Dynamic changes in urgency over the course of a decision

Gerard Derosiere[1], David Thura[2], Paul Cisek[3], Julie Duque[1]

[1] *Institute of Neuroscience, Laboratory of Neurophysiology, Université catholique de Louvain,Brussels, Belgium*

[2] *Lyon Neuroscience Research Center – Impact team - Inserm U1028 – CNRS UMR5225 – Lyon 1 University, Bron, France*

[3] *Department of Neuroscience, Université de Montréal, Montréal, QC H3T 1J4, Canada*

Article type: Brief Report

Corresponding author contact details:
Gerard Derosiere
CoActions Lab
Institute of Neuroscience
Université catholique Louvain
Av. Mounier, 53 - Bte B1.53.04
1200 Bruxelles, Belgium
Tel: + 32 (0)2 764 54 20
Email address: gerard.derosiere@uclouvain.be

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

Acknowledgements: This work was supported by grants from the “Fonds Spéciaux de Recherche” (FSR) of the Université Catholique de Louvain, the Belgian National Funds for Scientific Research (FRS-FNRS: MIS F.4512.14) and the “Fondation Médicale Reine Elisabeth” (FMRE). GD was a postdoctoral fellow supported by the FNRS. We thank Simon Van Hemelrijck and Julien Grandjean for their help in the acquisition of the data.

Abstract

While making decisions, humans and other animals always need to balance the desire to gather sensory information (to make the best choice) with the urge to act, facing a speed-accuracy tradeoff (SAT). Given the ubiquity of the SAT across species, extensive research has been devoted to understanding the computational mechanisms allowing its regulation at different timescales, including from one context to another, and from one decision to another. However, in dynamic environments, animals often need to change their SAT on even shorter timescales – *i.e.*, over the course of an ongoing decision – and very little is known about the mechanisms that allow such rapid adaptations. The present study aimed at addressing this issue. Human subjects performed a modified version of the tokens task, where an increase or a decrease in penalty occurring halfway through the trial promoted rapid SAT shifts, favoring speeded decisions either in the early or in the late stage of the trial. Importantly, these shifts were associated with stage-specific adjustments in the accuracy criterion exploited for committing to a choice and relatedly, with dynamic, non-linear changes in urgency. Those subjects who decreased the most their accuracy criterion at a given decision stage presented the highest gain in speed, but also the highest cost in terms of accuracy at that time. Altogether, the current findings offer a unique extension of former work, by revealing that dynamic changes in urgency allow the regulation of the SAT within the timescale of a single decision.

INTRODUCTION

Humans and other animals are motivated to make choices that maximize their reward rate. Paradoxically, while decision accuracy increases the likelihood of getting rewards, the long deliberation time necessary to make accurate choices has a cost that can ultimately reduce the reward rate (Carland et al., 2019). Hence, animals always need to balance the desire to gather sensory information (to make the best choice) with the pressure to act quickly, facing a speed-accuracy tradeoff (SAT; Balci et al., 2011; Bogacz et al., 2010). Given the central role of the SAT in decision-making, extensive research is being devoted to understanding the computational mechanisms at the basis of its regulation (Schall, 2019).

Models of decision-making have since long offered a theoretical account of how the brain may regulate the SAT (Ratcliff, 1985; Stone, 1960; Treisman & Williams, 1984). Traditional models postulate that decision-making involves an accumulation of sensory evidence (Bahl & Engert, 2019; Zylberberg et al., 2016), which drives neural activity up to a fixed level; once this critical threshold is reached, an action is selected (Alamia et al., 2019; Derosiere et al., 2018; Derosiere & Duque, 2020). In this view, to achieve a desired decision policy, the brain controls the height of the threshold, which reflects the accuracy criterion aimed for a given decision. Fast decisions involve low criteria, reducing the amount of evidence required for neural activity to reach the threshold, while longer and accurate deliberations imply higher criteria. Such adaptations were shown to occur both from one SAT context to another (Forstmann et al., 2008; Herz et al., 2016, 2017) and from one decision to another within the same context (Desender et al., 2019; Fischer et al., 2018; Purcell & Kiani, 2016), providing a key mechanism to trade speed with accuracy at different time-scales.

About a decade ago, different studies revealed that the amount of evidence required to commit to a choice decreases over the course of a decision, indicative of an accuracy criterion that wanes as time elapses (*i.e.*, rather than being fixed over time; *e.g.*, Cisek et al., 2009). To explain these data, some authors proposed to incorporate a time-dependent "urgency" signal in decision-making models, which – combined with sensory evidence – pushes neural activity upwards over time, effectively implementing a dropping accuracy criterion (Churchland et al., 2008; Ditterich, 2006; Drugowitsch et al., 2012; Standage et al., 2011; Thura et al., 2012). A given urgency signal is characterized by an initial state and a growing rate, which determine the initial height and the dropping rate of the criterion, respectively, and are thus central to the regulation of the SAT. In situations where speed is of essence, both the initial state (*e.g.*, Steinemann et al., 2018; Thura, 2020; Thura et al., 2014) and the growing rate (*e.g.*, Hanks et al., 2014; Murphy et al., 2016) of urgency are higher, implying a lower initial criterion that quickly decays over time, compared to when the emphasis is on accuracy.

Making decisions in dynamic environments sometimes requires adjusting the SAT on very short timescales – *i.e.*, not only from one context or decision to another but also during an ongoing decision (Gluth et al., 2012). For example, imagine a monkey foraging for fruits in a tree, calmly evaluating which looks tastier when, all of a sudden, a more dominant monkey shows up. In such a scenario, the foraging animal will have to speed up its decision, which may lead it to commit to a choice that does not meet its initially high standards. This situation illustrates how animals sometimes need to quickly change their decision policy and expedite a decision as it unfolds. Yet, very little is known about the computational mechanisms that allow such rapid adjustments.

Here, we address the hypothesis that human subjects can modify their SAT at specific stages of the deliberation process, by dynamically changing their accuracy criterion. As such, while the majority of former studies have conceptualized urgency as a signal that continuously and steadily grows over time within a trial, here we propose that individuals can in fact control the temporal dynamics of this growing signal and thus, of the criterion for committing to a choice. We tested this idea by assessing the behavior of 15 healthy participants in a modified version of the tokens task (Cisek et al., 2009), where penalty changes occurring halfway through the trial promoted rapid SAT shits, either in the early or in the late stage of the decision process.

MATERIALS AND METHODS

Participants

We tested 15 participants for this study (11 women; 24 ± 4.1 years old). All subjects were right-handed according to the Edinburgh Questionnaire (Oldfield, 1971) and had normal or corrected-to-normal vision. None of the participants had any neurological disorder or history of psychiatric illness or drug or alcohol abuse, or were on any drug treatments that could influence performance. Participants were financially compensated for their participation and earned additional money depending on their performance on the task (see below). The protocol was approved by the institutional review board of the Université catholique de Louvain, Brussels, Belgium, and required written informed consent.

Experimental setup

Experiments were conducted in a quiet and dimly lit room. Subjects were seated at a table in front of a 21-inch cathode ray tube computer screen. The display was gamma-corrected and its refresh rate was set at 100 Hz. The computer screen was positioned at a distance of 70 cm from the subject's eyes and was used to display stimuli during a decision-making task. Left and right forearms were placed on the surface of the table with both hands on a keyboard positioned upside-down. Left and right index fingers were located on top of the F12 and F5 keys, respectively (Figure 1.A).

Task

The task used in the current study is a variant of the "tokens task" (Cisek et al., 2009; Thura & Cisek, 2017) and was implemented by means of LabView 8.2 (National Instruments, Austin, TX). The sequence of stimuli is depicted in Figure 1.A. In between trials, subjects were always presented with a default screen consisting of three empty circles (4.5 cm diameter each), placed on a horizontal axis at a distance of 5.25 cm from each other. The central and lateral circles were light blue and dark blue, respectively, and were displayed on a white background for 2500 ms. Each trial started with the appearance of fifteen randomly arranged tokens (0.3 cm diameter) in the central circle. After a delay of 800 ms, the tokens began to jump, one-by-one every 200 ms from the center to one of the two lateral circles (*i.e.*, 15 token jumps; $Jump_1$ to $Jump_{15}$). The subjects were instructed to indicate by a left or right index finger keypress which lateral circle they thought would ultimately receive the majority of the tokens (F12 or F5 key-presses for left or right circle, respectively). They could respond as soon as they felt sufficiently confident, as long as it was after $Jump_1$ had occurred and before $Jump_{15}$. Once a response was provided, the tokens kept jumping every 200 ms until the central circle was empty. At this time, the selected circle was highlighted either in green or in red depending on whether the response was correct or incorrect, respectively, providing the subjects with a feedback of their performance; the feedback also included a numerical score displayed above the central circle (see below, *Reward, penalty and block types* section). In the absence of any response before $Jump_{15}$, the central circle was highlighted in red and a "Time Out" (TO) message appeared on top of the screen, together with a "0" (score) above the central circle. The feedback screen lasted for 500 ms and then disappeared at the same time as the tokens did (the circles always remained on the screen), denoting the end of the trial. Each trial lasted 6600 ms.

One key feature of the tokens task is that it allows one to compute, in each trial, the subject's accuracy criterion, based on the amount of sensory evidence that was available when the subject committed to her/his choice (*i.e.*, at decision time [DT]). To do so, the sum of log-likelihood ratios (SumLogLR) of individual token movements (*i.e.*, a first order estimate of sensory evidence) is usually calculated (Cisek et al., 2009). Based on the temporal profile of the accuracy criterion (*i.e.*, of the SumLogLR at DT), it is then possible to extract an urgency function, characterized by an initial level and a changing rate (*i.e.*, the intercept and the coefficients of the function, respectively). Hence, the tokens task provides us with the possibility to estimate how the accuracy criterion, as well as the initial level and the changing rate of urgency varies from one experimental condition to another. Further details regarding the computation of the accuracy criterion and of the urgency function are provided later, in the *Data analyses* section.

Reward, penalty and block types

As mentioned above, subjects received a feedback score at the end of each trial, which depended on whether they had selected the correct or the incorrect response. Correct responses led to positive scores (*i.e.*, a reward) while incorrect responses led to negative scores (*i.e.*, a penalty). Subjects knew that the sum of these scores would turn into a monetary reward at the end of the experiment.

In correct trials, the reward was equal to the number of tokens remaining in the central circle at the time of the response (in € cents). Hence, the potential reward for a correct response gradually decreased over time (Figure 1.B). For instance, a correct response provided between $Jump_5$ and $Jump_6$ led to a gain of 10 cents (10 tokens remaining in the central circle). However, it only led to a gain of 5 cents when the response was provided between $Jump_{10}$ and $Jump_{11}$ (5 tokens remaining in the central circle). The fact that the reward dropped over time produced an increasing urge to respond over the course of a trial, as evidenced from the urgency functions obtained in such a task (Derosiere et al., 2019).

Incorrect responses led to a negative score but here, the size of this penalty was not linearly proportional to the RT. Importantly, it differed in three block types (see Figure 1.B). The penalty for an incorrect response always equaled 7 cents in the first half of the trial (*i.e.*, up to $Jump_8$), regardless of the block type. However, in the second half of the trial (*i.e.*, after $Jump_8$), it could then either increase to 13 cents ($Penalty_{Increase}$ blocks), remain constant at 7 cents ($Penalty_{Constant}$ blocks) or decrease to 1 cent ($Penalty_{Decrease}$ blocks). The passage from the first half of the trial (called early-stage) to the second half (late-stage) was indicated to the subjects by a change in the color of the central circle, which always turned black at $Jump_8$.

We expected that the penalty shift would induce stage-specific adjustments of the SAT in the $Penalty_{Increase}$ and $Penalty_{Decrease}$ blocks, compared to the $Penalty_{Constant}$ condition. Particularly, in the $Penalty_{Increase}$ blocks, we expected that the prospect of a higher penalty at a late stage would promote faster decisions at the early stage, at the cost of accuracy. Hence, we expected subjects to trade accuracy for speed specifically when making early-stage decisions in the $Penalty_{Increase}$ blocks. Inversely, in the $Penalty_{Decrease}$ blocks, we predicted a tendency to make fast but less accurate decisions at a late-stage of the trial, after the drop in penalty.

Experimental procedure

Subjects performed three experimental sessions (one for each block type) conducted on separate days at a 24-h interval. Testing always occurred at the same time of the day for a given subject, to avoid variations that could be due to changes in chronobiological states (Derosière et al., 2015; Schmidt et al., 2006). The order of the sessions was counterbalanced across participants.

The three sessions always started with two short blocks of a simple RT task (SRT). In this task, subjects were presented with the same display as in the tokens task described above. However here, instead of jumping one by one, the 15 tokens jumped simultaneously into one of the two lateral circles (always the same one in a given block) and subjects were instructed to respond as fast as possible by pressing the appropriate key (*i.e.*, F12 and F5 for left and right circles, respectively). Because the target circle was known in advance of the block, this task did not require any choice to be made and was exploited to determine the subject's mean SRT for left and right index finger responses. We obtained this SRT by computing the difference between the key-press and the time at which the 15 tokens left the central circle (Cisek et al., 2009).

Next, subjects performed training blocks to become acquainted with the tokens task. In a first training block (20 trials, only run on the first session), we ran a version of the tokens task in which the feedback was simplified; the lateral circle turned either green or red, depending on whether subjects had provided a correct or incorrect response; no reward or penalty was provided here. Then, we ran two training blocks (20 trials each) in the condition subjects would be performing next during the whole session ($Penalty_{Increase}$, $Penalty_{Constant}$ or $Penalty_{Decrease}$).

The actual experiment involved 8 blocks of 80 trials (640 trials per session; 1920 trials per subject). Each block lasted about 8.5 minutes and a break of 5 minutes was provided between each of them. Each session lasted approximatively 120 minutes.

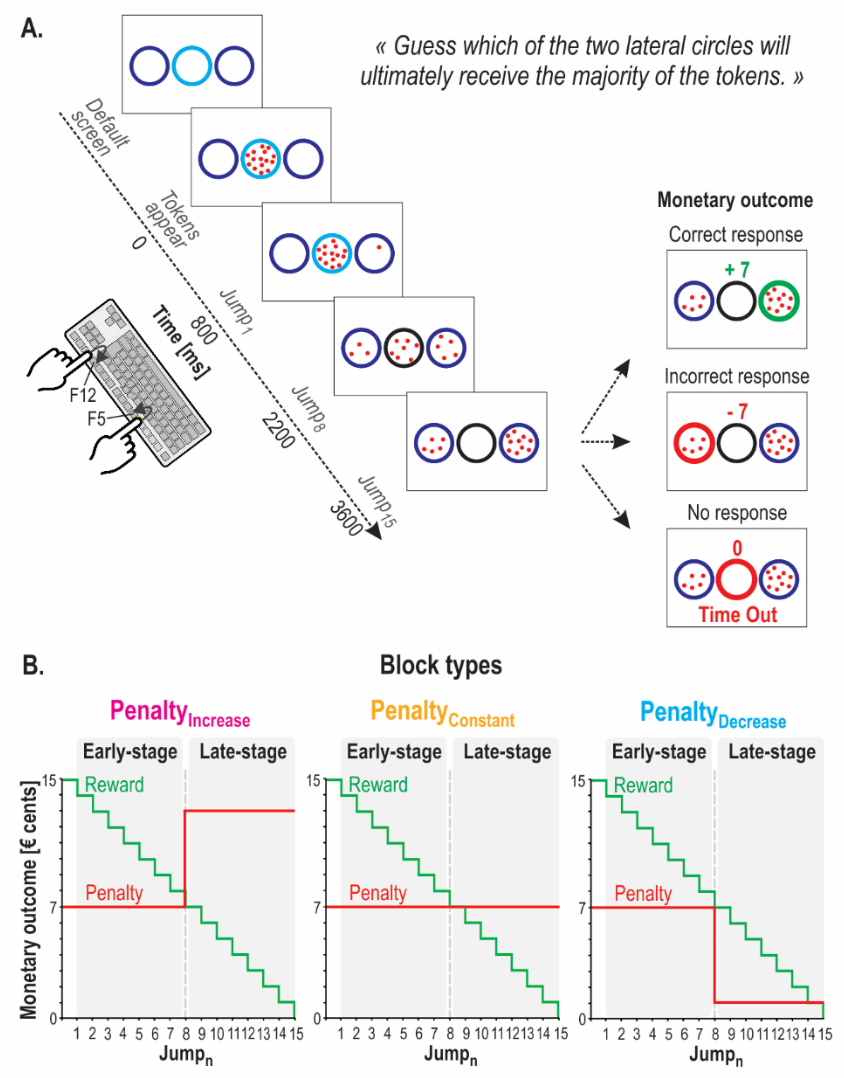

Figure 1: A. Schematic of the tokens task. In each trial, 15 tokens jumped one-by-one every 200 ms from the central circle to one of the lateral circles. The subjects had to indicate by a left or right index finger keypress (*i.e.*, F12 and F5 keys, respectively) which lateral circle they thought would receive more tokens at the end of the trial. For a correct response, the subjects won, in € cents, the number of tokens remaining in the central circle at the time of the response. Hence, the reward earned for a correct response decreased over time, as depicted in B. The example presented on upper inset at the right of panel A represents a correct response provided between Jump_5 and Jump_6 – *i.e.*, the score indicates that 7 tokens remained in the central circle at the moment the right circle was chosen. In contrast, as illustrated on the middle inset of A, subjects lost money if they chose the incorrect lateral circle: they received a negative score that depended on the block type, as indicated in B. In the absence of any response ("Time Out" trial, bottom inset), subjects were neither rewarded, nor penalized (score = 0). For representative purposes, the "Time Out" message is depicted below the circles in this example, while it was presented on top of the screen in the actual experiment. **B. Block types**. Incorrect responses led to a negative score, which differed in three block types. The penalty for an incorrect response always equaled 7 cents in the first half of the trial (*i.e.*, up to Jump_8), regardless of the block type. However, in the second half of the trial (*i.e.*, after Jump_8), it could then either increase to 13 cents ($\text{Penalty}_{\text{Increase}}$ blocks; magenta, left), remain constant at 7 cents ($\text{Penalty}_{\text{Constant}}$ blocks; yellow, center) or decrease to 1 cents ($\text{Penalty}_{\text{Decrease}}$ blocks; blue, right). The passage from the first half of the trial (called early-stage) to the second half (late-stage) was indicated to the subjects by a change in the color of the central circle, which always turned black at Jump_8 (see A).

Data analyses

Data were collected by means of LabView 8.2 (National Instruments, Austin, TX), stored in a database (Microsoft SQL Server 2005, Redmond, WA), and analyzed with custom Matlab scripts (MathWorks, Natick, MA). Detailed methods to analyze behavioral data from the tokens task have been described previously (Thura et al., 2014).

Decision time, accuracy and percentage of time outs

For each block type and each subject, we computed the average decision time (DT) and accuracy (% Correct choices), as well as the percentage of "time out" trials (%TO). To estimate the DT, we first calculated the reaction time (RT) during the tokens task by computing the difference between the time at which the subject pressed the key and the time of $Jump_1$. We then subtracted from this RT the mean simple reaction time (SRT) obtained for each subject. This procedure allowed us to remove from the individual RT obtained in the tokens task, the sum of the delays attributable to sensory processing of the stimulus display as well as to response initiation and muscle contraction, providing us the DT (Cisek et al., 2009; Derosiere et al., 2019).

For the analysis of DT and accuracy data, the dataset was split into two subsets according to whether decisions were made during the early stage (between $Jump_1$ and $Jump_8$; DTs ranging from 0 and 1400 ms) or during the late stage of the trial (between $Jump_8$ and $Jump_{15}$; DTs ranging from 1400 and 2800 ms). This allowed us to test for the effect of the block on the subjects' decision speed and accuracy, separately for responses provided either during the early- or late-stage of the trial. We predicted that, compared to the $Penalty_{Constant}$ condition, subjects' accuracy would be particularly low for responses provided during the early-stage in $Penalty_{Increase}$ blocks and during the late-stage in $Penalty_{Decrease}$ blocks, reflecting a propensity to trade decision accuracy for speed when the penalty is the lowest within a trial.

Accuracy criterion

As mentioned above, the tokens task allows us to estimate the subject's accuracy criterion, based on the amount of evidence that was available for the chosen circle in each trial at DT (*i.e.*, the SumLogLR at DT). As such, high (low) accuracy criteria imply the necessity to accumulate a large (small) amount of evidence before committing to a choice, and thus, high (low) SumLogLR at DT values. The SumLogLR of individual token movements was calculated as follows:

$$SumLogLR(n) = \sum_{k=1}^{n} log\frac{p(e_k|S)}{p(e_k|NS)} \tag{1}$$

Where $p(e_k|S)$ is the likelihood of a token event e_k (a token jumping into either the selected or non-selected lateral circle) during trials in which the selected target *S* is correct, and $p(e_k|NS)$ is its likelihood during trials in which the non-selected circle *NS* is correct. Hence, the SumLogLR at DT is proportional to the difference in the number of tokens that have moved in each direction at the time of commitment. To characterize the changes in accuracy criterion from one block condition to another during the early- and the late-stage of the trial, we split the SumLogLR at DT dataset into two subsets according to whether decisions were made in the former or in the latter stage. In accordance with previous studies (*e.g.*, Cisek et al., 2009; Murphy et al., 2016), we expected that the accuracy criterion would drop as the deadline to respond approached, thus leading to globally lower values in the late- relative to the early-stage of the trial. However, we predicted that, compared to the $\text{Penalty}_{\text{Constant}}$ condition, the criterion would be particularly low during the early-stage in $\text{Penalty}_{\text{Increase}}$ blocks and during the late-stage in $\text{Penalty}_{\text{Decrease}}$ blocks, reflecting the subjects' ability to adjust their criterion to a desired level at specific stages of the decision process.

Estimation of urgency functions

Models of decision-making incorporating an urgency signal posit that choices result from the combination of signals that reflect the available sensory evidence and the level of urgency that grows over time (*e.g.*, Ditterich, 2006; Drugowitsch et al., 2012). For instance, in a minimal implementation of the urgency-gating model (Cisek et al. 2009; Thura et al., 2014), evidence is multiplied by a linearly increasing urgency signal and then compared to the accuracy criterion. The result can be expressed as follows:

$$y_i = \left(N_i - N_{j \neq i}\right) \cdot [mt + b]^+ < AC \tag{2}$$

Where y_i is the "neural activity" for choices to target *i*, N_i is the number of tokens in target *i*, *t* is the number of seconds elapsed since the start of the trial, *m* and *b* are the slope and y-intercept of the urgency signal, and $[\ ]^+$ denotes half-wave rectification (which sets all negative values to zero). When y_i for any target crosses the accuracy criterion *AC*, that target is chosen.

A direct prediction of such urgency-based models is that decisions made with low levels of sensory evidence (*i.e.*, involving low accuracy criteria) should be associated with high levels of urgency and vice versa. That is, one core assumption is that a high urgency should push one to commit to a choice even if evidence for that choice is weak, effectively implementing a low accuracy criterion. Hence, the accuracy criterion values (SumLogLR at DT) can be exploited to estimate the level of urgency at DT (*e.g.*, Thura et al., 2014; Thura, 2020).

Here, we estimated the level of urgency based on the accuracy criterion values obtained for different DTs. We first grouped the trials in bins as a function of the DT, and calculated the average accuracy criterion for each bin. Nine bins were defined, with the first bin including decisions made between 600 and 800 ms, the second bin including decisions made between 800 and 1000 ms, and so on, until the last bin covering the period between 2200 and 2400 ms. The accuracy criterion values preceding 600 ms or following 2400 ms were not considered for this analysis because many subjects did not respond at these times (59.5 ± 0.04 % of the bins were missing values for these very early and very late times). Considering a model in which evidence is multiplied by an urgency signal, we estimated urgency values based on the accuracy criterion obtained at each bin, in each subject and each block condition, as follows:

$$U_{(t,p,s)} = \frac{T}{AC_{(t,p,s)}} \tag{3}$$

Above, *t* is the DT bin, *p* is the penalty condition, *s* is the subject number, *AC* is the accuracy criterion value (*i.e.*, SumLogLR at DT), *T* is a constant representing a fixed threshold (which we fixed to 1), and *U* is the estimated urgency value. We then fitted regression models over the obtained urgency values. A linear and a second-order polynomial model were fitted and the Akaike Information Criterion (AIC) was obtained for each subject and each block condition, allowing us to compare the two models to each other. We predicted that, compared to the $\text{Penalty}_{\text{Constant}}$ condition, urgency would be particularly high during the early-stage in $\text{Penalty}_{\text{Increase}}$ blocks and during the late-stage in $\text{Penalty}_{\text{Decrease}}$ blocks, and that the polynomial model would thus better capture the dynamic changes in urgency in these two block conditions compared to the linear one (*i.e.*, lower AIC values for polynomial fits).

Statistical analyses

Statistica software was used for all analyses (version 7.0, Statsoft, Oklahoma, United-States). The DT, accuracy and SumLogLR at DT data were analyzed using two-way repeated-measure ANOVAs (ANOVA_{RM}) with BLOCK ($\text{Penalty}_{\text{Increase}}$, $\text{Penalty}_{\text{Constant}}$, $\text{Penalty}_{\text{Decrease}}$) and STAGE (early-stage, late-stage) as within-subject factors. The %TO data were analyzed using a one-way ANOVA_{RM} with BLOCK ($\text{Penalty}_{\text{Increase}}$, $\text{Penalty}_{\text{Constant}}$, $\text{Penalty}_{\text{Decrease}}$) as a within-subject factor. Finally, the AIC values obtained from the urgency fits were analyzed using a two-way ANOVA_{RM} with BLOCK ($\text{Penalty}_{\text{Increase}}$, $\text{Penalty}_{\text{Constant}}$, $\text{Penalty}_{\text{Decrease}}$) and MODEL (linear, polynomial) as a within-subject factors. When appropriate, LSD post-hoc tests were used to detect paired differences. Results are presented as mean ± SE.

RESULTS

Decision time, accuracy and %TO

The average DT was not significantly different in the $\text{Penalty}_{\text{Increase}}$, the $\text{Penalty}_{\text{Constant}}$ and the $\text{Penalty}_{\text{Decrease}}$ blocks (1503 ± 32, 1527 ± 25 ms and 1505 ± 19 respectively), as indicated by the absence of main effect of BLOCK ($F_{1,14}$ = 0.73, p = .489). Importantly though, the $ANOVA_{RM}$ revealed a significant BLOCK*STAGE interaction on the DT data ($F_{1,14}$ = 10.04, p = .0005; see Figure 2.A-B). In fact, the DT of responses provided during the early stage was significantly lower in the $Penalty_{Increase}$ (990 ± 43 ms) than in both the $Penalty_{Constant}$ (1061 ± 37 ms; p = .004) and the $Penalty_{Decrease}$ blocks (1063 ± 32 ms; p = .003). Conversely, the DT of responses provided during the late stage was significantly lower in the $Penalty_{Decrease}$ (1946 ± 15 ms) than in both the $Penalty_{Constant}$ (1994 ± 24 ms; p = .046) and the $Penalty_{Increase}$ blocks (2016 ± 30 ms; p = .005). Importantly though, DTs were similar for the $Penalty_{Increase}$ and the $Penalty_{Constant}$ blocks during the late stage (p = .353); DTs were also comparable for the $Penalty_{Decrease}$ and the $Penalty_{Constant}$ blocks during the early-stage (p = .925). These findings indicate that subjects increased their decision speed at very specific stages during the trial in the $Penalty_{Increase}$ and $Penalty_{Decrease}$ blocks: they made faster decisions specifically during the early-stage of the $Penalty_{Increase}$ blocks and during the late-stage of the $Penalty_{Decrease}$ blocks compared to the $Penalty_{Constant}$ block type.

The average accuracy was not significantly different in the $Penalty_{Increase}$, the $Penalty_{Constant}$ and the $Penalty_{Decrease}$ blocks (84.9 ± 1.8, 86.6 ± 1.0 % and 84.6 ± 1.2, respectively; no main effect of BLOCK: $F_{1,14}$ = 1.48, p = .244), but here again, the BLOCK*STAGE interaction was significant ($F_{1,14}$ = 5.83, p = .008; see Figure 2.C-D). As such, responses provided during the early stage were associated with a lower accuracy in the $Penalty_{Increase}$ (82.4 ± 2.5 %) than in both the $Penalty_{Constant}$ (85.7 ± 1.7 %; p = .027) and the $Penalty_{Decrease}$ blocks (85.5 ± 1.8 %; p = .038). Furthermore, responses provided during the late stage were associated with a lower accuracy in $Penalty_{Decrease}$ (83.6 ± 1 %) than in both the $Penalty_{Constant}$ (87.6 ± 0.9 %; p = .009) and the $Penalty_{Increase}$ blocks (87.4 ± 1.4 %; p =.013). Importantly though, accuracy was similar for the $Penalty_{Increase}$ and the $Penalty_{Constant}$ blocks during the late stage of trials (p = .878), while it was comparable for the $Penalty_{Decrease}$ and the $Penalty_{Constant}$ blocks during the early stage (p = .879). Hence, consistent with the DT findings, subjects decreased their decision accuracy at specific stages of the trial in the $Penalty_{Increase}$ and $Penalty_{Decrease}$ blocks: compared to $Penalty_{Constant}$ blocks, they made more errors during the early stage in the $Penalty_{Increase}$ blocks and during the late stage in the $Penalty_{Decrease}$ blocks.

In summary, these behavioral observations indicate that subjects traded decision accuracy for speed specifically at times where the penalty was lowest relative to the rest of the trial (see Figure 2.E-F). In $Penalty_{Increase}$ blocks, the perspective of a rise in penalty promoted faster but less accurate decisions specifically in the first half of the trial, while in $Penalty_{Decrease}$ blocks, the penalty drop promoted faster but less accurate decisions in the second half of the trial.

Consistent with a constant trade-off between decision speed and accuracy, we observed a significant correlation between the block-related shift in DT and in accuracy (R = .66, p < .0001; Figure 2.G), when considering together early-stage (*i.e.*, [100-($Penalty_{Increase}$/$Penalty_{Constant}$*100)]) and late-stage decisions ([100-

($Penalty_{Decrease}$/$Penalty_{Constant}$*100)]). That is, the more subjects favored speed in one block type with respect to the $Penalty_{Constant}$ condition, the more they lost in accuracy in that block, regardless of the stage of the decision process.

Finally, the $ANOVA_{RM}$ revealed a significant effect of BLOCK on the timeout (%TO) data ($F_{2,28} = 15.95$, $p < .0001$; see Figure 2.H). The %TO was indeed higher in the $Penalty_{Increase}$ (7.2 ± 1.3 %) than in both the $Penalty_{Constant}$ (4.1 ± 0.8 %; $p = .004$) and the $Penalty_{Decrease}$ blocks (1.5 ± 0.2 %; $p < .0001$). In addition, it was lower in the $Penalty_{Decrease}$ than in the $Penalty_{Constant}$ blocks ($p = .016$). Hence, the lower the penalty was during the late-stage of the trial, the less the subjects were inclined to be cautious and to avoid responding, consistent with the reduced decision accuracy observed in the late-stage of $Penalty_{Decrease}$ blocks.

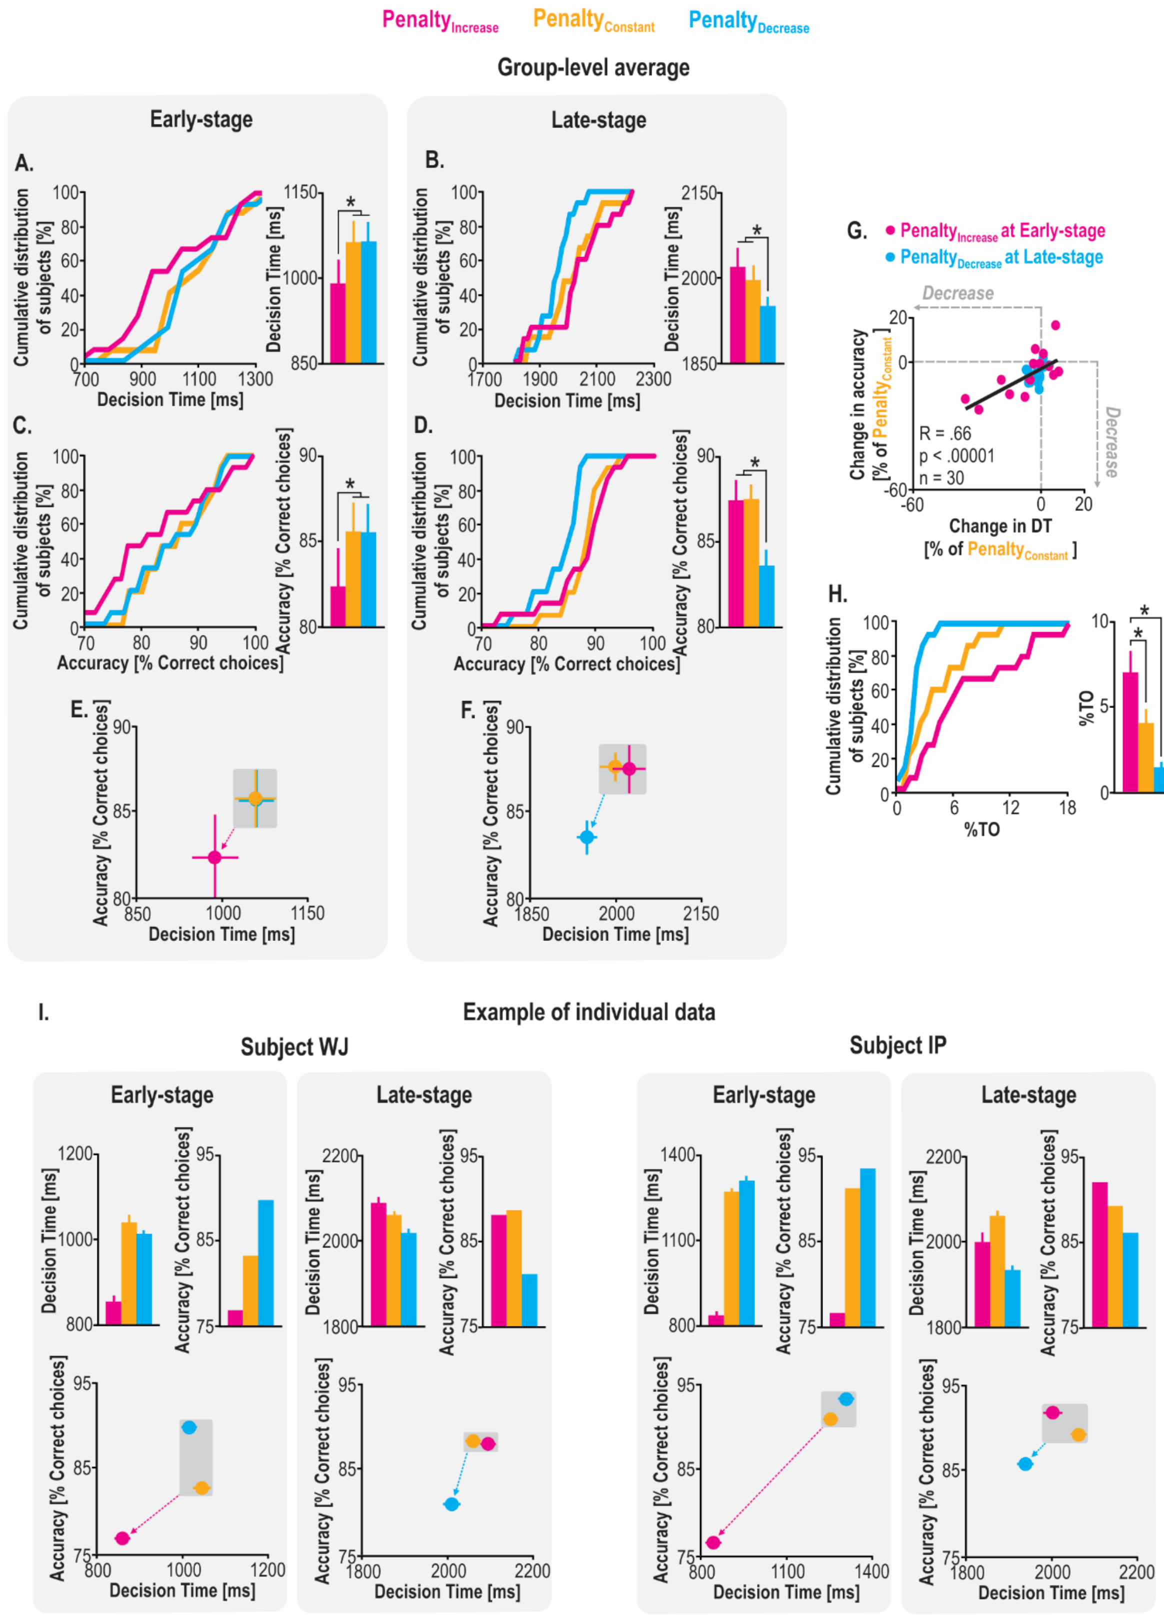

Figure 2: A. Decision Time (DT). Cumulative distribution of subjects and mean DT measured at the early stage of the trial in the Penalty$_{Increase}$ (magenta traces), the Penalty$_{Constant}$ (yellow traces) and the Penalty$_{Decrease}$ blocks (blue traces). **B.** Same as *A.* for the late stage. **C. Accuracy (% of correct choices).** Cumulative distribution of subjects and mean accuracy measured at the early stage of the trial in the Penalty$_{Increase}$ (magenta traces), the Penalty$_{Constant}$

(yellow traces) and the $Penalty_{Decrease}$ blocks (blue traces). **D.** Same as *C.* for the late stage. **E and F. Shift in DT (x-axis) and accuracy (y-axis).** The graphs summarizes the effects represented in A and C, and B and D, respectively. A decrease in DT and in accuracy is observed specifically in the early stage of the $Penalty_{Increase}$ block (magenta dot) and in the late stage of the $Penalty_{Decrease}$ block (blue dot). **G. Relationship between the shift in DT (x-axis) and accuracy (y-axis).** A significant correlation was found between the block-related shift in DT (*i.e.*, [100-($DT_{PenaltyIncrease}$/$DT_{PenaltyConstant}$*100)] for early-stage decisions and [100-($DT_{PenaltyDecrease}$/$DT_{PenaltyConstant}$*100)] for late-stage decisions, magenta and blue circles, respectively) and the block-related shift in accuracy. Both early-stage and late-stage data are shown, leading to an *n* of 30 points. As apparent on the graph, the relationship was especially present for early-stage data. **H. Percentage of time out trials (%TO).** Cumulative distribution of subjects and mean %TO in the $Penalty_{Increase}$ (magenta traces), the $Penalty_{Constant}$ (yellow traces) and the $Penalty_{Decrease}$ (blue traces) blocks. * Between-block difference at p < .05. Error bars represent SE. **I. Example of individual data.** DT and accuracy data are represented for the three block conditions and each stage

Accuracy criterion

The accuracy criterion based on which subjects made their decision was estimated using the SumLogLR at DT: the higher the SumlogLR at DT, the higher the accuracy criterion. Overall, decisions made during the early stage were based on a higher accuracy criterion than those made during the late stage of trials (1.50 ± 0.08 and 1.27 ± 0.07 a.u., respectively), as confirmed by the $ANOVA_{RM}$ showing a main effect of the factor STAGE on the SumLogLR at DT ($F_{1,14} = 76.34$, $p < .0001$). Hence, subjects' accuracy criterion decreased over the course of the decision process, putatively indicating an increasing urge to respond as the central circle was emptying, in agreement with previous studies (*e.g.,* Gluth et al., 2012; Thura and Cisek, 2014, 2017).

Interestingly, the accuracy criterion also depended on the BLOCK under consideration, as revealed by a significant BLOCK*STAGE interaction ($F_{2,28} = 3.46$, $p = .045$). For early-stage decisions, it was lower in the $Penalty_{Increase}$ (1.40 ± 0.13 a.u.) than in the $Penalty_{Constant}$ blocks (1.57 ± 0.09 a.u.; p = .014), with a drop of 10.15 ± 7.56 % (Figure 3.A). For late-stage decisions, the criterion tended to be lower in the $Penalty_{Decrease}$ (0.85 ± 0.02 a.u.) than in the $Penalty_{Constant}$ blocks (0.96 ± 0.02 a.u.; p = .128), with a significant drop of 10.15 ± 2.4 % (*i.e.*, t-test against 0: $t_{14} = 4.23$ $p = .0008$; Figure 3.B). These effects were stage-specific: the criterion was comparable in the $Penalty_{Decrease}$ (1.51 ± 0.13 a.u.) and $Penalty_{Constant}$ blocks for early-stage decisions (1.57 ± 0.09 a.u.; p = .369), as well as in the $Penalty_{Increase}$ (0.97 ± 0.04 a.u.) and $Penalty_{Constant}$ blocks for late-stage decisions (0.96 ± 0.02 a.u.; p = .814). Altogether, these findings indicates that subjects were able to lower their criterion for committing to a choice at specific stages of the trial, in a dynamic way.

These stage-specific adjustments in accuracy criterion appeared to have a significant impact on decision speed and accuracy. Indeed, we observed a positive correlation between the block-related adjustments in criterion and the block-related shift in DT (R = .68, p < .0001; Figure 3.C) as well as in the actual accuracy (R = .93, p < .0001; Figure 3.D), when considering together early-stage (*i.e.*, [100-

($Penalty_{Increase}$/$Penalty_{Constant}$*100)] and late-stage decisions [100-($Penalty_{Decrease}$/$Penalty_{Constant}$*100)]). That is, the subjects who decreased the most their criterion in one block type with respect to the $Penalty_{Constant}$ condition were those who presented the highest gains in decision speed, but also the highest costs in terms of accuracy.

Urgency functions

A direct prediction of urgency-based models is that decisions made with a low accuracy criterion are associated with a high level of urgency and vice versa. Hence, we used the temporal profile of the accuracy criterion, obtained for decisions made between 600 and 2400 ms (presented in Figure 3.D), to estimate urgency functions. Linear and polynomial models were fitted over the rectified SumLogLR at DT values and Akaike Information Criterion (AIC) were obtained for each model (Figure 3.E and F).

Interestingly, the AIC values were lower for polynomial than for linear models (-0.41 ± 1.44 and 0.90 ± 1.5, respectively; ΔAIC = -1.31 ± 0.54), as revealed by a significant effect of the factor MODEL ($F_{1,14}$ = 6.51, p = .029). Thus, on average, the polynomial model better captured the changes in urgency that occurred over the time course of a trial. Importantly though, the superiority of the polynomial model (with respect to the linear one) was not ubiquitous across all block conditions, as suggested by the BLOCK*FIT interaction ($F_{2,28}$ = 2.83, p = .005; Figure 3.G). Indeed, post-hoc tests revealed lower AIC values for polynomial than linear fits for the $Penalty_{Increase}$ data (-0.63 ± 3.17 and 2.45 ± 3.15, respectively; p = .00005; ΔAIC = -3.08 ± 1.02), but not for the $Penalty_{Constant}$ (-2.85 ± 2.32 and -1.84 ± 2.53, respectively, p = .128; ΔAIC = -1.01 ± 0.69) or the $Penalty_{Decrease}$ data (2.26 ± 3.25 and 2.1 ± 3.51, respectively, p = .799; ΔAIC = 0.16 ± 0.43).

In a final analysis, we extracted the urgency value predicted by the polynomial model for DTs of 600 and 2400 ms (*i.e.*, corresponding to decisions made in the early and in the late stage of the trial, respectively). This approach allowed us to confirm the observations made on the accuracy criterion. Indeed, an ANOVA revealed a significant BLOCK*STAGE interaction on the estimated urgency values ($F_{2,28}$ = 8.33, p= .001). At 600 ms, urgency was significantly higher in the $Penalty_{Increase}$ (1.05 ± 0.31 a.u.) than in both the $Penalty_{Constant}$ (0.60 ± 0.09 a.u.; p = .019), and the $Penalty_{Decrease}$ blocks (0.57 ± 0.11 a.u.; p = .013; Figure 3.H), while at 2400 ms, it was higher in the $Penalty_{Decrease}$ (1.71 ± 0.14 a.u.) than in the $Penalty_{Constant}$ (1.36 ± 0.10 a.u.; p = .066) and the $Penalty_{Increase}$ blocks (1.16 ± 0.08 a.u.; p = .004; Figure 3.I). Importantly, the effects were again stage-specific: urgency was comparable in the $Penalty_{Decrease}$ and $Penalty_{Constant}$ blocks for early-stage decisions (p = .873), as well as in the $Penalty_{Increase}$ and $Penalty_{Constant}$ blocks for late-stage decisions (p = .262).

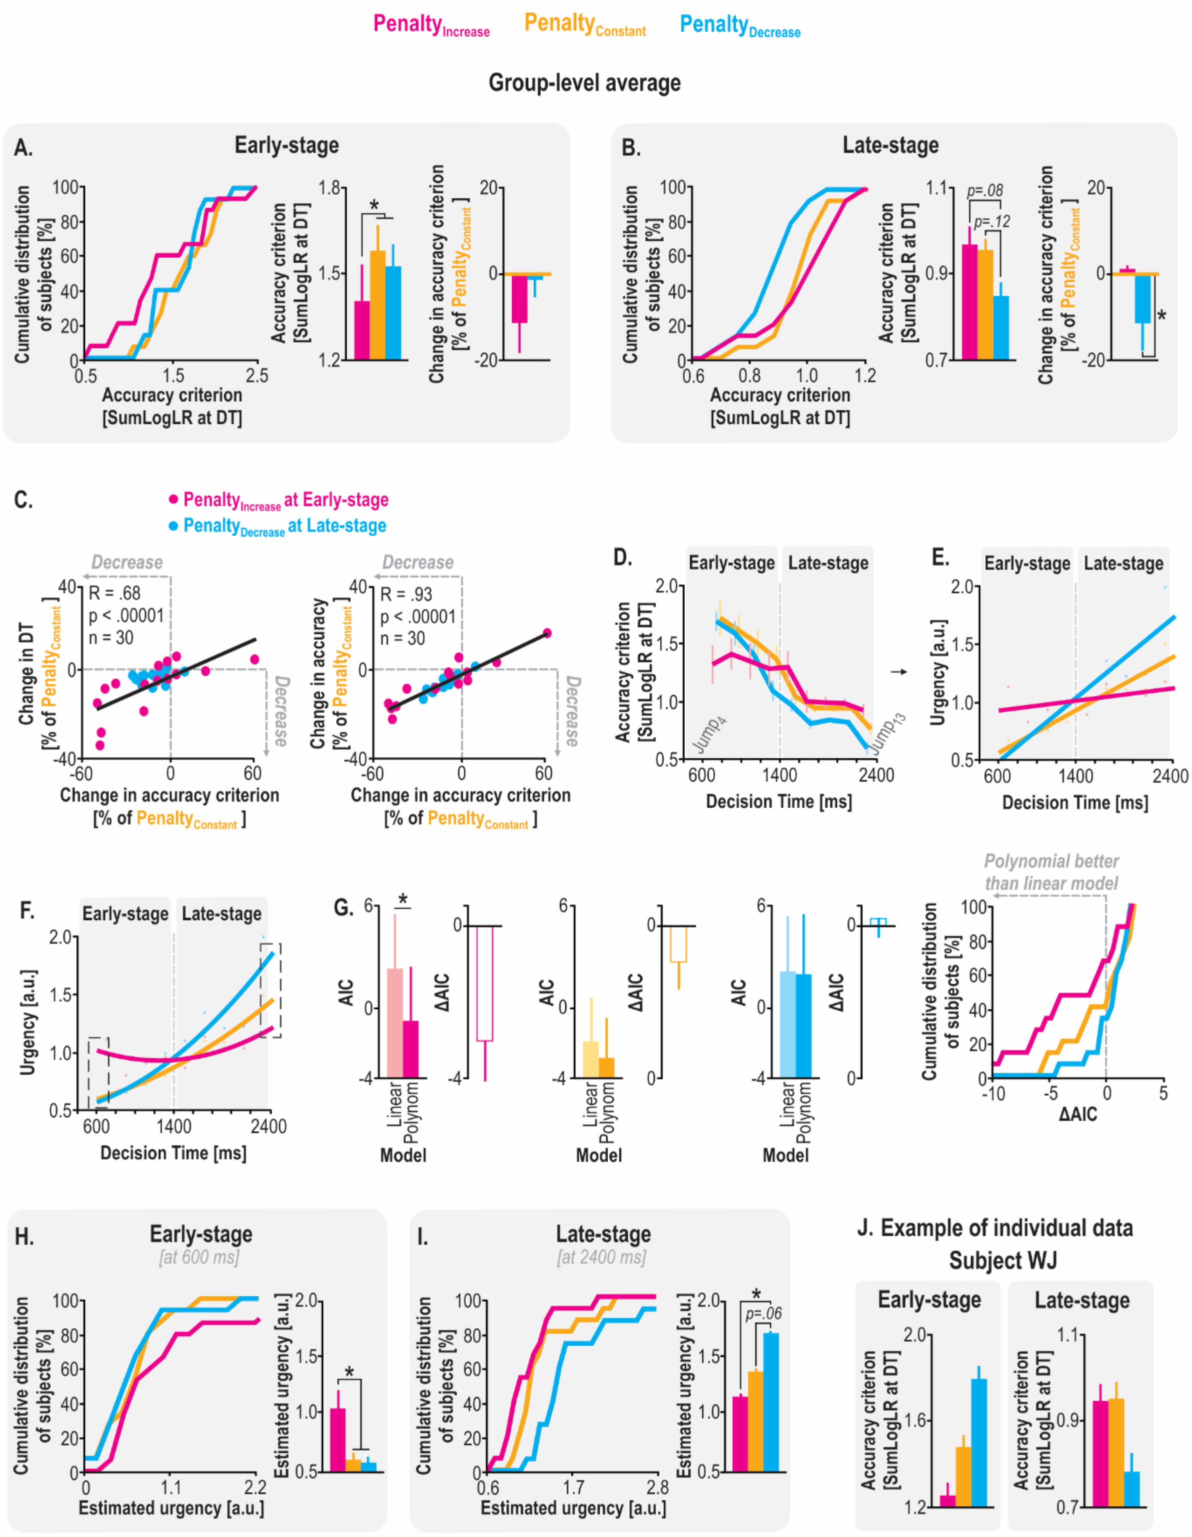

Figure 3: A. Decision Time (DT). Cumulative distribution of subjects and mean accuracy criterion values measured at the early stage of the trial in the $\text{Penalty}_{\text{Increase}}$ (magenta traces), the $\text{Penalty}_{\text{Constant}}$ (yellow traces) and the $\text{Penalty}_{\text{Decrease}}$ blocks (blue traces). **B.** Same as *A.* for the late stage. **C. Relationships between the shift in accuracy criterion (x-axis) and the shift in DT and accuracy (y-axes; left and right, respectively).** A significant positive correlation was found between the block-related adjustments in criterion (*i.e.*, [100-($\text{Criterion}_{\text{PenaltyIncrease}}$/$\text{Criterion}_{\text{PenaltyConstant}}$*100)] for early-stage decisions and [100-

(Criterion$_{PenaltyDecrease}$/Criterion$_{PenaltyConstant}$*100)] for late-stage decisions) and the block-related shift in DT and accuracy. Both early-stage and late-stage data are shown, leading to an *n* of 30 points. **D. Temporal evolution of the accuracy criterion as a function of DT. E and F. Urgency functions, obtained using a linear and polynomial model, respectively. G. Comparison of the Akaike Information Criterion (AIC) for linear and polynomial models.** The graphs show that the AIC value was significantly higher for polynomial fits when applied on the Penalty$_{Increase}$ data (magenta bars), but not for Penalty$_{Constant}$ and Penalty$_{Decrease}$ data (yellow and blue, respectively). The cumulative distribution of subjects obtained for the ΔAIC (*i.e.*, ΔAIC = AIC$_{Polynomial}$ - AIC$_{Linear}$) is presented on the right, highlighting that the polynomial model outperformed the linear one for most of the single-subject data (magenta trace). **H. Cumulative distribution of subjects and mean urgency estimated in the early stage (*i.e.,* at 600 ms) using the polynomial fit. I. Same as H for the late stage** (*i.e.*, estimation made at 2400 ms). * = Between-block significant difference at p < .05. Error bars represent SE. **J. Example of individual data.** Accuracy criterion values data are represented for the three block conditions and each stage.

DISCUSSION

In dynamic environments, humans and other animals often need to change their choice SAT while a decision is ongoing. Yet, very little is known about the computational mechanisms that allow these rapid changes of decision policy. In the present study, we addressed the hypothesis that human subjects can shift their SAT at specific stages of the deliberation process, by dynamically adjusting their accuracy criterion. Participants performed a modified version of the tokens task (Cisek et al., 2009), where an increase or a decrease in penalty occurring halfway through the trial promoted rapid SAT shifts, either in the early or in the late decision stage. Our results reveal that subjects traded accuracy for speed specifically at times where the penalty was the lowest within a trial. Interestingly, these changes were accompanied by stage-specific adjustments in accuracy criterion; in fact, those who decreased the most their criterion presented the highest gains in decision speed, but also the highest costs in terms of accuracy.

Several studies have now revealed the flexibility with which humans can adapt their choice SAT at different time-scales, including from one context to another (Forstmann et al., 2008; Herz et al., 2016, 2017) and from one decision to another (Desender et al., 2019; Fischer et al., 2018; Purcell & Kiani, 2016). The current findings offer a unique extension of this work, by showing that the SAT can be modulated on an even shorter time-scale – *i.e.*, over the course of a single decision. In Penalty$_{Increase}$ blocks, decisions were faster but less accurate in the first half of the trial (*i.e.*, compared to the Penalty$_{Constant}$ condition), while in Penalty$_{Decrease}$ blocks, such SAT shifts occurred in the second half of the trial. The occurrence of a shift in the first half of the trial in Penalty$_{Increase}$ blocks indicates the operation of a proactive, anticipatory process, through which the prospect of a future rise in penalty determined the decision policy to adopt for early-stage decisions. As such, in the current task, subjects likely chose a policy for modifying their SAT before the trial had even started (or before the block of trials). Given that each block (and even each session) always involved the same type of penalty change, subjects could determine what decision policy they should adopt in this specific setting and apply it during deliberation. Whether rapid shifts in SAT can occur reactively (*e.g.*, following online, unpredictable cues) remains an open question, worthy of future investigation.

Individuals' accuracy criterion dropped over time in all block conditions, consistent with the idea of an urgency signal pushing subjects towards commitment as time elapses (Derosiere et al., 2019; Murphy et al., 2016; Thura et al., 2014). Moreover, the temporal dynamics of this drop depended on whether the penalty increased or decreased halfway through the decision process. In the $\text{Penalty}_{\text{Increase}}$ blocks, subjects lowered their accuracy criterion specifically in the early decision stage (*i.e.*, relative to $\text{Penalty}_{\text{Constant}}$ blocks) while in $\text{Penalty}_{\text{Decrease}}$ blocks, they did so in the late decision stage. In fact, the adjustment of the accuracy criterion and of urgency was more pronounced in the early decision stage (*i.e.*, in the $\text{Penalty}_{\text{Increase}}$ relative to the $\text{Penalty}_{\text{Constant}}$ blocks), than in the late one (*i.e.*, in the $\text{Penalty}_{\text{Decrease}}$ relative to the $\text{Penalty}_{\text{Constant}}$ blocks, where differences were marginally significant). This idea is substantiated by the finding that a polynomial model captured more variance of the changes in urgency when considering the $\text{Penalty}_{\text{Increase}}$ data (*i.e.*, compared to a linear model), but not when considering the $\text{Penalty}_{\text{Decrease}}$ data. Hence, participants seemed more effective at adjusting their level of urgency (and, relatedly, their accuracy criterion) for early- compared to late-stage decisions. One possible explanation for this is that urgency was inherently lower for early decisions than for late ones, leaving more room for volitional regulation. Alternatively, it may be the case that the incentive to adjust urgency was stronger in the early stage of $\text{Penalty}_{\text{Increase}}$ blocks than in the late stage of $\text{Penalty}_{\text{Decrease}}$ ones. As such, because of the natural aversion of humans to risk (Weber et al., 2004; Zhang et al., 2014), the prospect of a future rise in penalty might have been more salient than the sudden drop in penalty, thus leading to stronger changes in urgency in the former block condition.

In conclusion, the present study builds on former work on the computational mechanisms underlying the SAT policy. Consistent with past research, we show that the accuracy criterion progressively drops over time during the decision process, in line with an increased urge to commit as the time left to respond diminishes. Most importantly, we provide evidence that rapid shifts in SAT can occur over the course of an ongoing decision and that these changes are related to dynamic adjustments of the accuracy criterion and, relatedly, of urgency. Future work is needed to extend the current observations to situations involving reactive SAT shifts, which may emerge in response to online sensory cues.

REFERENCES

Alamia, A., Zénon, A., VanRullen, R., Duque, J., & Derosiere, G. (2019). Implicit visual cues tune oscillatory motor activity during decision-making. *NeuroImage*, *186*, 424–436. https://www.sciencedirect.com/science/article/pii/S1053811918321050?via%3Dihub

Bahl, A., & Engert, F. (2019). Neural circuits for evidence accumulation and decision making in larval zebrafish. *Nature Neuroscience*. https://doi.org/10.1038/s41593-019-0534-9

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception, and Psychophysics*, *73*(2), 640–657. https://doi.org/10.3758/s13414-010-0049-7

Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology*, *63*(5), 863–891. https://doi.org/10.1080/17470210903091643

Carland, M. A., Thura, D., & Cisek, P. (2019). The Urge to Decide and Act: Implications for Brain Function and Dysfunction. In *Neuroscientist* (Vol. 25, Issue 5, pp. 491–511). https://doi.org/10.1177/1073858419841553

Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, *11*(6), 693–702. https://doi.org/10.1038/nn.2123

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, *29*(37), 11560–11571. https://doi.org/10.1523/JNEUROSCI.1844-09.2009

Derosiere, G., Billot, M., Ward, E. T., & Perrey, S. (2015). Adaptations of motor neural structures' activity to lapses in attention. *Cerebral Cortex*, *25*(1), 66–74. https://doi.org/10.1093/cercor/bht206

Derosiere, G., & Duque, J. (2020). Tuning the Corticospinal System : How Distributed Brain Circuits Shape Human Actions. *The Neuroscientist*. https://doi.org/10.1177/1073858419896751

Derosiere, G., Klein, P. A., Nozaradan, S., Zénon, A., Mouraux, A., & Duque, J. (2018). Visuomotor correlates of conflict expectation in the context of motor decisions. *Journal of Neuroscience*, *38*(44), 9486–9504. https://doi.org/10.1523/JNEUROSCI.0623-18.2018

Derosiere, G., Thura, D., Cisek, P., & Duque, J. (2019). Motor cortex disruption delays motor processes but not deliberation about action choices. *Journal of Neurophysiology*, *122*(4), 1566–1577. https://doi.org/10.1152/jn.00163.2019

Desender, K., Boldt, A., Verguts, T., & Donner, T. H. (2019). Confidence predicts speed-accuracy tradeoff for subsequent decisions. *ELife*, *8*. https://doi.org/10.7554/eLife.43499

Ditterich, J. (2006). Evidence for time-variant decision making. *European Journal of Neuroscience*, *24*(12), 3628–3641. https://doi.org/10.1111/j.1460-9568.2006.05221.x

Drugowitsch, J., Moreno-Bote, R. N., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, *32*(11), 3612–3628. https://doi.org/10.1523/JNEUROSCI.4010-11.2012

Fischer, A. G., Nigbur, R., Klein, T. A., Danielmeier, C., & Ullsperger, M. (2018). Cortical beta power reflects decision dynamics and uncovers multiple facets of post-error adaptation. *Nature Communications*, *9*(1). https://doi.org/10.1038/s41467-018-07456-8

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E. J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(45), 17538–17542. https://doi.org/10.1073/pnas.0805903105

Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding when to decide: Time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *Journal of Neuroscience*, *32*(31), 10686–10698. https://doi.org/10.1523/JNEUROSCI.0727-12.2012

Hanks, T. D., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *ELife*, *2014*(3), 1–17. https://doi.org/10.7554/eLife.02260

Herz, D. M., Tan, H., Brittain, J. S., Fischer, P., Cheeran, B., Green, A. L., Fitzgerald, J., Aziz, T. Z., Ashkan, K., Little, S., Foltynie, T., Limousin, P., Zrinzo, L., Bogacz, R., & Brown, P. (2017). Distinct mechanisms mediate speed-accuracy adjustments in cortico-subthalamic networks. *ELife*, *6*, 1–25. https://doi.org/10.7554/eLife.21481

Herz, D. M., Zavala, B. A., Bogacz, R., & Brown, P. (2016). Neural Correlates of Decision Thresholds in the Human Subthalamic Nucleus. *Current Biology*, *26*(7), 916–920. https://doi.org/10.1016/j.cub.2016.01.051

Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature Communications*, *7*(1), 1–14. https://doi.org/10.1038/ncomms13526

Oldfield, R. C. (1971). The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia*, *9*(1), 97–113. https://doi.org/10.1016/0028-3932(71)90067-4

Purcell, B. A., & Kiani, R. (2016). Neural Mechanisms of Post-error Adjustments of Decision Policy in Parietal Cortex. *Neuron*, *89*(3), 658–671. https://doi.org/10.1016/j.neuron.2015.12.027

Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review*, *92*(2), 212–225. http://www.ncbi.nlm.nih.gov/pubmed/3991839

Schall, J. D. (2019). Accumulators, Neurons, and Response Time. In *Trends in Neurosciences* (Vol. 42, Issue 12, pp. 848–860). Elsevier Ltd. https://doi.org/10.1016/j.tins.2019.10.001

Schmidt, C., Peigneux, P., Muto, V., Schenkel, M., Knoblauch, V., Münch, M., De Quervain, D. J. F., Wirz-Justice, A., & Cajochen, C. (2006). Encoding difficulty promotes postlearning changes in sleep spindle activity during napping. *Journal of Neuroscience*, *26*(35), 8976–8982. https://doi.org/10.1523/JNEUROSCI.2464-06.2006

Standage, D., You, H., Wang, D. H., & Dorris, M. C. (2011). Gain modulation by an urgency signal controls the speed-accuracy trade-off in a network model of a cortical decision circuit. *Frontiers in Computational Neuroscience*, *5*(February), 1–14. https://doi.org/10.3389/fncom.2011.00007

Steinemann, N. A., O'Connell, R. G., & Kelly, S. P. (2018). Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nature Communications*, *9*(1). https://doi.org/10.1038/s41467-018-06117-0

Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, *25*(3), 251–260. https://doi.org/10.1007/BF02289729

Thura, D. (2020). Decision urgency invigorates movement in humans. *Behavioural Brain Research*, *382*, 112477. https://doi.org/10.1016/j.bbr.2020.112477

Thura, D., Beauregard-Racine, J., Fradet, C. W., & Cisek, P. (2012). Decision making by urgency gating: Theory and experimental support. *Journal of Neurophysiology*, *108*(11), 2912–2930. https://doi.org/10.1152/jn.01071.2011

Thura, D., & Cisek, P. (2017). The Basal Ganglia Do Not Select Reach Targets but Control the Urgency of Commitment. *Neuron*, *95*(5), 1160–1170.e5. https://doi.org/10.1016/j.neuron.2017.07.039

Thura, D., Cos, I., Trung, J., & Cisek, P. (2014). Context-dependent urgency influences speed-accuracy trade-offs in decision-making and movement execution. *Journal of Neuroscience*, *34*(49), 16442–16454. https://doi.org/10.1523/JNEUROSCI.0162-14.2014

Treisman, M., & Williams, T. C. (1984). A theory of criterion setting with an application to sequential dependencies. *Psychological Review*, *91*(1), 68–111. https://doi.org/10.1037/0033-295X.91.1.68

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting Risk Sensitivity in Humans and Lower Animals: Risk as Variance or Coefficient of Variation. *Psychological Association*, *111*(2), 430–445. https://doi.org/10.1037/0033-295X.111.2.430

Zhang, R., Brennan, T. J., & Lo, A. W. (2014). The origin of risk aversion. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(50), 17777–17782. https://doi.org/10.1073/pnas.1406755111

Zylberberg, A., Fetsch, C. R., & Shadlen, M. N. (2016). The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *ELife*, *5*(OCTOBER2016). https://doi.org/10.7554/eLife.17688