations Have the Same Psychological Effect. *Comput. Brain Behav.* (2020).

doi:10.1007/s42113-020-00074-y

59. Purcell, B. A. & Palmeri, T. J. Relating accumulator model parameters and neural dynamics. *J. Math. Psychol.* **76,** 156–171 (2017).
60. Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36,** 955–968 (2002).
61. 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).
62. 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).
63. 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).
64. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10,** 433–436 (1997).
65. Wetherill, G. B. & Levitt, H. Sequential estimation of points on a psychometric function. *Br. J. Math. Stat. Psychol.* **18,** 1–10 (1965).
66. 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).
67. 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).
68. Widmann, A. & Schröger, E. Filter effects and filter artifacts in the analysis of electrophysiological data. *Front. Psychol.* **3,** 233 (2012).
69. 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).
70. 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).
71. 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).
72. 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).

73. Miller, J., Ulrich, R. & Schwarz, W. Why jackknifing yields good latency estimates. *Psychophysiology* **46,** 300–312 (2009).
74. Pfurtscheller, G. & Lopes da Silva, F. H. Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clin. Neurophysiol.* **110,** 1842–1857 (1999).
75. Busemeyer, J. R. & Diederich, A. *Cognitive Modeling*. (SAGE, 2010).
76. Nelder, J. A. & Mead, R. A Simplex Method for Function Minimization. *The Computer Journal* **7,** 308–313 (1965).
77. Burnham, K. P. & Anderson, D. R. Multimodel Inference: Understanding AIC and BIC in Model Selection. *Sociol. Methods Res.* **33,** 261–304 (2004).

Acknowledgments

The authors thank Sarita Tamang and Genevieve Price for data collection. This work was supported by a research grant to SPK and RGO'C from the U.S. National Science Foundation under grant number BCS-1358955, a research grant from Science Foundation Ireland to SPK (grant no. 15/CDA/3591), and a European Research Council (ERC) Starting Grant to RGO'C under the European Union's Horizon 2020 research and innovation programme (grant no. 638289). EAC 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

SPK and RGO'C conceived the research. SPK had the data collected and analysed them. SPK and EAC developed and fit the computational models. SPK, EAC and RGO'C wrote the paper.

Competing interests

The authors declare no competing interests.

Figure legends

Figure 1. Prior-cued motion discrimination task and behavioural results.

a. 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 8.4 ± 1.7 %). In the Deadline regime, a shorter deadline was individually set (mean 485 ± 61 ms) to be close to the median response time (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 both LoCoh and Deadline relative to Easy; see Supplementary Table 3 for all pairwise comparisons). **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 $\eta^2=0.83$; both LoCoh and Deadline > Easy; see Supplementary Table 3 for all pairwise comparisons), with a significant interaction between the two ($F(4,76)=27.2$, $p<0.001$, partial $\eta^2=0.59$; valid-invalid difference bigger for LoCoh and Deadline relative to Easy). Error bars indicate S.E.M. after between-subject variance is factored out. Individual subjects are shown as background dots to convey inter-individual variability in performance.

Figure 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 mark the first time point at which the temporal slope (collapsed across cue validity) reaches a significantly positive value (382 ms post-cue for Deadline and 676 ms for both Easy and LoCoh). Given that MB is measured in 294-ms windows centred at each timepoint, these buildup onsets are too early to be driven by evidence encoding. Note that y-axis is flipped so that increasing motor preparation (decreasing MB) is upwards. Amplitudes at the pre-evidence time-point (-59 ms) used for model constraints are shown to the right. **b.** Mean pre-response MB amplitude contralateral to the response, for each prior condition and regime, with correct responses divided into 5 equally-sized RT bins. The mean amplitude across conditions represents the 'Motor Threshold,' indicated by a green dashed line in panels 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 is taken to mark the time at which motor preparation gives way to motor execution processes, and is

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 9 conditions. Topographies show 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 are quantitatively constrained by pre-evidence MB amplitudes in panel a, and post-decision motor execution times by the pre-response execution potential latencies in panel 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 buildup rate, the correct alternative happens to start lower yet ultimately wins the race due to cumulative evidence. Supplementary Fig. 3 explicitly lays out all assumed mathematical relationships and offers a plausible systems-level implementation. Error bars in panels a-c show S.E.M. after factoring out between-subject variance.

Figure 3. Comparison of real versus mode-simulated behaviour.

a. Empirical RT distributions for correct (thick) and incorrect (thin) responses under the Easy, LoCoh, and Deadline regimes, averaged across subjects. The proportion of misses (no response within 1600 ms of evidence onset) for each condition is shown as a bar on the right. Mean RT reproduced from Fig 1b is marked above the RT distributions, and mean error rates reproduced from Fig 1c are plotted in the insets, to aid comparison with model-simulated data. Error bars show standard error of the mean, in some cases smaller than the marker size. At the bottom, conditional accuracy functions are shown for all conditions, quantifying the accuracy in each of 7 equal-sized RT bins. To follow up on a Regime x Validity x RT-bin interaction ($F(24, 456) = 14.3$, $p<0.001$, partial $\eta^2=0.43$), a Validity x RT-bin repeated-measures ANOVA was carried out in each Regime and revealed significant main effects of both factors in all 3 regimes, and a significant Validity x 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 2nd-slowest bin for LoCoh and Easy (but not Deadline, which lacked slow RTs in general), and the fastest RT bin had significantly lower accuracy than the 2nd-fastest in all 3 regimes (Supplementary Table 6). **b.** Simulated behavioural data from the Neurally-informed (NI) model, plotted in an identical way to the real behaviour. **c.** Simulated behavioural data from the constant-bound, Drift Diffusion Model (DDM). See Supplementary Fig. 4 for the version of this figure for models without drift

rate biases. Supplementary Fig. 5 shows conditional accuracy functions for each individual subject, demonstrating that grand average behavioural patterns are representative of those exhibited by individual decision makers.

Figure 4. Parameter estimates showing regime effects of the DDM and neurally-informed (NI) model.

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, normalised with respect to the value in the Easy regime to facilitate cross-model comparison. **c.** Urgency rate estimated by the NI model. Inset shows 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. Error bars show the 95% confidence intervals 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 just the across-subject 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-subject variations in the quality of the EEG signals used in the neural constraints.

Figure 5. Real and simulated average decision signal dynamics and amplitudes.

All data are for correct trials only. **a.** Stimulus-locked and **b.** response-locked MB amplitude reflecting motor preparation contralateral and ipsilateral to the (correct) response executed. Waveforms are here averaged across cue validities to maximise signal robustness (see Supplementary **Fig. 6** for traces broken out by validity). **c.** Amplitudes of contralateral and ipsilateral MB just prior to response, broken out by prior cue validity. **d.** Stimulus-locked and **e.** response-locked centro-parietal positivity (CPP) for each regime. Note that although the CPP continues to rise and peak beyond the point where we measure its pre-response amplitude (a 59-ms window centered -97.5 ms prior to response, i.e. the earliest motor execution onset), this likely represents post-decision accumulation [42,43], 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 5 RT bins for each prior cue type and Regime highlights an overall inverted U-shaped relationship with RT. Error bars in all panels of the real data show S.E.M. after between-subject variability has been factored out. **h-n.** The same plots are 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 (see Methods).

Figure 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 centro-parietal positivity (CPP). Note that since the color 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 (see main text). Amplitudes in the pre-evidence time-interval [-59 0] ms are shown to the right, with error bars showing S.E.M.

Figure 7. Relationship between drift rate biases and CPP amplitude modulations.

a. Simulation of average evidence accumulation trajectories (upward-building 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/upwards for trials with valid/invalid cues. For a given drift rate, the latter bound shifts

result in the accumulator reaching higher levels for invalid than valid trials. However, a bias in drift rate tends to counteract this effect because the slower buildup on invalid trials makes the accumulator cross the collapsing bound later and hence lower down. In the LoCoh regime, the bounds are high and the drift rate is low, so that a relatively small bias in the drift rate is enough to fully cancel-out the effect of starting point bias on 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 quantifying 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; hence, 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 illustrated in panel a, or observed in both the real and simulated CPP (Fig. 5f,m).

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 full neurally-informed (NI) model and 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 Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) metrics which respectively penalise less and more strictly for complexity for reference. Given the purpose of the study to establish the full range of effects with convergent support from neural data, we use AIC for model selection throughout. Lower values of G^2, AIC and BIC indicate a better fit. For the NI model, 'tnd' refers to accumulation onset only, because post-decision motor times were constrained by the neural data, whereas for the DDM it refers to the total non-decision time.

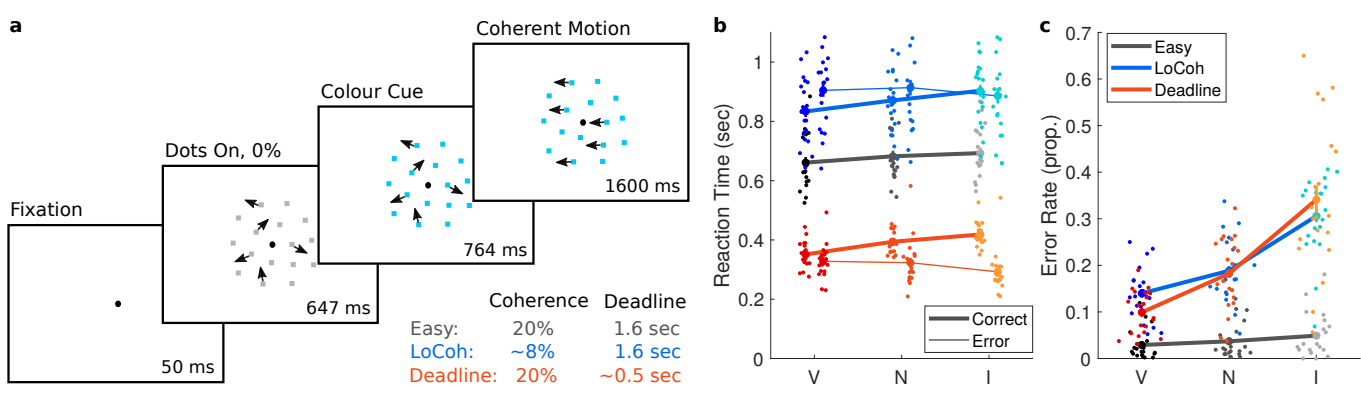
a
Coherent Motion
Colour Cue
Dots On, 0%
Fixation
1600 ms
764 ms
647 ms
50 ms
Coherence
Deadline
Easy: 20% 1.6 sec
LoCoh: ~8% 1.6 sec
Deadline: 20% ~0.5 sec
b
Reaction Time (sec)
Correct
Error
V
N
I
c
Error Rate (prop.)
Easy
LoCoh
Deadline
V
N
I

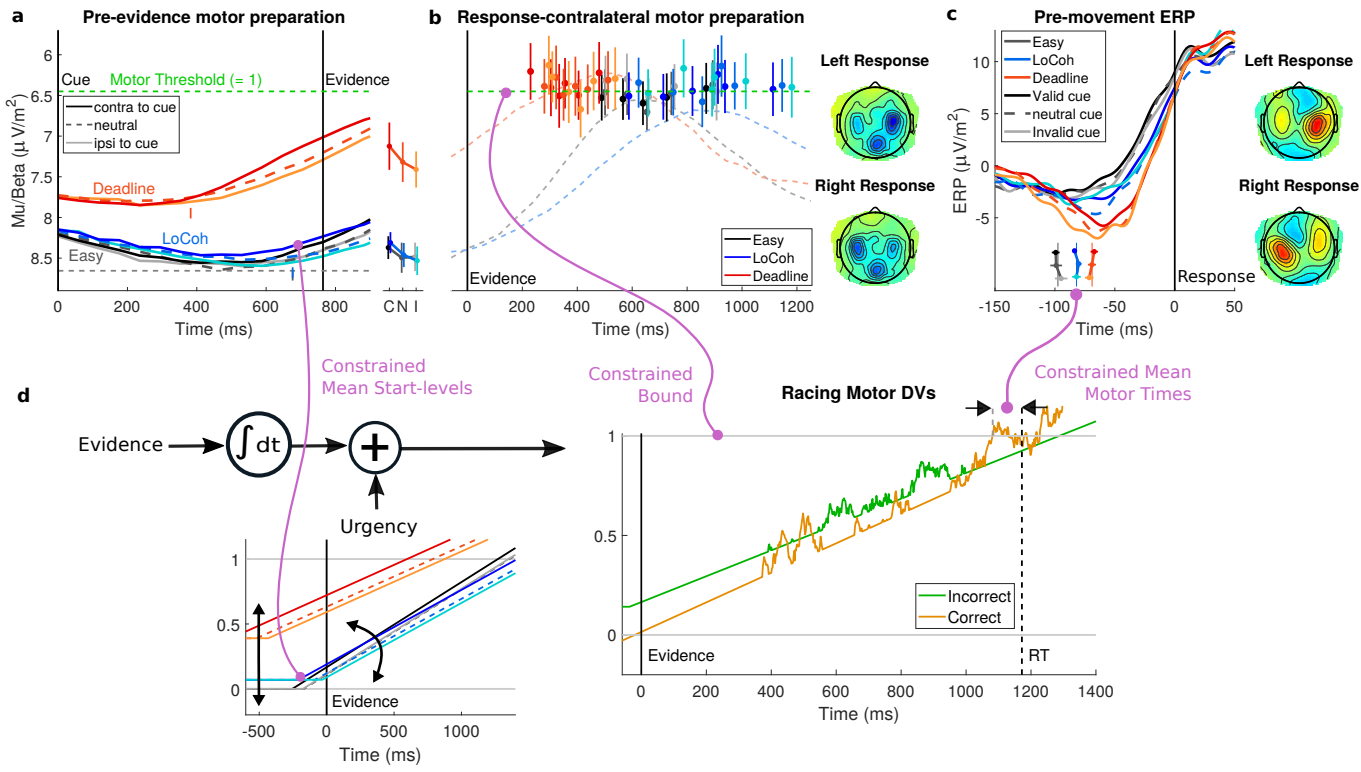
a
Pre-evidence motor preparation
Cue
Motor Threshold (= 1)
Evidence
contra to cue
neutral
ipsi to cue
Deadline
LoCoh
Easy
Mu/Beta (μV/m²)
Time (ms)
C N I
b
Response-contralateral motor preparation
Left Response
Right Response
Easy
LoCoh
Deadline
Evidence
Time (ms)
c
Pre-movement ERP
Easy
LoCoh
Deadline
Valid cue
neutral cue
Invalid cue
ERP (μV/m²)
Left Response
Right Response
Response
Time (ms)
d
Evidence
∫dt
Urgency
Constrained
Mean Start-levels
Constrained
Bound
Racing Motor DVs
Constrained Mean
Motor Times
Incorrect
Correct
Evidence
RT
Time (ms)

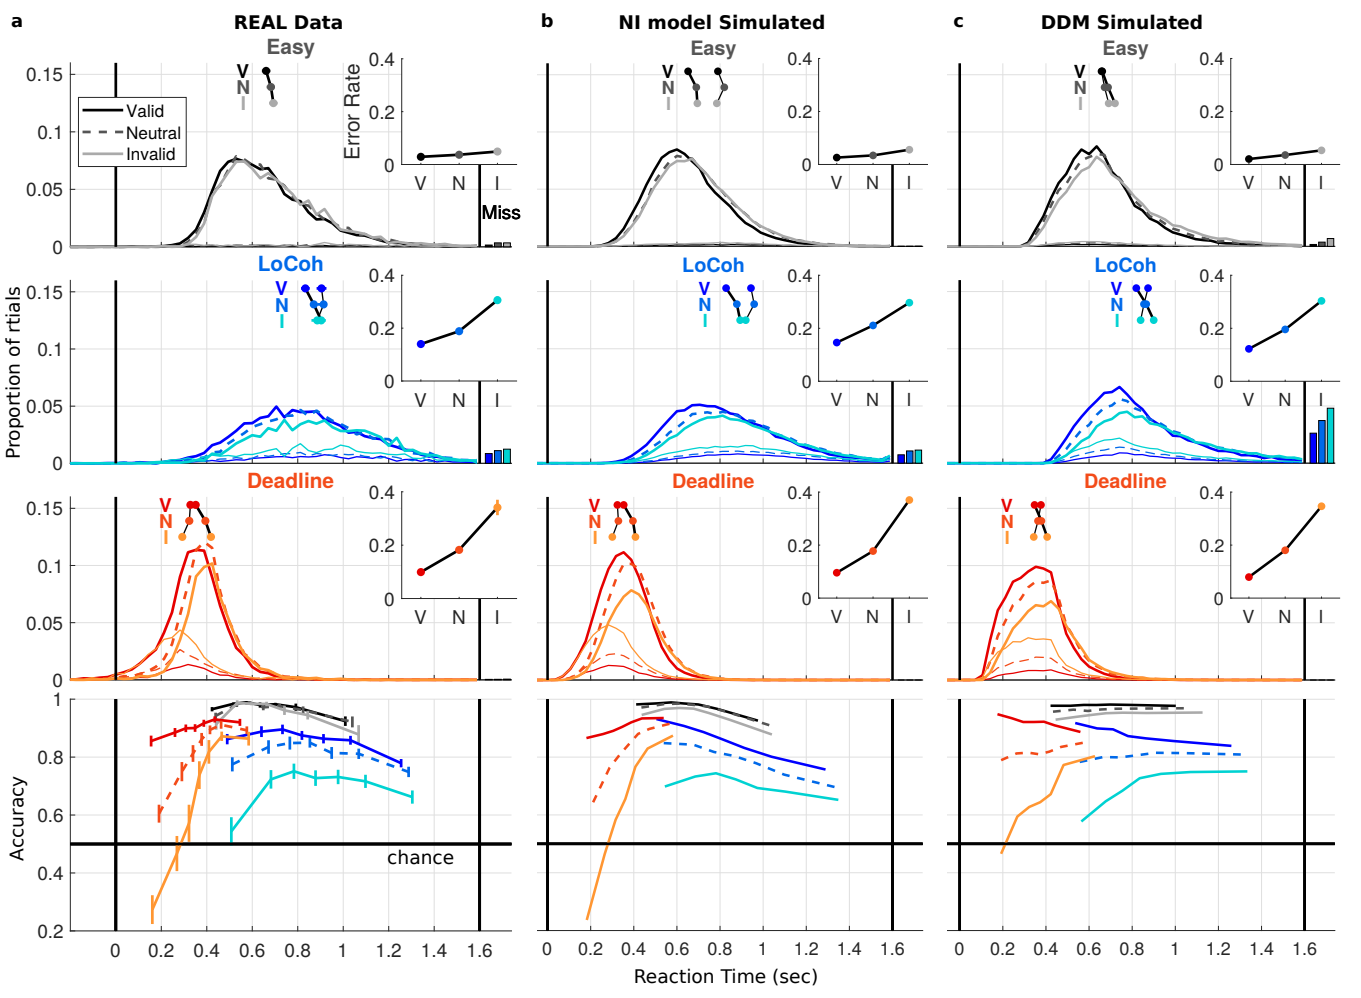
a
REAL Data
b
NI model Simulated
c
DDM Simulated
Easy
LoCoh
Deadline
Valid
Neutral
Invalid
Error Rate
Miss
V
N
I
Proportion of rtials
Accuracy
chance
Reaction Time (sec)

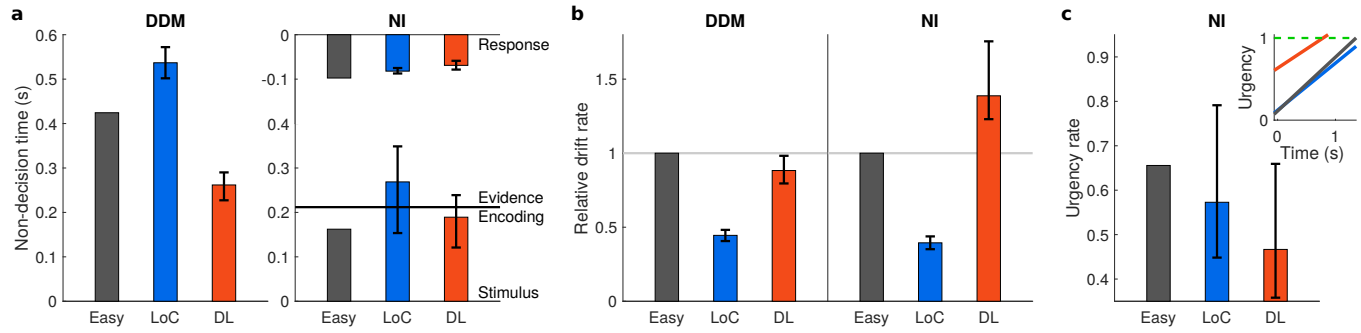
a
DDM
NI
Non-decision time (s)
0.6
0.5
0.4
0.3
0.2
0.1
0
-0.1
Response
Evidence
Encoding
Stimulus
Easy
LoC
DL
b
DDM
NI
Relative drift rate
1.5
1
0.5
0
c
NI
Urgency rate
0.9
0.8
0.7
0.6
0.5
0.4
Urgency
Time (s)

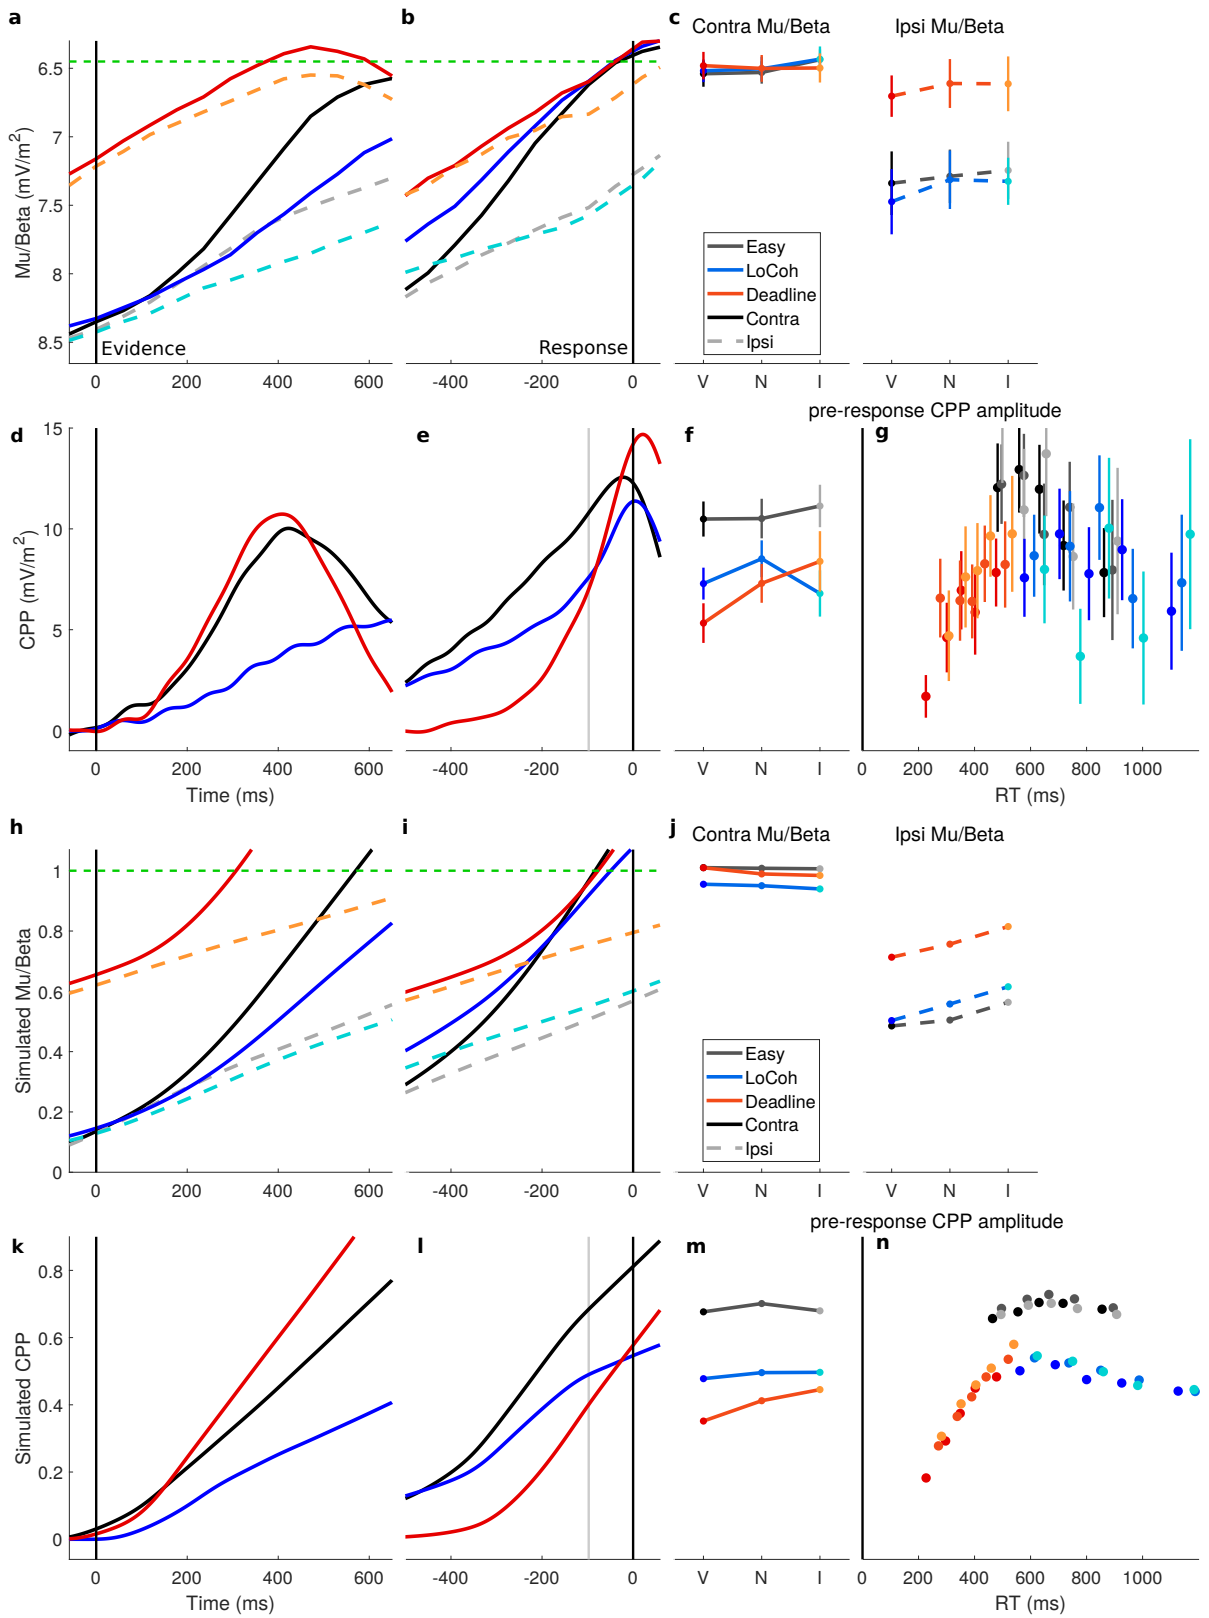
a
b
c
Contra Mu/Beta
Ipsi Mu/Beta
Mu/Beta (mV/m^2)
Evidence
Response
Easy
LoCoh
Deadline
Contra
Ipsi
V
N
I
pre-response CPP amplitude
d
e
f
g
CPP (mV/m^2)
Time (ms)
RT (ms)
h
i
j
Simulated Mu/Beta
k
l
m
n
Simulated CPP

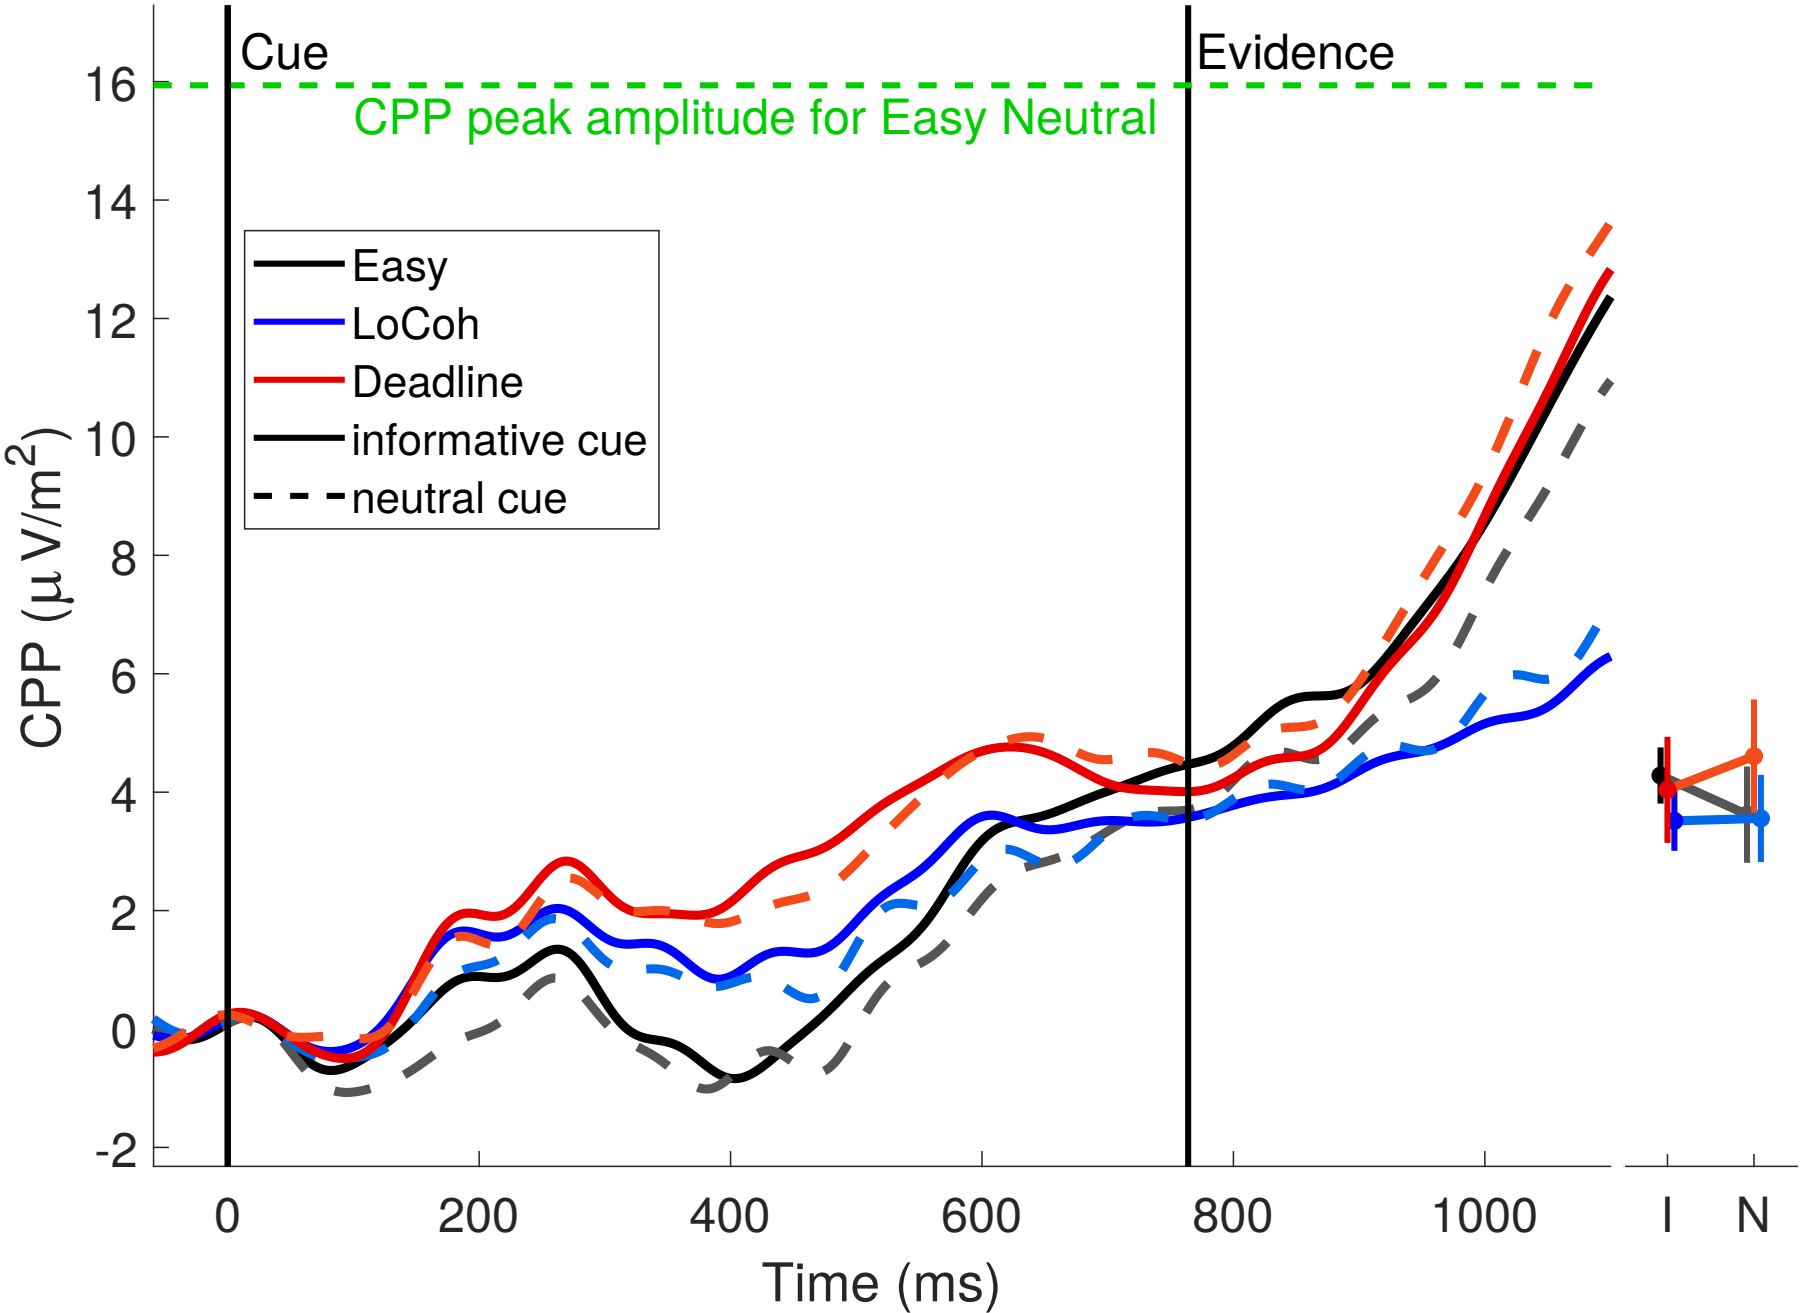
Cue
Evidence
CPP peak amplitude for Easy Neutral
Easy
LoCoh
Deadline
informative cue
neutral cue
CPP (μV/m2)
16
14
12
10
8
6
4
2
0
-2
0
200
400
600
800
1000
I
N
Time (ms)

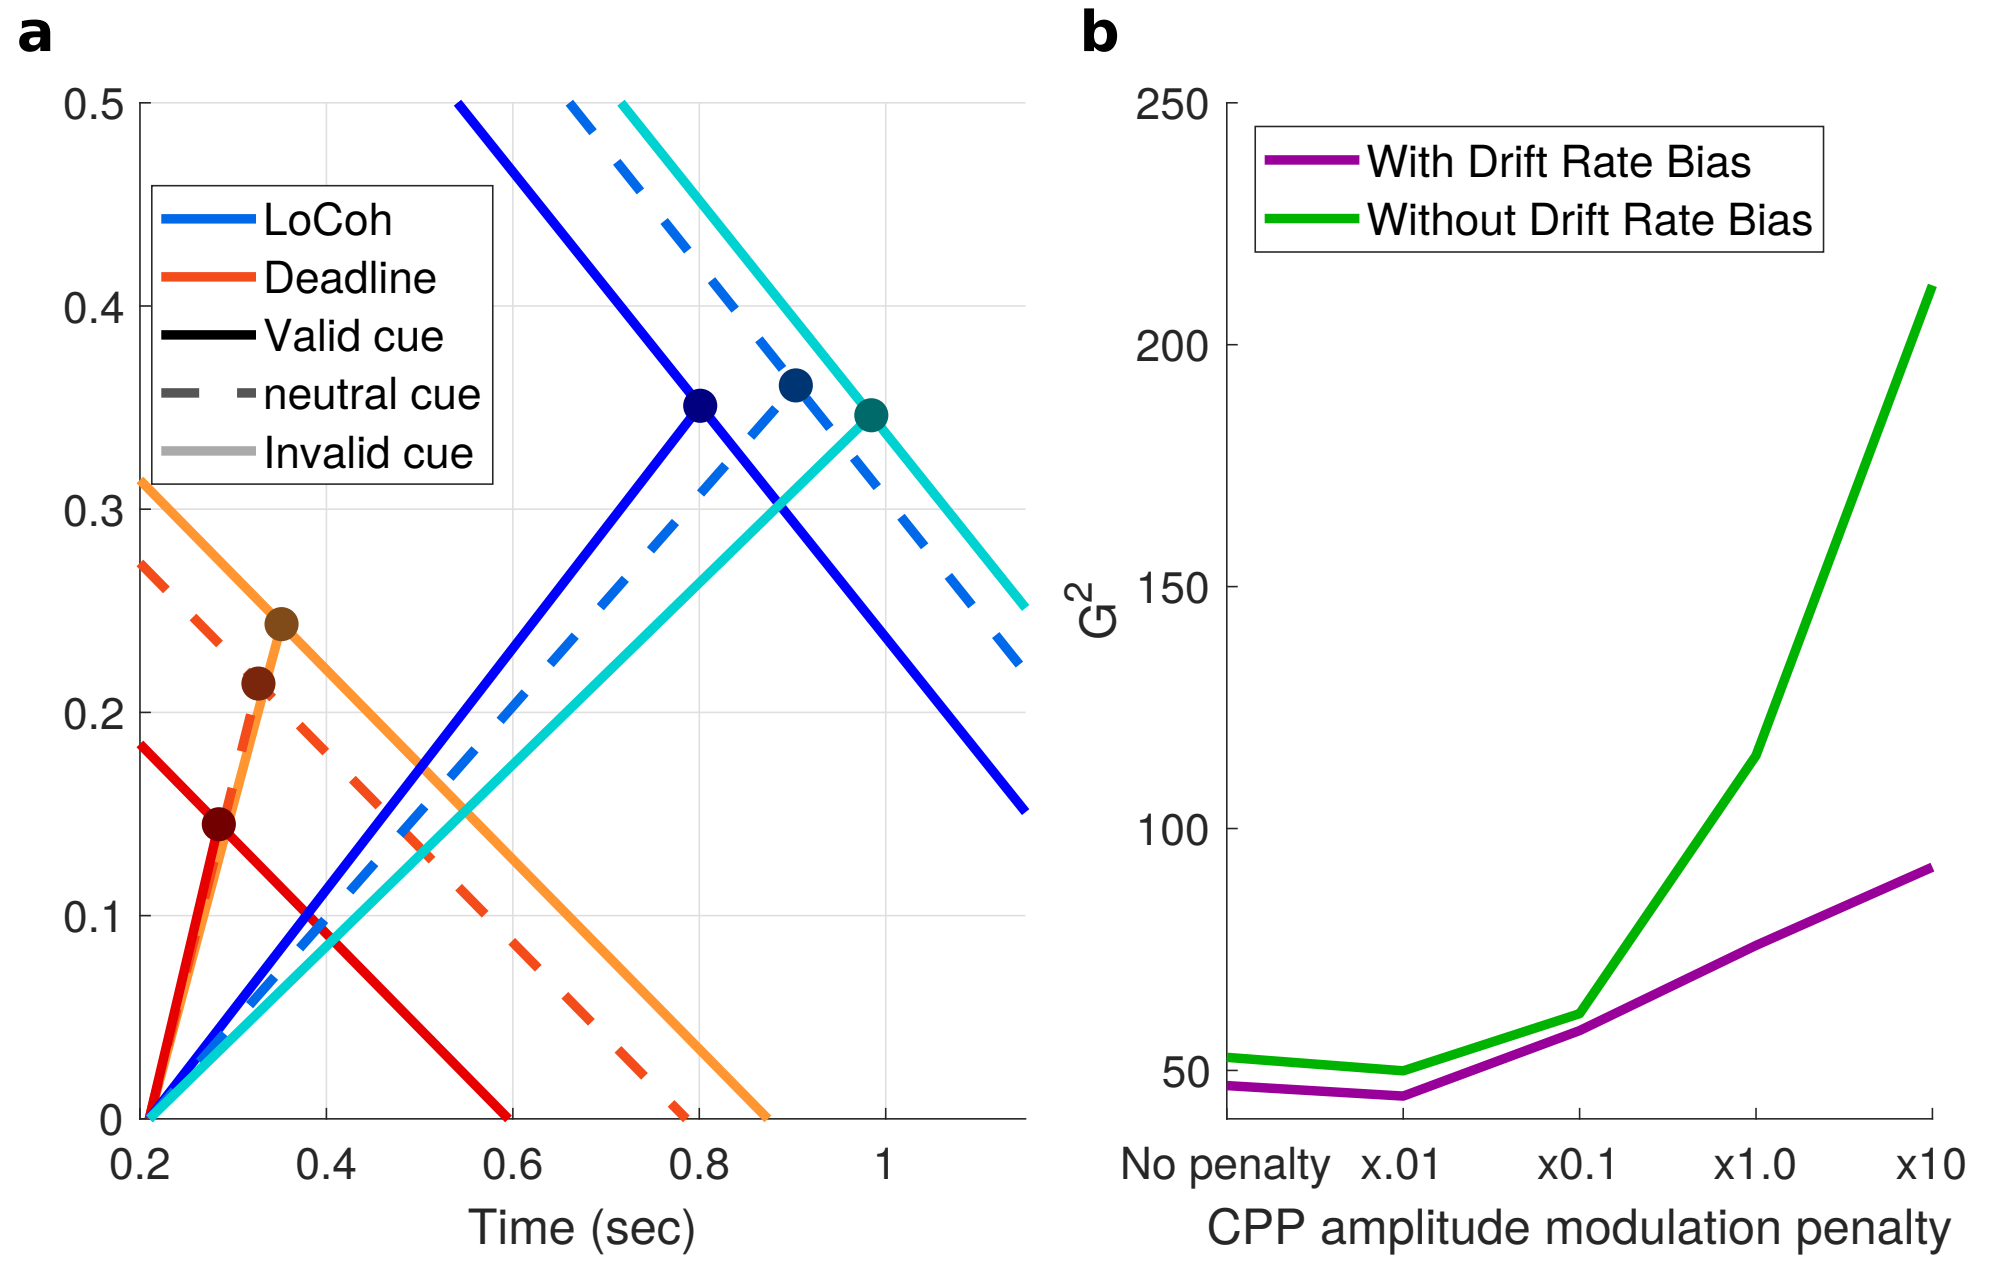
a
b
LoCoh
Deadline
Valid cue
neutral cue
Invalid cue
0.5
0.4
0.3
0.2
0.1
0
0.2
0.4
0.6
0.8
1
Time (sec)
With Drift Rate Bias
Without Drift Rate Bias
250
200
150
100
50
G2
No penalty
x.01
x0.1
x1.0
x10
CPP amplitude modulation penalty

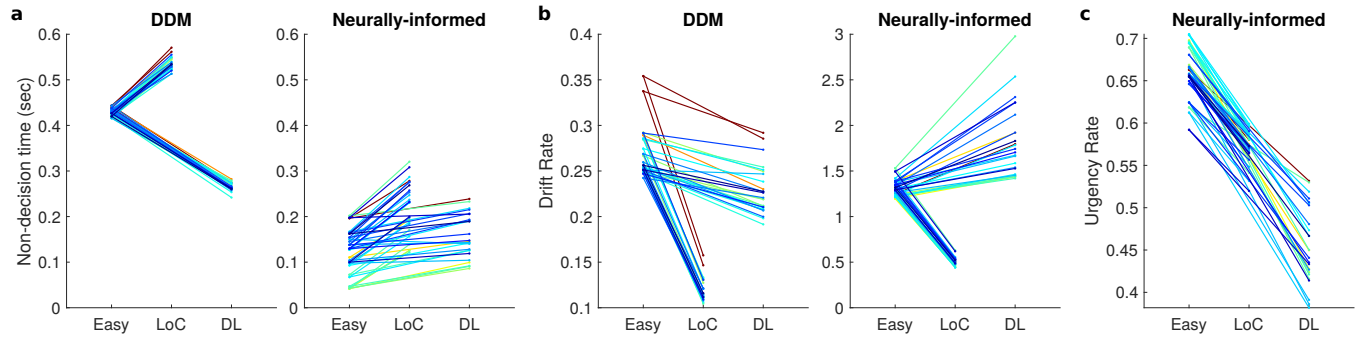
a
DDM
Neurally-informed
Non-decision time (sec)
0.6
0.5
0.4
0.3
0.2
0.1
0
Easy
LoC
DL
b
DDM
Neurally-informed
Drift Rate
0.4
0.35
0.3
0.25
0.2
0.15
0.1
3
2.5
2
1.5
1
0.5
0
c
Neurally-informed
Urgency Rate
0.7
0.65
0.6
0.55
0.5
0.45
0.4

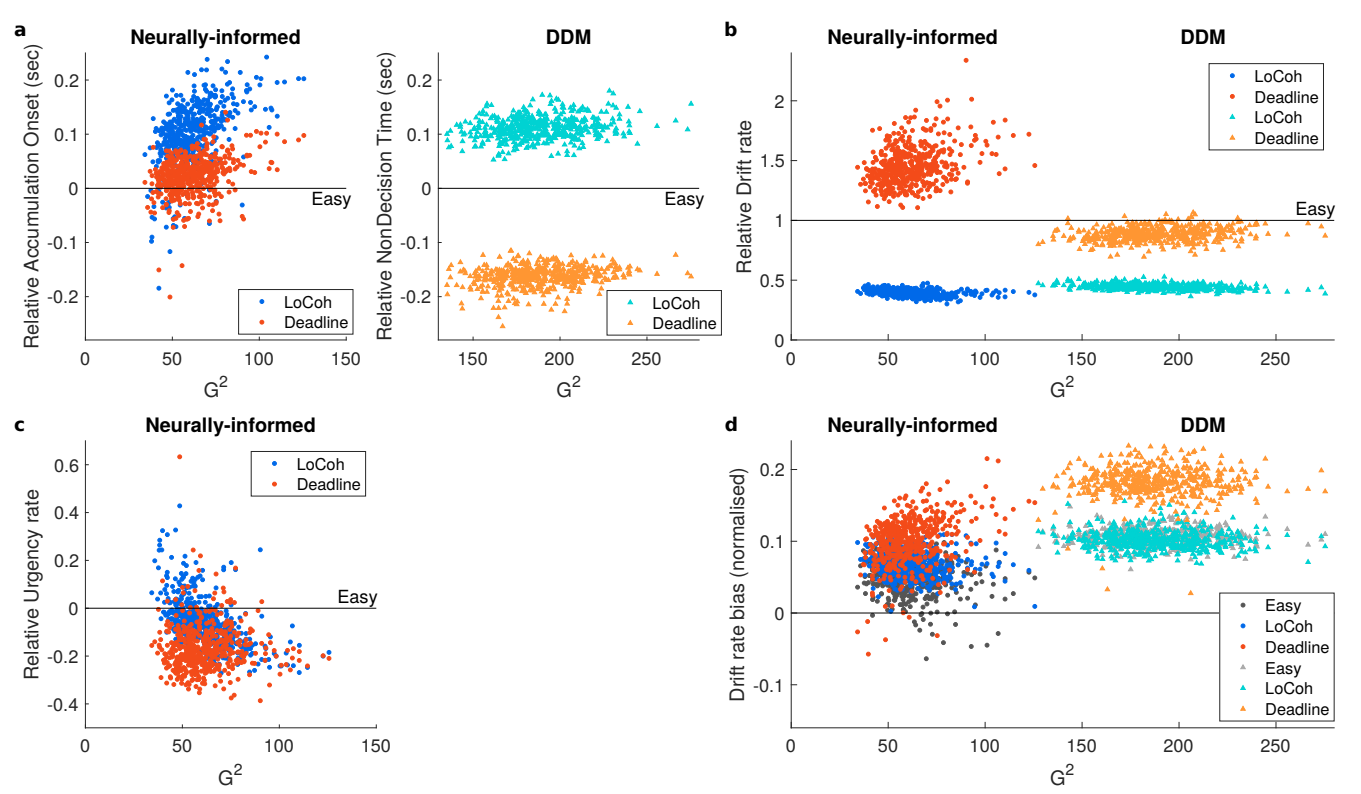
a
Neurally-informed
DDM
Relative Accumulation Onset (sec)
Relative NonDecision Time (sec)
Easy
LoCoh
Deadline
G2
b
Neurally-informed
DDM
Relative Drift rate
LoCoh
Deadline
LoCoh
Deadline
Easy
G2
c
Neurally-informed
Relative Urgency rate
LoCoh
Deadline
Easy
G2
d
Neurally-informed
DDM
Drift rate bias (normalised)
Easy
LoCoh
Deadline
Easy
LoCoh
Deadline
G2

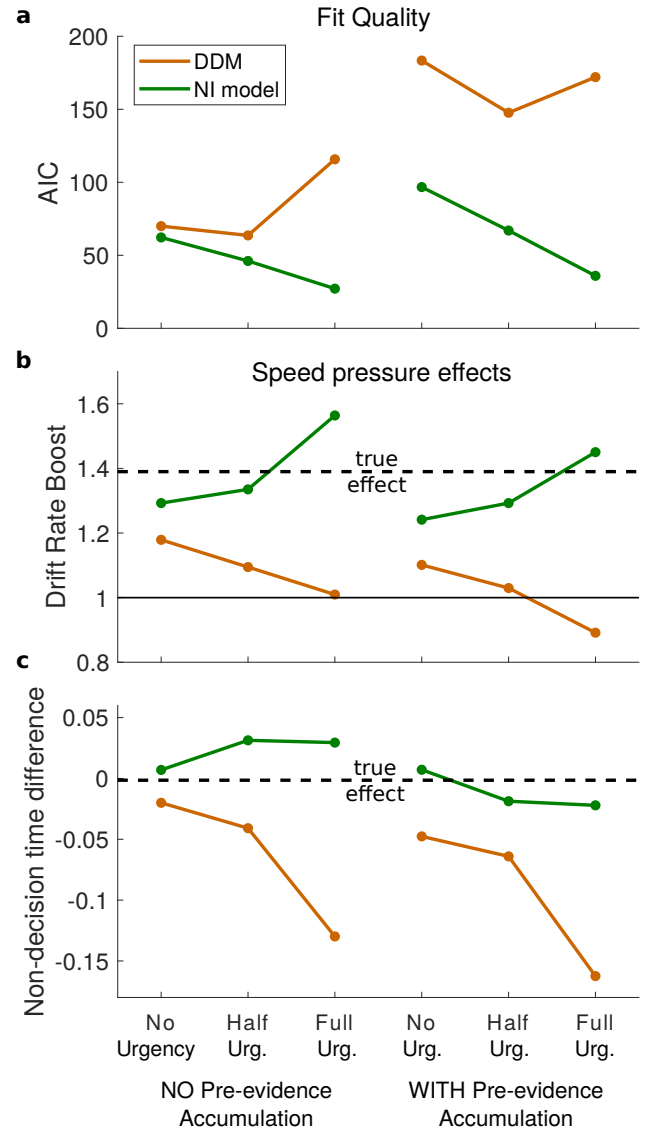
a
Fit Quality
DDM
NI model
AIC
200
150
100
50
0
b
Speed pressure effects
Drift Rate Boost
1.6
1.4
1.2
1
0.8
true
effect
c
Non-decision time difference
0.05
0
-0.05
-0.1
-0.15
true
effect
No
Urgency
Half
Urg.
Full
Urg.
No
Urg.
Half
Urg.
Full
Urg.
NO Pre-evidence
Accumulation
WITH Pre-evidence
Accumulation

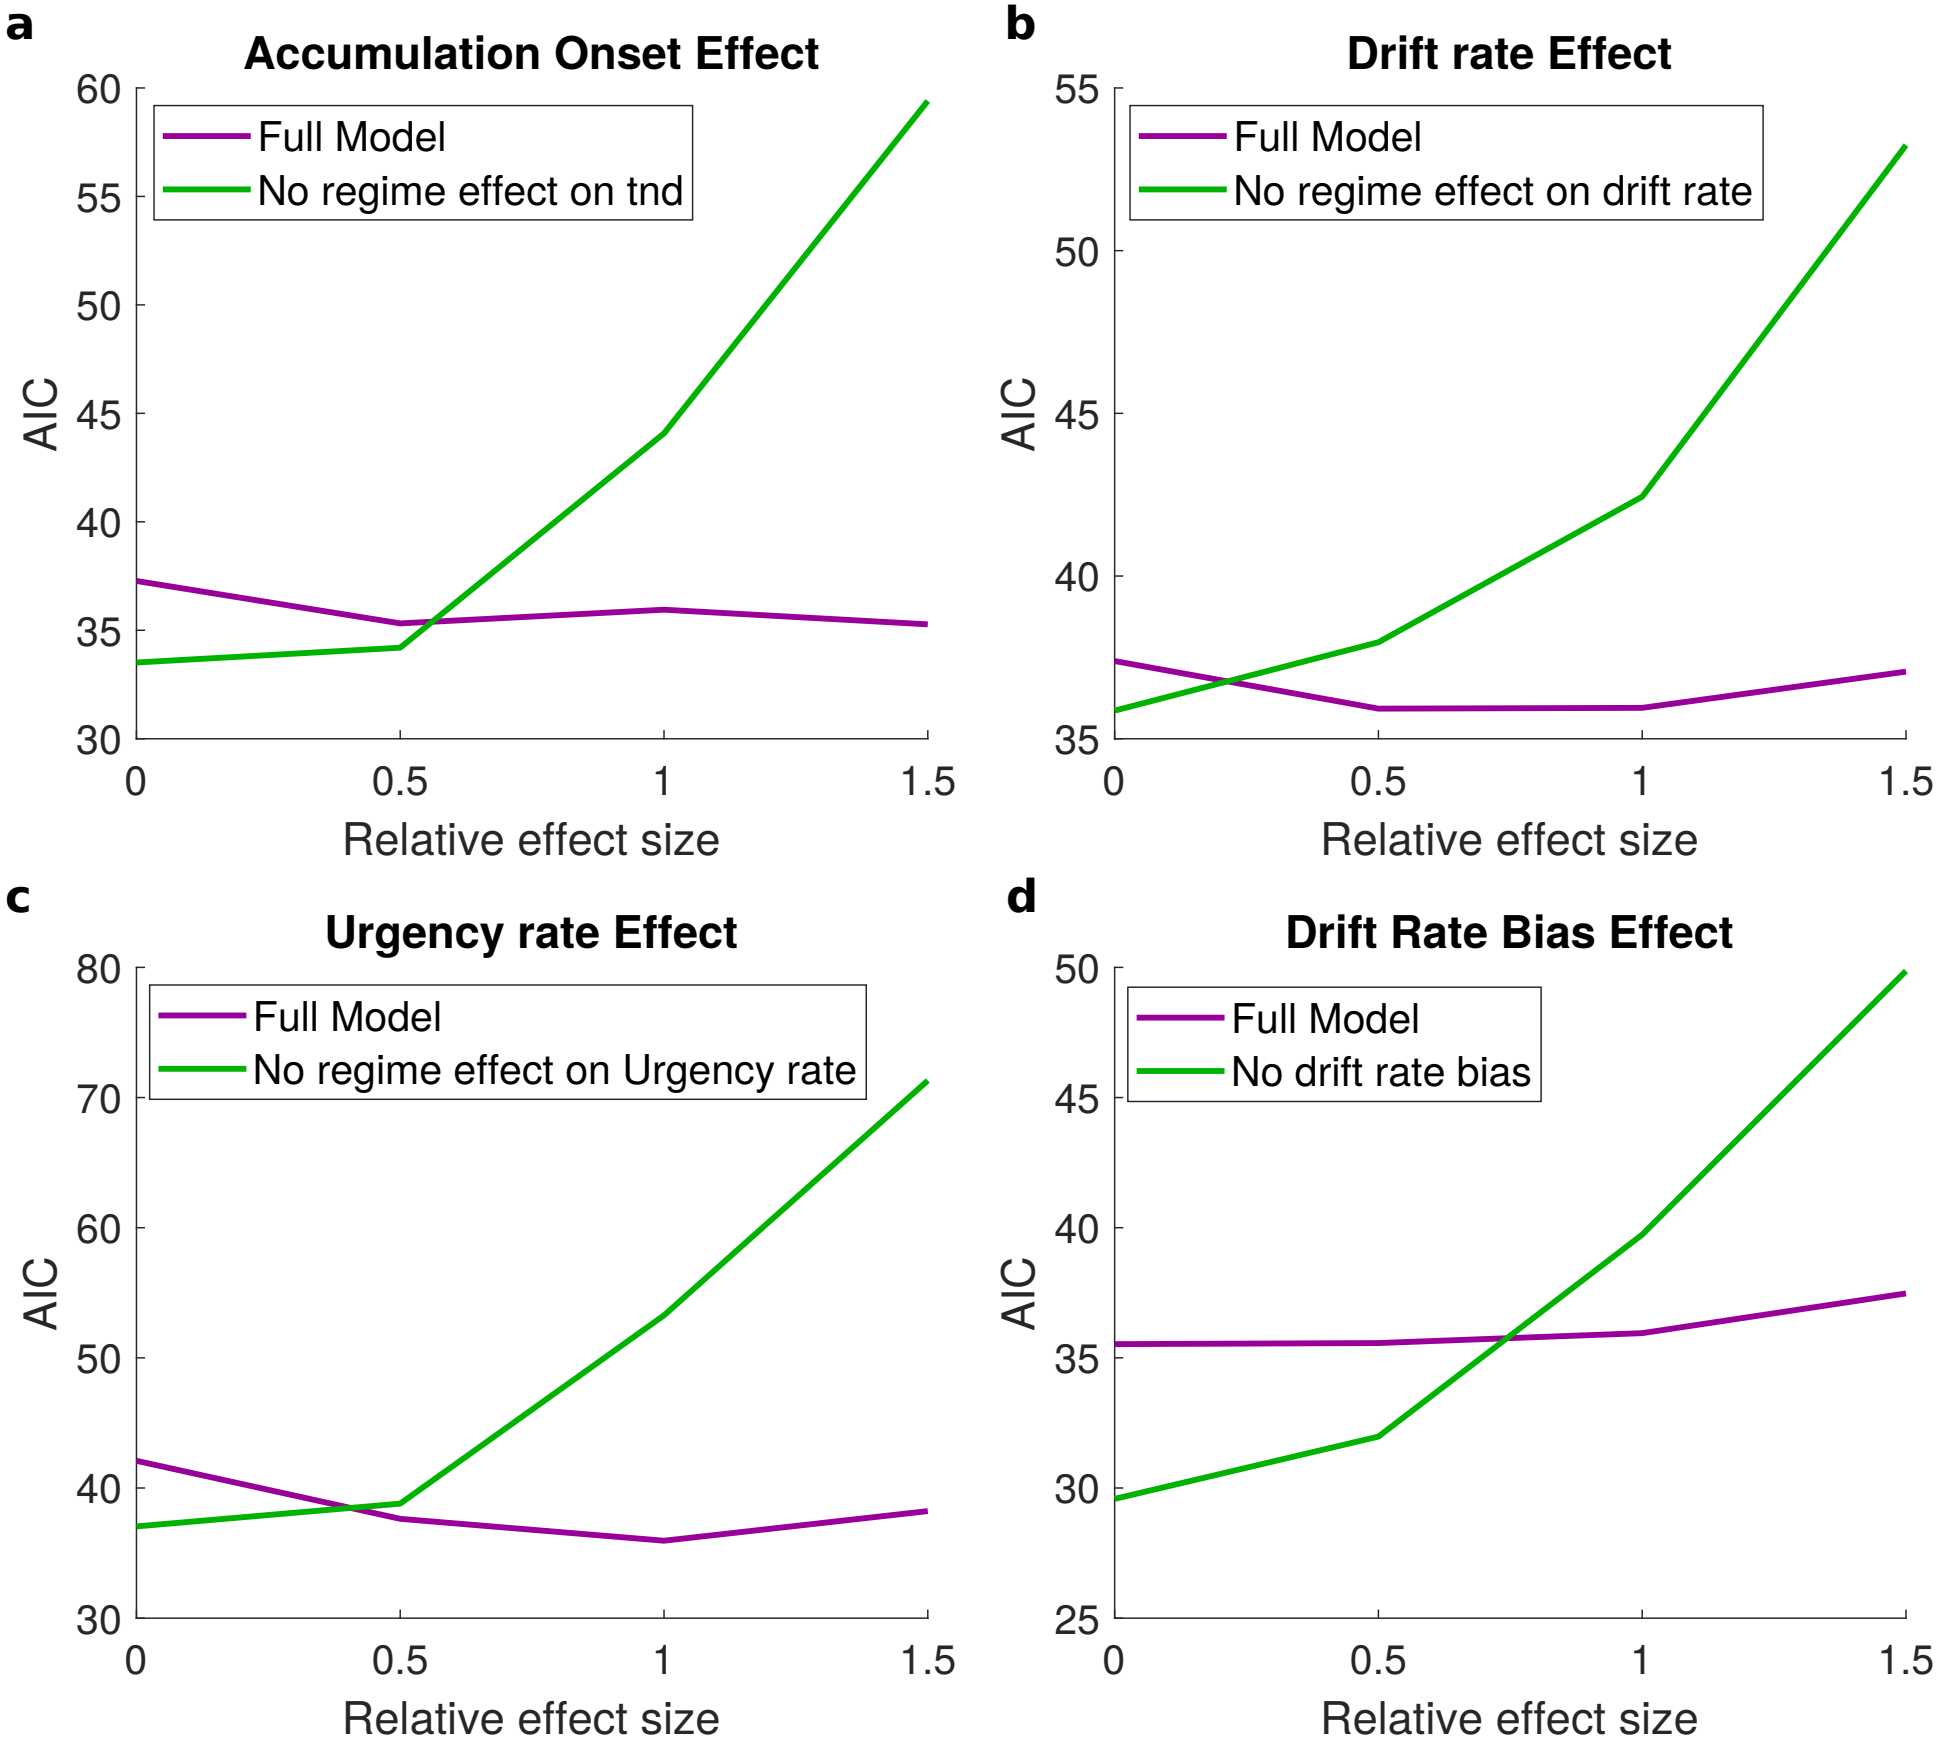
a
Accumulation Onset Effect
Full Model
No regime effect on tnd
AIC
Relative effect size
b
Drift rate Effect
Full Model
No regime effect on drift rate
AIC
Relative effect size
c
Urgency rate Effect
Full Model
No regime effect on Urgency rate
AIC
Relative effect size
d
Drift Rate Bias Effect
Full Model
No drift rate bias
AIC
Relative effect size

Kelly, Corbett and O'Connell (2020): Supplementary Information

Supplementary Tables

Supplementary Table 1: estimated parameter values for the neurally-informed model.

	Zv	Zn	Zi	U	D	ΔD	Tac	Te	Tm	S	Sz	St	Su
Easy	0.129	0.073	0.078	0.656	1.319	0.077	0.162	0.212	0.098	0.522	0.404	0.225	0.166
LoCoh	0.156	0.088	0.056	0.573	0.520	0.073	0.269	.	0.082	.	.	.	.
Deadline	0.695	0.606	0.565	0.467	1.830	0.114	0.189	.	0.069	.	.	.	.

Neurally-constrained parameter values are indicated in red, free parameters in black. **Zv**, **Zn** and **Zi** refer to the starting level of motor preparation relative to a threshold of 1 (acting as the scaling parameter) for valid, neutral and invalid trials, respectively. **U** refers to the urgency buildup rate, **D** the drift rate, **ΔD** the drift rate bias (the amount added to mean drift rate D for valid trials and subtracted for invalid trials), **Tac** the accumulation onset time relative to evidence presentation, **Te** the onset time of evidence encoding relative to evidence presentation, **Tm** the post-threshold motor-execution delay, **S** the within-trial Gaussian noise, **Sz** the starting point variability, **St** the non-decision time variability, and **Su** the urgency rate variability. Note that Te, S, Sz, St and Su each take a single value across all regimes.

Supplementary Table 2: estimated parameter values for a 'full' diffusion model fit to the behavioural data.

	b	z	d	Δd	tnd	sz	st	sd
Easy	0.070	0.003	0.254	0.027	0.424	0.046	0.311	0.033
LoCoh	0.064	0.008	0.113	0.025	0.537	.	.	.
Deadline	0.037	0.010	0.224	0.042	0.262	.	.	.

b refers to the bound, **z** the starting point bias, **d** the drift rate, **Δd** the drift rate bias (the amount added to mean drift rate **d** for valid trials and subtracted for invalid trials), **tnd** the non-decision time, **sz** the starting point variability, **st** the non-decision time variability, and **sd** the drift rate variability. Within-trial Gaussian noise is fixed at a value of 0.1, acting as the scaling parameter for the DDM. The three variability parameters each take a single value across all regimes.

Supplementary Table 3: Pairwise comparisons following up main effects and interactions found in rmANOVAs (Validity x Regime x Motion Direction) for mean correct RT and error rate (Fig. 1b).

Condition		df	t-value	p-value	CI	
Mean correct RT: Main effect of Validity						
Valid	Neutral	19	-6.72	2.01E-06	-0.044	-0.023
Valid	Invalid	19	-6.94	1.29E-06	-0.073	-0.039
Neutral	Invalid	19	-4.58	0.0002	-0.033	-0.012
Mean correct RT: Main effect of Regime						
Easy	LoCoh	19	-7.88	2.10E-07	-0.24	-0.14
Easy	Deadline	19	14.16	1.50E-11	0.25	0.33
Mean correct RT: Effect of Regime on Valid-Invalid difference						
Easy	LoCoh	19	2.36	0.029	0.004	0.071
Easy	Deadline	19	2.74	0.013	0.008	0.062
Error rate: Main effect of Validity						
Valid	Neutral	19	-9.30	1.68E-08	-0.057	-0.036
Valid	Invalid	19	-8.97	2.93E-08	-0.176	-0.109
Neutral	Invalid	19	-7.33	6.01E-07	-0.124	-0.069
Error rate: Main effect of Regime						
Easy	LoCoh	19	-16.63	8.80E-13	-0.19	-0.15
Easy	Deadline	19	-10.63	1.94E-09	-0.20	-0.14
Error rate: Effect of Regime on Valid-Invalid difference						
Easy	LoCoh	19	-16.63	6.91E-08	-0.19	-0.15
Easy	Deadline	19	-10.63	2.65E-06	-0.20	-0.14

Supplementary Table 4: Pairwise comparisons of Regime effects on post-decision motor time.

Condition		df	t-value	p-value	CI	
Stimulus-locked Mu/Beta slope						
Easy	LoCoh	19	-4.63	0.00018	-22.88	-8.63
Easy	Deadline	19	-4.76	0.00014	-41.03	-15.96

Supplementary Table 5: results of Validity x RT-bin rmANOVA conducted on the behavioural accuracy for each Regime (Conditional Accuracy Functions; Fig 3a, bottom).

Conditional Accuracy Function:

	df1	df2	F	p	partial η^2
Easy					
Main effect Validity	2	38	9.64	0.00041	0.34
Main effect RT bin	6	114	15.42	1.08E-05	0.45
Validity x RT bin Interaction	12	228	2.23	0.058	0.10
LoCoh					
Main effect Validity	2	38	69.89	1.85E-13	0.79
Main effect RT bin	6	114	12.39	1.89E-06	0.39
Validity x RT bin Interaction	12	228	4.83	0.0006	0.20
Deadline					
Main effect Validity	2	38	42.94	1.77E-10	0.69
Main effect RT bin	6	114	65.45	0	0.78
Validity x RT bin Interaction	12	228	29.48	5.22E-15	0.61

Supplementary Table 6: Follow-up to main effects of RT-bin in Supplementary Table 4, testing for the inverted-U nature of conditional accuracy functions by testing neighboring RT bins at both the slow and fast ends.

Condition	df	t-value	p-value	CI	
RT bin 7 vs bin 6					
Easy	19	3.79	0.0012	0.023	0.081
LoCoh	19	5.54	2.43E-05	0.041	0.090
Deadline	19	1.30	0.21	-0.008	0.033
RT bin 1 vs bin 2					
Easy	19	5.38	3.40E-05	0.027	0.060
LoCoh	19	5.63	2.00E-05	0.056	0.121
Deadline	19	5.18	5.31E-05	0.073	0.172

Supplementary Table 7: Results of 3x2 rmANOVA (Regime x Cue informativeness) of CPP amplitude at a pre-evidence baseline interval of -59 to 0 ms with respect to evidence onset. Bayes factors are included for the comparison of the full model with a version with each term excluded (BFexcl), and also for the full model relative to the null model (BF01 at bottom).

	df1	df2	F	p	partial η^2	BFexcl
Cue information	1	19	0.001	0.97	6.18E-05	7.83
Regime	2	38	0.23	0.61	0.012	12.52
Cue x Regime	2	38	0.98	0.37	0.049	89.98
BF01 Null Vs Full Model = 271.3						

Supplementary Table 8: Results of 3x2 rmANOVA (Regime x Cue informativeness) of CPP signal slope in the first 118 ms post-evidence.

	df1	df2	F	p	partial η^2	BFexcl
Cue information	1	19	0.17	0.69	0.009	7.06
Regime	2	38	1.35	0.27	0.066	3.69
Cue x Regime	2	38	0.58	0.57	0.030	28.09
BF01 Null Vs Full Model = 67.04						

Supplementary Table 9: Pairwise comparisons following main effects of Regime on CPP slope.

Condition		df	t-value	p-value	CI	
Stimulus-locked CPP slope						
Easy	LoCoh	19	7.02	1.09E-06	0.020	0.036
Easy	Deadline	19	-1.15	0.27	-0.014	0.004
Response-locked CPP slope: 118 ms window ending at -97.5 ms (Easy motor time)						
Easy	LoCoh	19	0.42	0.68	-0.011	0.017
Easy	Deadline	19	-2.24	0.038	-0.035	-0.001
Response-locked CPP slope: 118 ms window ending at regime-specific motor time						
Easy	LoCoh	19	-0.02	0.98	-0.015	0.014
Easy	Deadline	19	-3.30	0.0038	-0.041	-0.009

Supplementary Table 10: Pairwise comparisons following main effects of Regime on Mu/Beta slope.

Condition		df	t-value	p-value	CI	
Stimulus-locked Mu/Beta slope						
Easy	LoCoh	19	2.41	2.65E-02	0.16	2.29
Easy	Deadline	19	3.32	0.0036	0.80	3.55
Response-locked Mu/Beta slope						
Easy	LoCoh	19	2.36	0.029	0.16	2.59
Easy	Deadline	19	2.23	0.038	0.09	2.78

Supplementary Table 11: Results of one-way rmANOVAs testing for effects of Validity on pre-response CPP amplitude in each Regime

Condition	df1	df2	F	p	partial η^2	BF01
Easy	2	38	0.62	0.45	0.03	4.65
LoCoh	2	38	0.54	0.59	0.03	5.06
Deadline	2	38	8.57	0.00085	0.31	

Supplementary Figures

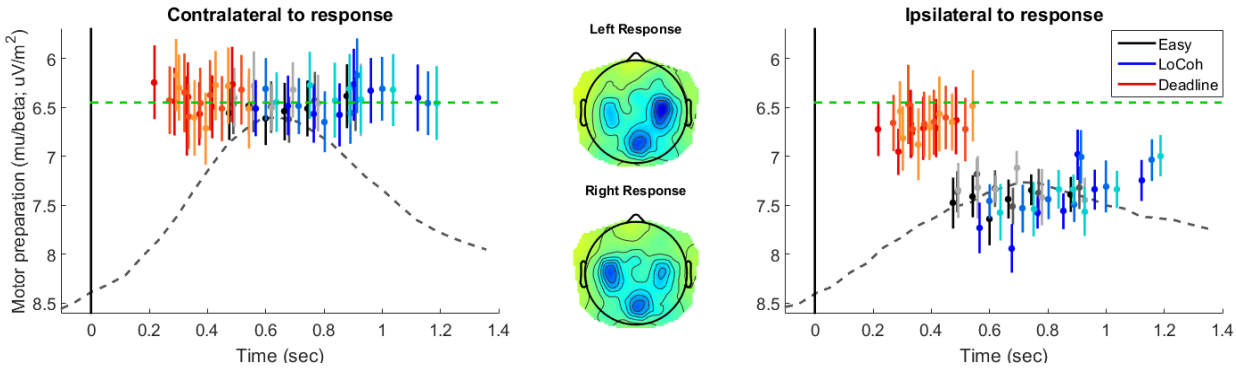

Supplementary Fig. 1: Motor preparation for the responding hand reaches a fixed threshold level contralateral to the responding hand, with lower levels reached by the non-responding side. Mu/Beta (MB) amplitudes contralateral (left) and ipsilateral (right) to the executed response, measured just prior to the time of response with the data for each prior condition and regime divided into 5 equally-sized RT bins for correct responses and one for errors. As a reference, the average stimulus-locked motor preparation waveforms for neutrally-cued, correct trials in the Easy regime are plotted as dashed traces in the background. Note that the y-axis is flipped so that increasing motor preparation, reflected in decreasing spectral amplitude, is upwards. Across an expansive range of RT, the contralateral MB amplitude exhibits remarkable stereotypy, indicating a fixed-threshold relationship to action execution at the motor level, whereas by comparison, MB amplitude ipsilateral to response reaches a point consistently lower than threshold and varies extensively across conditions and RT. The mean level of motor preparation at response across all conditions is indicated by a green dashed line. This is set as the threshold level in the neurally-informed model. Error bars for each point (mean ± standard error) were computed after factoring out between-subject variance.

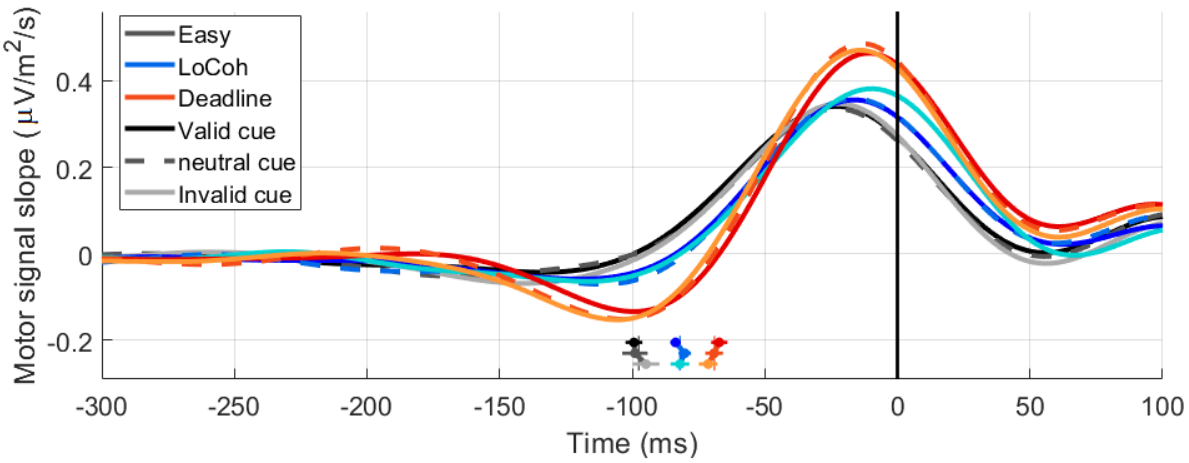

Supplementary Fig. 2: The first temporal derivative of the same motor signals as shown in Fig. 2c following low-pass filtering, quantifying momentary signal slope. The time point at which the first-derivative signal for each condition passes through zero, indicated below the waveforms for all conditions (with error bars showing S.E.M.), marks the inflection point for the motor signal, where negative-going motor preparation activity gives way to a sharp potential assumed to reflect motor execution.

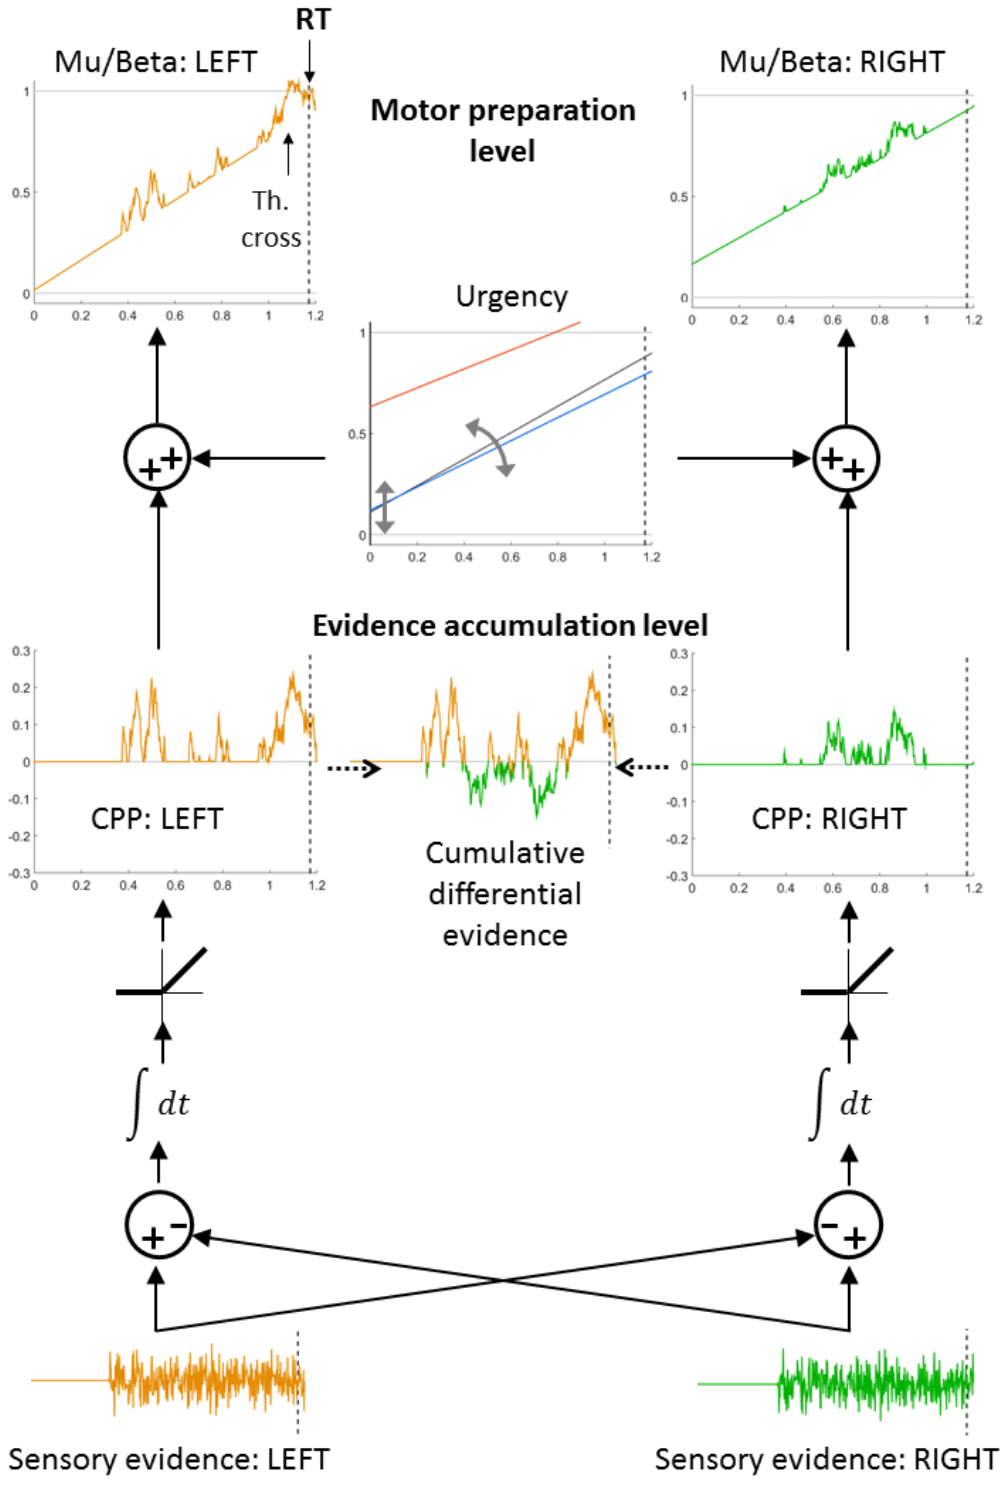

Supplementary Fig. 3: Multi-level implementation of the bounded evidence accumulation model with dynamic urgency. Fig. 2D describes the essential properties of the neurally-informed (NI) model that govern how it is constructed, constrained and validated by neural decision signals - specifically, that the ultimate bounded decision variables at the motor preparation level have fixed bounds and are driven by the sum of cumulative evidence plus

regime-adjusted, biased and temporally growing urgency. Here, we lay out a detailed implementation that serves the dual purpose of making all assumed mathematical relationships in the model explicit, and proposing one possible way these functions could plausibly be implemented in a neural information processing architecture. The essential characteristics that enable simulation of key decision variables for comparison with empirical neural signals are that the process generating the CPP represents pure cumulative differential evidence with no urgency component, with its scalp signature reflecting the absolute value of this, whereas motor preparation signals reflected in Mu/Beta (MB) are driven by the sum of cumulative evidence plus urgency. In this architecture, similar to previous implementations of a diffusion process for differential evidence [1], the evidence representations for the two competing directions of motion (left and right) are subtracted and accumulated by two separate accumulators. As described previously [2], to account for the fact that the CPP always builds with the same surface-positive polarity irrespective of which choice alternative the cumulative evidence favours, we assumed that it reflects the summed activity of two accumulator populations that together represent a single signed quantity of cumulative differential evidence, with one population encoding positive values and the other representing negative values. This scheme is similar, for example, to the manner in which left/right motion direction is encoded by two separate sensory neural populations neither of which can go negative in their firing rate [3]. Note that the application of a half-wave rectification operation on the differential evidence representations is employed here to make explicit the nature of the CPP representation for the purposes of simulation, and we allow that this same computation may of course be implemented differently in practice. Each of these accumulator units is fed into a motor preparation process at which an action-triggering threshold is set, and a common source of evidence-independent, dynamically growing urgency adds to both of these competing motor preparation signals, albeit with independent variability. The traces in this schematic are from a single example trial in the LoCoh regime in which evidence favors the leftward direction and results in a correct response being made at 1.172 sec. We chose a trial in which the evidence was particularly uncertain to aid our explanation. The urgency signal plot shows the average urgency signal component for each of the regimes, and the gray arrows inside it signify that on each trial, the starting value and slope of the urgency signal varies randomly around that mean trajectory. Note that the architecture shown here suggests that the evidence accumulation process reflected in the CPP is an intermediate processing level that receives the evidence directly, and in turn directly feeds its output into the motor level; this has not been empirically verified, and is the focus of ongoing neurophysiology investigations. It is not essential to the signal's use in the neurally-informed modelling approach in this study. In other words, what matters to the signal's use in the NI modelling approach is its well established functional characteristics, not the details of how these characteristics arise from the underlying neural circuitry.

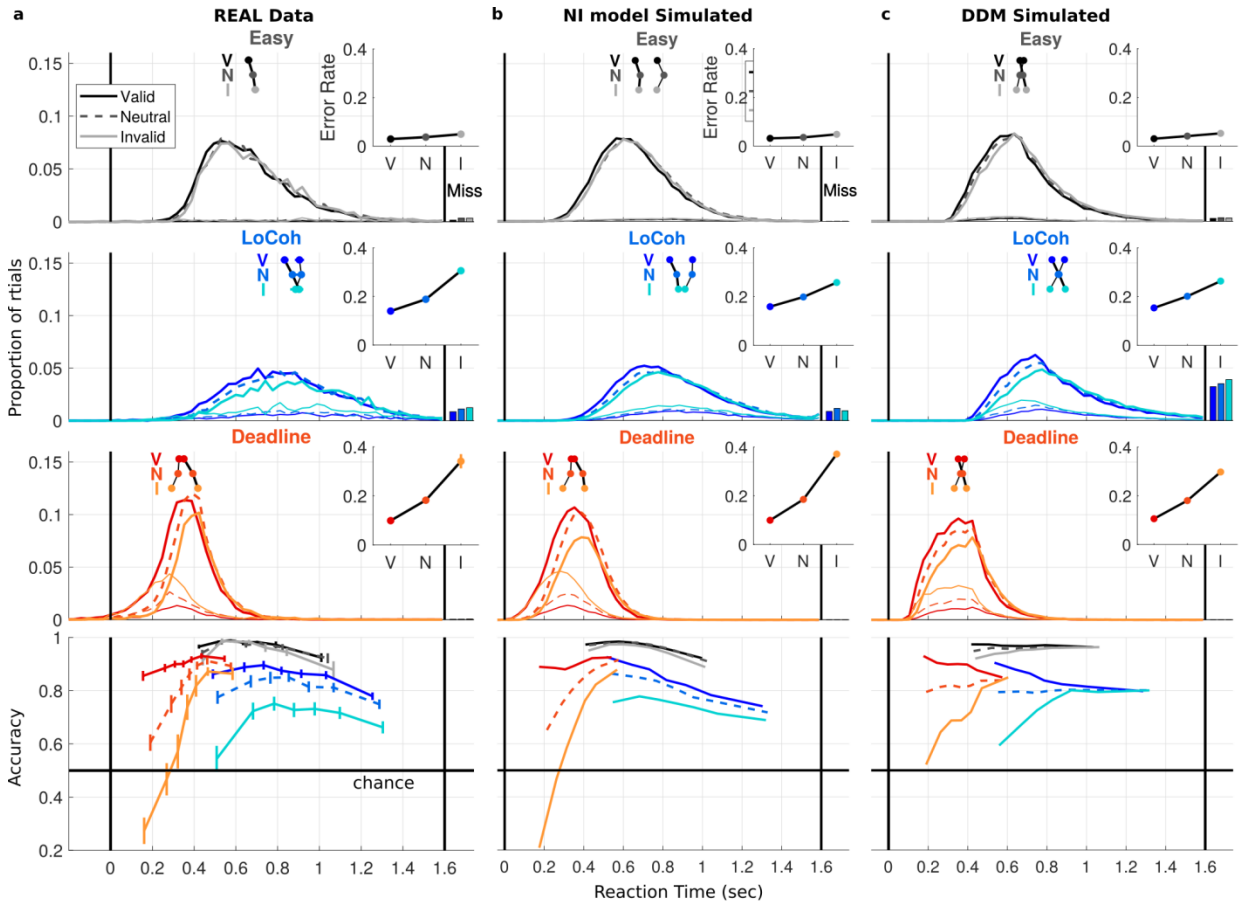

Supplementary Fig. 4: Same as Fig. 3 but showing models with no drift rate biases allowed.

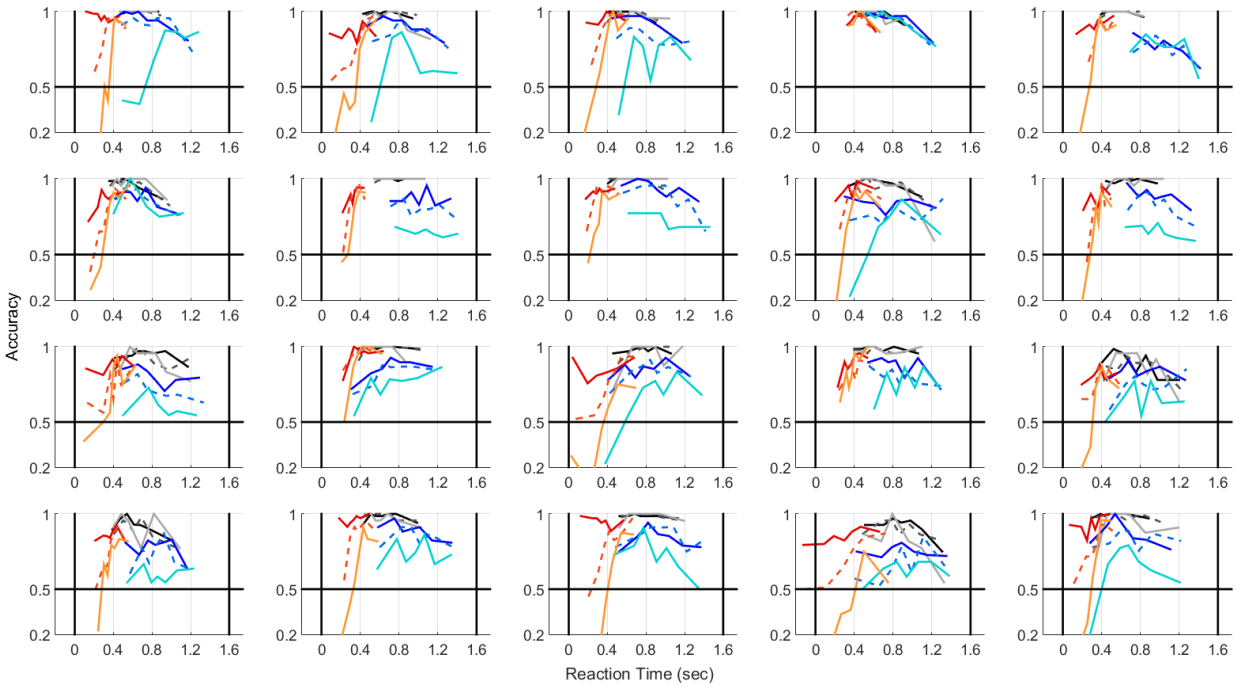

Supplementary Fig. 5: Conditional Accuracy Functions for each of the 20 individual subjects, highlighting consistency with the grand average pattern in terms of the effects of Regime and Prior Probability on RT and accuracy, and the tendency for accuracy to exhibit an overall inverted-U shape as a function of RT, consistent with starting point variability combined with dynamically increasing urgency.

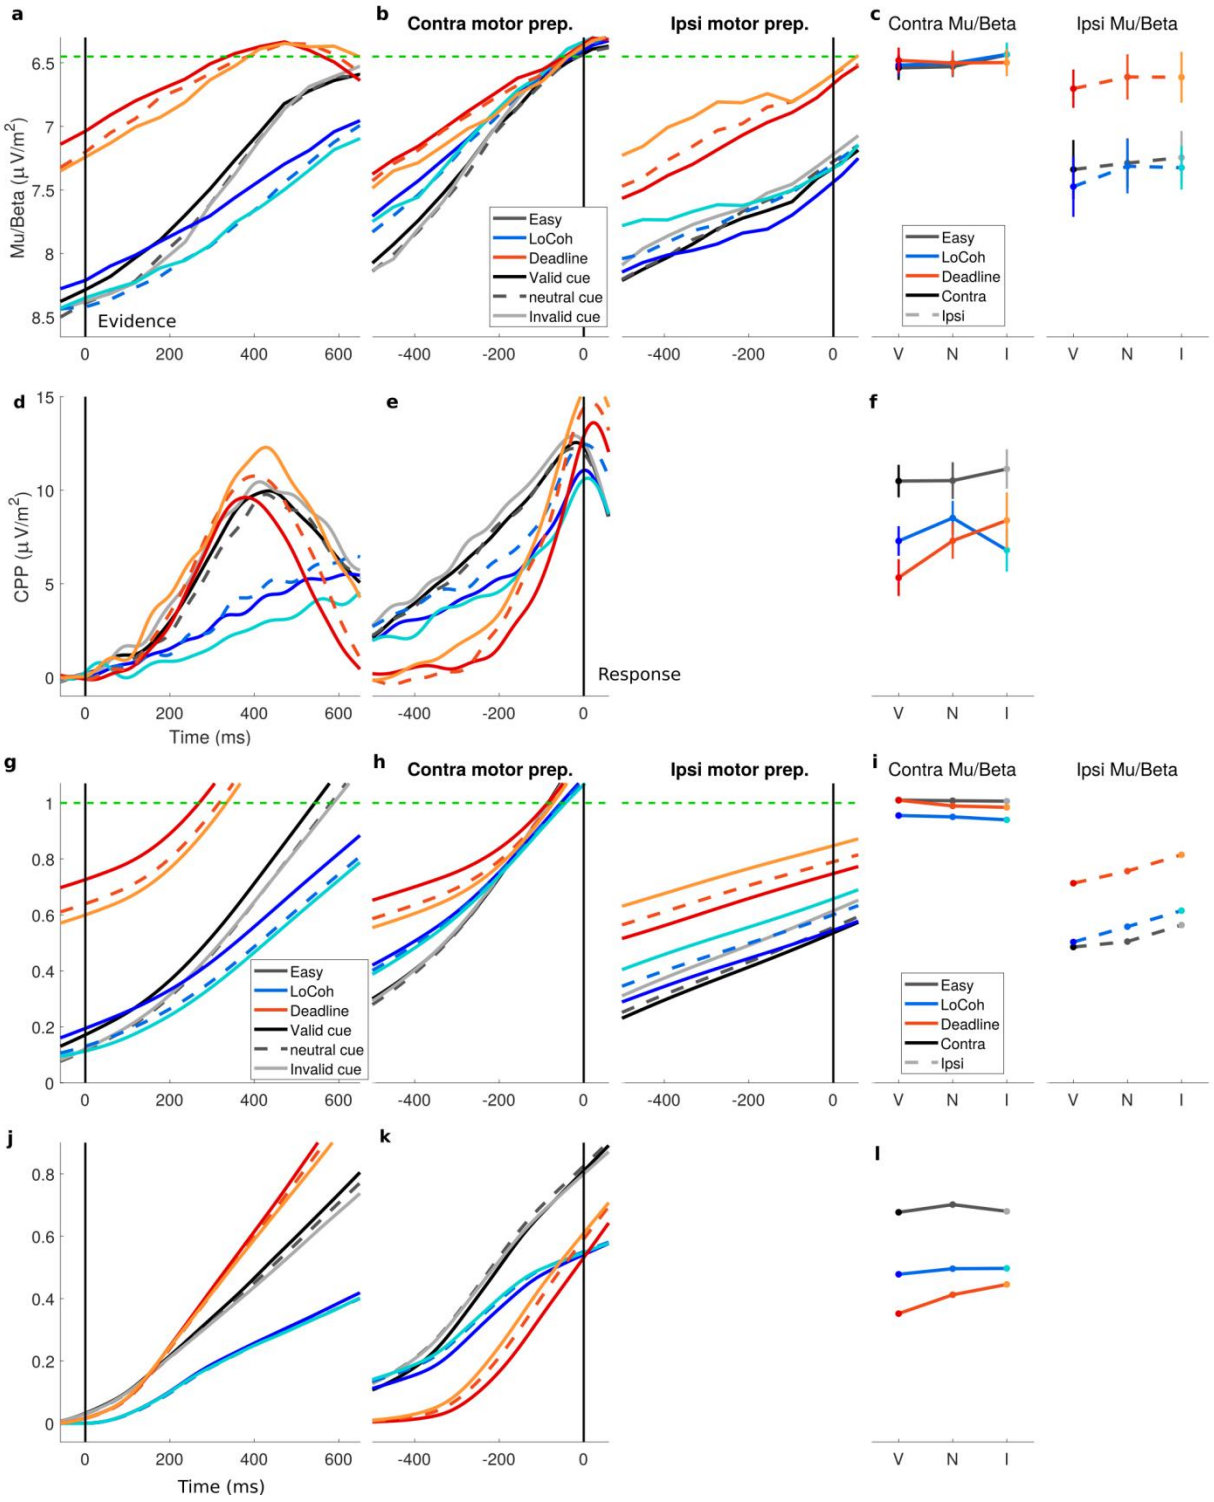

Supplementary Fig. 6: Real and simulated waveforms as in Fig. 5 but with each regime's traces further broken out by prior cue validity. All data are for correct trials only. **a**. Stimulus-locked and **b**. response-locked MB amplitude reflecting motor preparation contralateral (left) and ipsilateral (right) to the (correct) response executed. **c**. Amplitudes of contralateral (left) and ipsilateral (right) MB just prior to response, broken out by prior cue validity. **d**. Stimulus-locked

and **e**. response-locked centro-parietal positivity (CPP) for each regime and validity. **f**. CPP amplitude in a 59-ms window centered -97.5 ms prior to response, broken out by prior cue validity. **g-l**. The same plots are produced for simulated decision signal trajectories from the best fitting NI model.

The simulations demonstrate another instance where an intuitive link between a model parameter and decision signal characteristic does not confer a reliable correspondence in trial-averaged data [4]. Specifically, despite a clear drift rate bias in the model fit in the Deadline regime indicating that evidence is accumulated more steeply when validly than when invalidly cued, the simulated temporal slope of the accumulator signal ranks in the opposite way, due to other contingencies inherent in the model. This further highlights the importance of simulating decision variable dynamics rather than relying on correspondences between the model parameters themselves and the signal features. In this particular case, it demonstrates that neural evidence for the drift rate bias across regimes was better sought in CPP pre-response amplitude (Fig. 7) rather than buildup rate.

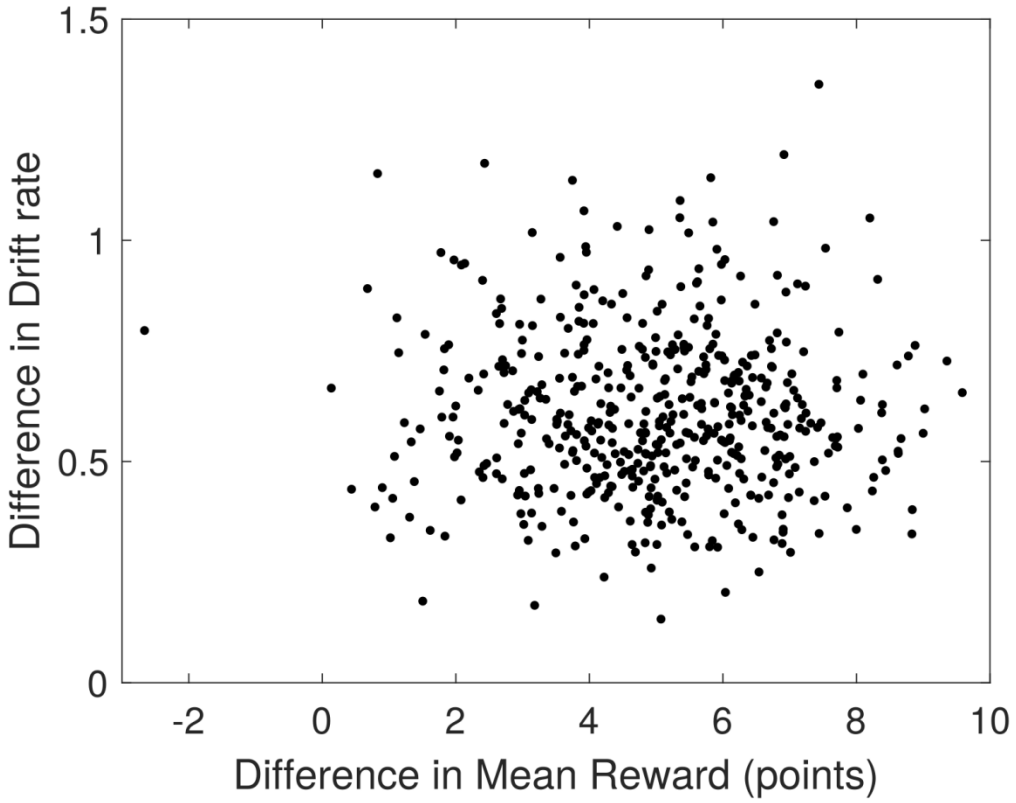

Supplementary Fig 7. Examining the potential relationship between drift rate enhancement and greater expected rewards in the Deadline regime. The size of the increase in drift rate (Deadline regime minus Easy) indicated by the NI model is plotted against the increase in mean reward per trial (Deadline minus Easy) for all bootstrapped-resampled datasets. Recall that in the Easy regime, being correct and on-time earned 40 points per trial, and on average across the 20 subjects, 12% higher rewards per trial were earned under the Deadline regime. Not a single one of the 500 bootstrap samples failed to produce a higher drift rate under the Deadline regime (all differences positive), and there was no significant relationship found between the size of this drift rate boost and the relative increase in expected reward (r=-0.0044, p=0.92, 95% CI = [-0.092, 0.083]).

In addition to this, we asked whether the Deadline regime was inherently more rewarding even without boosting drift rate. By simulating average reward per trial for different adjustments to the estimated model parameters, we determined that if the subjects had performed exactly as in the Easy regime with no adaptation (setting the median deadline of the Deadline regime on the RT distributions simulated in the Easy regime), they would have earned an average of only 8.0 points per trial despite the greater number of points awarded per correct trial. If they had adapted all parameters in the Deadline regime exactly as was found in the NI model fit with the exception of holding the drift rate equal to its setting in the Easy regime, they would have earned an average of 40.6 points per trial, the same as the rewards obtained in practice in the Easy regime. Note also that if the drift rate boost were an automatic

enhancement due to higher expected reward then one might expect all process parameters to be boosted, when in fact urgency rate was decreased under the Deadline regime, not increased. These points taken together suggest that the Deadline regime settings by themselves do not automatically confer increased rewards, and that rather than an automatic processing enhancement due to greater expected rewards, the boost likely represents a strategic adaptation to gain as many rewards as possible. In other words, it is the strategically boosted drift rate that caused the small additional average rewards, rather than the other way around.

References

1. Mazurek, M. E., Roitman, J. D., Ditterich, J. & Shadlen, M. N. A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13,** 1257–1269 (2003).
2. 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.e6 (2018).
3. Dayan, P. & Abbott, L. F. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems (Computational Neuroscience Series)*. (The MIT Press, 2005).
4. Purcell, B. A. & Palmeri, T. J. Relating accumulator model parameters and neural dynamics. *J. Math. Psychol.* **76,** 156–171 (2017).