

When Humans Behave Like Monkeys: Feedback Delays And Extensive Practice Increase The Efficiency Of Speeded Decisions

Nathan J. Evans^{ab} and Guy E. Hawkins^b

^a Department of Psychology, University of Amsterdam, The Netherlands

^b School of Psychology, University of Newcastle, Australia

Keywords:

Comparative cognition — Decision making — Decision strategy — Collapsing thresholds — Cognitive modeling

GEH was supported by an Australian Research Council Discovery Early Career Researcher Award (DE170100177) and Discovery Project (DP180103613).

Correspondence concerning this article may be addressed to: Nathan Evans (nathan.j.evans@uon.edu.au)

Abstract

The study of non-human primates has been foundational in understanding the neural origins of human decision processes, yet the approach rests on the assumption that one can validly extrapolate from the animal to the human. In the context of decision making, this requires constancy across species in physiological *and* cognitive processes. The former cannot be experimentally validated and therefore remains assumed, and recent findings have called into question the latter: non-human primates become increasingly urgent as the time spent making a decision increases, but humans do not; from a normative perspective, monkeys are making closer-to-optimal decisions than humans. Rather than presuming species differences, here we test an alternative hypothesis: previously overlooked differences in methodological procedures from the two research traditions implicitly reinforced fundamentally different decision strategies across the two species. We show that when humans experience decision contexts matched to those experienced by non-human primates – extensive task practice, or time-based penalties – they display increasing levels of urgency as decision time grows longer, in precisely the same manner as non-human primates. Our findings indicate that previously observed differences in decision strategy between humans and non-human primates are eliminated when the decision environment is more closely matched across species, placing a constraint on the interpretation and mapping of neurophysiological results in non-human primates to humans when there are fundamental differences in the task design.

Introduction

The study of non-human primates has been foundational in understanding the neural origins of human decision processes (e.g., Churchland, Kiani, & Shadlen, 2008; Juan, Shorter-Jacobi, & Schall, 2004; Schall & Hanes, 1993; Schall, 1999; Shadlen & Newsome, 1996; Yang & Shadlen, 2007), with the optimality of different decision strategies being a key area of focus within both human and non-human primate decision-making. Within the human decision-making literature, theoretical models of the decision process have traditionally assumed that when people are making decisions, their level of urgency to make a decision remains constant throughout the entirety of the process: that is, they do *not* become increasingly desperate to select an option as more time ticks away. These “constant urgency” strategies are optimal (minimise decision time for a given level of accuracy) in decision environments where the incoming sensory information does not vary across successive decisions (Wald & Wolfowitz, 1948), and have had a wealth of success in explaining human decision-making (Brown & Heathcote, 2008; Ho et al., 2014; Evans & Brown, 2017; Evans, Rae, Bushmakin, Rubin, & Brown, 2017; Hawkins et al., 2014; Ratcliff & Rouder, 1998; Ratcliff, Smith, Brown, & McKoon, 2016). However, the vast majority of decisions, in experimental and real-life situations, have sensory information that fundamentally differs between decisions. In these environments, the optimal decision strategy is to become more urgent as the time spent making the decision grows longer. Such an “increasing urgency” strategy achieves high decision accuracy for those decisions where high quality sensory information is available, and wastes little time on decisions with poorer sensory information (Drugowitsch, Moreno-Bote, Churchland, Shadlen, & Pouget, 2012; Thura, Beauregard-Racine, Fradet, & Cisek, 2012).

Based on these claims of optimality, several recent studies have attempted to compare “constant urgency” and “increasing urgency” accounts of decision-making using the

dominant theoretical modeling framework in the decision-making literature: the diffusion model (Ratcliff, 1978). The diffusion model proposes that incoming sensory information is encoded as “evidence” for each decision alternative, and a decision is triggered when the evidence for one of these alternatives reaches some threshold level of evidence. In this model, a “constant urgency” strategy corresponds to thresholds that remain *fixed* over the course of a decision, and an “increasing urgency” strategy corresponds to thresholds that *dynamically decrease* over the course of a decision, commonly referred to as *collapsing thresholds*. In general, previous quantitative assessments of standard decision-making tasks have found strong evidence in favour of fixed thresholds for human decision-making, even when the quality of sensory information varied across decisions (Evans, Hawkins, Boehm, Wagenmakers, & Brown, 2017; Hawkins, Forstmann, Wagenmakers, Ratcliff, & Brown, 2015; Hawkins, Wagenmakers, Ratcliff, & Brown, 2015; Voskuilen, Ratcliff, & Smith, 2016; though see also van Maanen, Fontanesi, Hawkins, & Forstmann, 2016; Khodadadi, Fakhari, & Busemeyer, 2017; Leimbach et al., 2018; Murphy, Boonstra, & Nieuwenhuis, 2016 for data from expanded judgement tasks that attempt to directly measure the threshold and have found evidence for collapsing thresholds). In contrast, research comparing the fixed and collapsing thresholds theories in non-human primates has found considerable evidence in favour of collapsing thresholds, both qualitatively and quantitatively (Chandrasekaran, Peixoto, Newsome, & Shenoy, 2017; Churchland et al., 2008; Ditterich, 2006a, 2006b; Dru-gowitsch et al., 2012; Hawkins, Forstmann, et al., 2015; Hawkins, Wagenmakers, et al., 2015; Thura & Cisek, 2016). Overall, previous findings suggest humans and non-human primates use different strategies to make speeded decisions, at least under some contexts: humans appear to use the less efficient fixed decision thresholds, and non-human primates appear to use the more efficient collapsing thresholds.

There are a number of potential explanations as to why humans and non-human

primates might differ in their decision strategy, and specifically their changing urgency over the duration of a decision. The first and most obvious hypothesis is that there are genuine differences in default decision strategies across species, which gave rise to the behavioural differences observed in previous research. Generally speaking, previous research has concluded that humans make overly cautious decisions and consequently perform sub-optimally (Evans & Brown, 2017; Evans, Bennett, & Brown, 2018; Rahnev & Denison, 2018; Starns & Ratcliff, 2012), whereas non-human primates typically make very rapid decisions (Hawkins, Forstmann, et al., 2015; Purcell, Schall, Logan, & Palmeri, 2012; Ratcliff, Cherian, & Segraves, 2003; Ratcliff, Hasegawa, Hasegawa, Smith, & Segraves, 2007), highlighting the possibility of cross-species differences in decision-making.

A second hypothesis is that differences in the implementation of ostensibly identical experimental tasks between humans and non-human primates might induce differences in decision strategy, and hence observed behaviour. Human decision-makers typically complete short experimental sessions with little-to-no prior practice, and when they do undertake considerable practice it can influence various parameters of decision-making models (Balci et al., 2011; Dutilh, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2009; Starns & Ratcliff, 2010). In contrast, non-human primates necessarily require large amounts of training to shape their behaviour to a form that allows them to complete an experimental task to a sufficient standard. Furthermore, the standard decision paradigm in human research provides the participant with immediate feedback on each decision, after which the task proceeds to the next trial with only a minimal inter stimulus interval. In contrast, much recent evidence for collapsing thresholds in non-human primates (Ditterich, 2006a, 2006b; Drugowitsch et al., 2012; Hawkins, Forstmann, et al., 2015; Hawkins, Wagenmakers, et al., 2015; Huang & Rao, 2013) originated from the use of data from a single high-profile study (Roitman & Shadlen, 2002) where feedback information, and its associated reward

for correct responses, was withheld until a response was registered or a minimum time post-stimulus onset had elapsed, whichever came *later*. The extended practice afforded to non-human primates may allow them to explore a wider range of decision strategies than is possible for humans to explore in a single experimental session, and the use of a feedback delay may cause the normative advantages of a collapsing threshold to become even more extreme. Indeed, previous research in humans has shown that task demands can influence the overall level of the decision threshold (Evans et al., 2018) and how it changes throughout the course of a decision (Malhotra, Leslie, Ludwig, & Bogacz, 2017), indicating that the structure of the task can influence decision strategies (see also Palestro, Weichart, Sederberg, & Turner, 2018).

Our study aims to provide key insight into which of these two explanations provide the best explanation of previously observed differences between human and non-human primate decision strategies. A great deal of previous behavioural and neurophysiological research has focused on the similarities between human and non-human primate decision-making, and non-human primate research and models have been critical to understanding human cognition and neurophysiology. Therefore, it is crucial to establish whether there are legitimate cross-species differences, or whether there have been unintentional differences in the decision environments examined in previous research that led to the appearance of cross-species differences.

We test these hypotheses in an experiment where human participants experienced a task setup that mirrored the critical components of the non-human primate studies: extended practice in a standard decision-making task, and where different groups of participants completed the task with or without a feedback delay. We performed rigorous quantitative model comparison to compare these theoretical accounts, using the “fixed thresholds” and “collapsing thresholds” models of previous studies (Hawkins, Forstmann,

et al., 2015; Voskuilen et al., 2016), Bayesian hierarchical methods to apply the models to data (Shiffrin, Lee, Kim, & Wagenmakers, 2008), the Deviance Information Criterion (DIC; Spiegelhalter, Best, Carlin, & Van Der Linde, 2002) to compare the relative ability of the theories to account for the data, and the estimated decision thresholds to assess the magnitude of the increase in urgency over time. We note that in no way are we attempting to assess whether the collapsing thresholds diffusion model is the most appropriate model for experimental tasks with feedback delays. Rather, our aim is to assess whether incorporating a feedback delay in an experimental task leads to evidence in favour of collapsing thresholds over fixed thresholds. This would provide a potential explanation for previously observed cross-species differences in decision strategy.

Method

Participants

Twenty one undergraduate students from the University of Newcastle completed the 10 sessions of the experiment online (in coordination with the experimenter), and were reimbursed with a \$25 gift voucher for each session, in accordance with the University of Newcastle Human Research Ethics Committee. Participants were pseudo-randomly split into two groups (to ensure equal sample sizes) that differed in terms of the minimum time that elapsed between stimulus onset and feedback delivery: 11 participants were in the ‘no delay’ feedback condition, and 10 participants were in the ‘1 second delay’ feedback condition.

Task and procedure

Participants completed a random dot motion task (RDM) with a white-noise algorithm (Pilly & Seitz, 2009), which is a widely used experimental paradigm to examine rapid decision-making behaviour and its neurophysiological underpinnings. The RDM task

consisted of a display with a black background and 40 white dots (each 1.5 pixels in radius) positioned within an unmarked circular aperture (75 pixels radius) at the centre of the display. Participants were tasked with choosing whether the dots appeared to move toward the top-left ('z' key) or top-right ('/' key) of the display. On each trial, a proportion of the dots (defined by the "motion coherence") moved in one of the two directions, with the direction randomly chosen on each trial. These coherently moving dots were shifted $\sqrt{18}$ pixels in their direction of motion (3 pixels left/right, 3 pixels up) per frame (66.7ms), where the specific subset of coherently moving dots was re-sampled on every frame; this prevented the coherently moving dots from being too easily detectable. All non-coherently moving dots were repositioned on each frame to a new, randomly sampled x, y coordinate within the circular aperture. At any point throughout a trial, any dot that moved outside of the aperture was randomly replaced within it. Before each trial participants were shown a fixation cross in the centre of the display, with the display time (in milliseconds) sampled from an exponential distribution with mean 700 and an upper truncation of 4,800, and then shifted by 200.

Participants received correct/error feedback following every decision. For all participants there was a 1,000 millisecond inter-stimulus interval, and an additional 1,000 millisecond timeout for errors. For participants in the '1 second delay' condition, there was also a 1,000 millisecond feedback delay: responses faster than 1,000 milliseconds were recorded but the trial did not terminate and provide feedback until 1,000 milliseconds of trial time had elapsed; this mirrors the procedure of Roitman and Shadlen (2002). Participants were not explicitly informed about the feedback delay, meaning that they could not adjust their strategy accordingly a-priori. Rather, they learnt about the feedback delay as they gained experience with the task, which mirrors the reinforcement structure of Roitman and Shadlen (2002). Participants in the 'no delay' condition received immedi-

ate feedback following each response. Essentially, this meant that the task was identical for both conditions after 1,000 milliseconds of decision time had elapsed, with the only difference being the treatment of responses made before 1,000 milliseconds.

Design and data analysis. Our experiment had a single manipulation: the discrimination difficulty, represented by the motion coherence in the RDM (6 levels: 0%, 2.5%, 5%, 10%, 20% and 40% coherently moving dots, each appearing equally often). This manipulation is important because collapsing thresholds are the optimal decision strategy when the rate of evidence accumulation differs randomly from one decision to the next (Drugowitsch et al., 2012; Thura et al., 2012). This is experimentally manipulated by randomising the ease of decisions, motion coherence of the RDM, across trials. Participants completed 432 trials per session (8 blocks of 54 trials each), for a total of 4,320 trials across the 10 sessions.

Before analysis, the first block of trials (i.e., 54 trials) from each session was removed to allow participants to reacclimate to the task. In addition, trials with responses slower than 7,000ms were removed, with approximately 1.7% of trials failing to meet this criterion. Lastly, the first session was excluded for one of the participants in the ‘no delay’ condition as their performance was at chance accuracy.

The exact specifications of the fixed and collapsing thresholds models can be found in the supplementary materials. One important thing to note is that we defined the fixed threshold diffusion model within the “full diffusion” framework (i.e., with random between-trial variability parameters for drift rate, starting point, and non-decision time), and the collapsing thresholds diffusion model within the “simple diffusion” framework (i.e., no random between-trial variability in any parameters). Our reasons for defining the collapsing thresholds model in the simple diffusion framework were twofold: the inclusion of between-trial variability parameters makes the model analytically intractable, and previous studies have claimed that between-trial variability parameters add unnecessary complexity to the

collapsing thresholds model (Ditterich, 2006b; Palmer, Huk, & Shadlen, 2005; Shadlen & Kiani, 2013) that may cause it to lose in model selection (O’Connell, Shadlen, Wong-Lin, & Kelly, 2018). Indeed, when we used conventional model selection methods to compare fixed and collapsing thresholds models that were both endowed with between-trial variability parameters, the collapsing thresholds model performed more poorly than what we report below (see supplementary materials for details¹).

In the main text, we compare the models with the Deviance Information Criterion (DIC; Spiegelhalter et al., 2002), using the minimum deviance as the point estimate for calculating the effective number of parameters. Importantly, unlike many other methods of comparing models (e.g., Akaike Information Criterion [AIC; Akaike, 1974], Bayesian Information Criterion [BIC; Schwarz et al., 1978]), DIC takes into account the complexity of the entire mathematical function of the model rather than reducing the estimate of a model’s complexity to its number of free parameters (see Myung, 2000 and Evans, Howard, Heathcote, & Brown, 2017 for the importance of accounting for functional form complexity). The DIC metric is placed on the deviance scale, meaning that *smaller* values indicate a better model.

In addition to these “model-based” analyses, we also performed a series of analyses on the raw behavioural data. Specifically, we performed 5 default Bayesian ANOVAs (Rouder, Morey, Speckman, & Province, 2012, through the software package JASP; JASP Team, 2018) using the within-subjects factor of session² and the between-subjects factor of feedback delay group, each on a different summary statistic calculated from the data: accuracy, proportion of responses faster than 1 second (i.e., prior to the feedback delay

¹Note that our analyses were based upon the previous developments and methods of several studies (Evans, Submitted; Hawkins, Forstmann, et al., 2015; Holmes, 2015; Ratcliff & Tuerlinckx, 2002; Ratcliff, Van Zandt, & McKoon, 1999; Schwarz et al., 1978; Smith, 2000; Stone, 1960; Turner, Sederberg, Brown, & Steyvers, 2013; Voss & Voss, 2007), which are cited and discussed in the supplementary materials.

²Note that the first session was excluded from the ANOVA analyses owing to the removal of the at-chance participant for this session, which created a “missing data” problem.

time), mean response time, variance in response time, and the non-parametric skew (i.e., $\frac{\text{mean}-\text{median}}{\text{standard deviation}}$) in response time.³ Note that the final 3 summary statistics were calculated and analysed with correctness of the response (i.e., correct/error) as a within-subjects factor, as differences between correct and error response time distributions are often a factor of interest in the response time literature (Ratcliff, 1978; Ratcliff & Rouder, 1998). Most importantly, previous studies have suggested that one of the behavioural characteristics of collapsing thresholds is a reduction in the skew of the response time distribution (Hawkins, Forstmann, et al., 2015; Hawkins, Wagenmakers, et al., 2015; Evans, Hawkins, et al., 2017), meaning that our assessment of how skew differs between feedback delay groups and changes over practice could provide an insightful “model-free” assessment to compliment our model-based analyses. We also obtained the predictions of both models for all 5 summary statistics in an attempt to understand which specific aspects of the data can/cannot be captured by each model.

Results

Model-based analyses

The estimated decision thresholds for the constant urgency, fixed thresholds model (orange, dashed lines) and the decreasing urgency, collapsing thresholds model (blue, solid lines) can be seen in Figure 1. Means of the group-level mean posterior distributions for each model can be seen in the supplementary materials.

The feedback delay had a clear impact on the estimated decision thresholds – and therefore, the change in urgency over decision time – between the two groups. In the no delay to feedback group, the estimated thresholds for the two models were similar though not identical for most experimental sessions, displaying a fixed level of urgency

³To avoid over-complicating the model-free analyses, all 5 summary statistics averaged across the motion coherence factor.

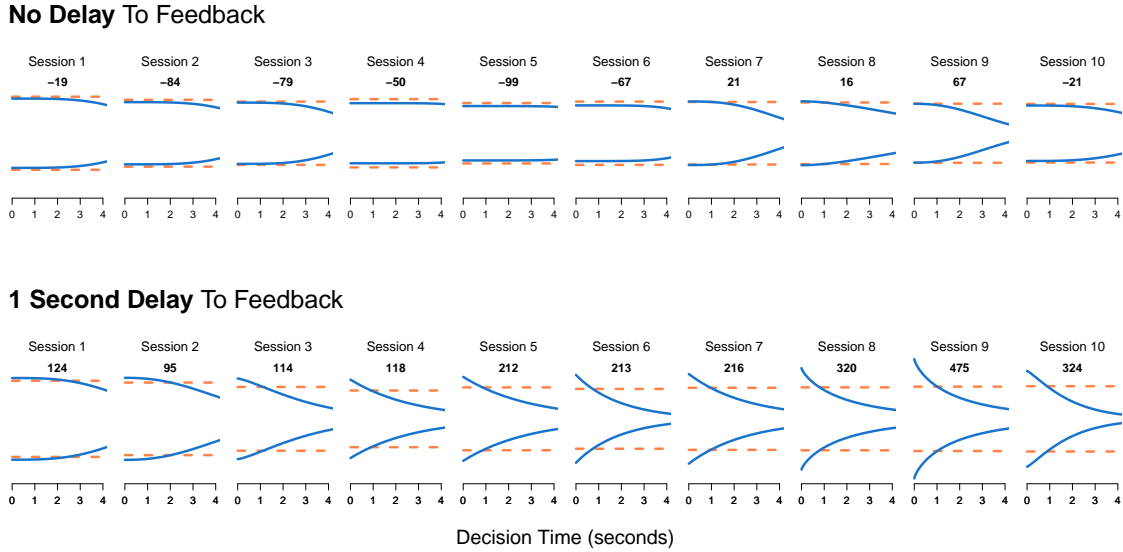


Figure 1. Estimated decision thresholds from the fixed and collapsing thresholds models (orange and blue lines, respectively), where the x-axis displays elapsed decision time and the y-axis displays the amount of evidence required to trigger a decision. There was a drastic difference in the decision strategy used by participants that never experienced a delay between their response and correct/error feedback (upper row) and those that experienced a minimum 1 second delay (lower row), and the severity of this difference increased with time on task (i.e., session number). The DIC difference between the fixed and collapsing thresholds is shown above each panel, where negative values indicate evidence for the fixed thresholds model and positive values indicate evidence for the collapsing thresholds model. Estimated thresholds for each model were calculated using the median of the group-level mean posterior distributions.

over decision time. In contrast, in the delayed feedback group the estimated thresholds for the two models differ a great deal in most sessions, with the collapsing thresholds model indicating a very strong increase in urgency over time (i.e., decreasing thresholds). The DIC values confirm the statistical reliability of these effects, with a strong preference for the fixed thresholds model for 7/10 sessions in the no delay to feedback group, and a strong preference for the collapsing thresholds model in all 10 sessions in the delayed feedback group.

Task practice also had a clear impact on the estimated decision thresholds, especially

for the delayed feedback group, which can be seen in Figure 1. For the no delay group, the estimated thresholds appear to shift from fixed to slightly collapsing over sessions, which is supported by the DIC values (sessions 7, 8, and 9 show strong evidence for the collapsing thresholds model). For the delayed feedback group, the collapse in the estimated thresholds became much more pronounced over practice, beginning with a minor collapse and moving to an extreme collapse. Therefore, extended task practice appears to increase the likelihood of using increasing urgency decision strategies.

Lastly, Figure 2 provides Quantile-Probability (QP) plots for each model (same colours as Figure 1) and the empirical data (black). The Figure shows that both models provide an adequate account of the trends observed in the data for each group and session, suggesting a general descriptive adequacy. This assures us that the best model, according to DIC, is not simply the superior variant of two models that both poorly account for the empirical data. Overall, the models show a small degree of misfit to the slowest error responses (final quantile of the distribution, shown as the uppermost dots/lines of each panel, towards the left). These responses make up a very small proportion of the data, and therefore have very little weight in the overall ability of the models to quantitatively account for the data. Interestingly though, the biggest discrepancy between the predictions of these models appears to be for the slowest errors, which falls in line with previous studies that have suggested that constant urgency and increasing urgency models are most clearly distinguished, both qualitatively and quantitatively, in the tails of the response time distributions (Hawkins, Forstmann, et al., 2015; Hawkins, Wagenmakers, et al., 2015). Therefore, although the final quantile of the error response distributions may have few responses associated with it, and therefore be of little weight in an absolute sense, it may be the key source of quantitative differentiation between these models, and therefore be of large relative weight.

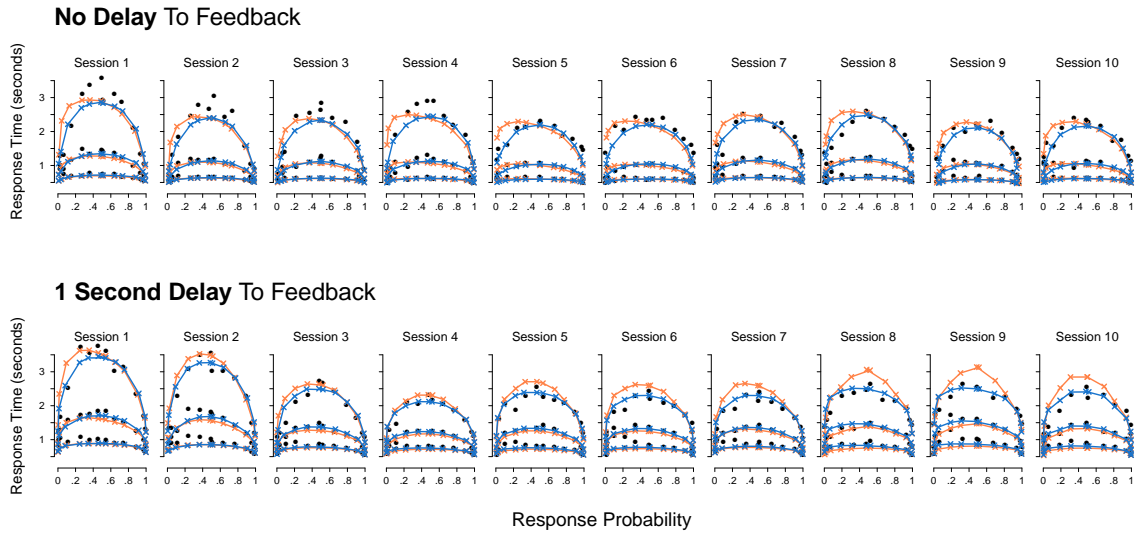


Figure 2. Quantile-Probability (QP) plots for the empirical data (black dots), and the fixed and collapsing thresholds models (orange and blue crosses and lines, respectively). The x-axis plots the probability of a response, and the y-axis plots the associated response time, with the dots and crosses representing the 0.1, 0.5, and 0.9 quantiles for each response (correct/error) of each condition (coherence percentage), with a separate line for each of the 3 different quantiles that connects all responses and conditions. Both models appear to have a general descriptive adequacy for each session of each group, with some minor misfit to the final quantile (0.9) of the error responses (x-axis less than 0.5), which make up a very minor portion of the data. The models also appear to be most clearly distinguished in the slowest responses (0.9 quantile).

Model-free analyses

The results of the Bayesian ANOVAs for the 5 summary statistics calculated from the raw behavioural data can be seen in Table 1. We summarize these analyses in terms of the evidence in favour of the main effect of each experimental factor and the interaction of factors, reflected by the inclusion Bayes factor. For mean response time, there was decisive evidence in favour of an effect of correctness, and weak evidence in favour of an effect of session, suggesting that correct and error mean response times differed, and that mean response time potentially varied over sessions. For the variance of response times,

there was strong evidence in favour of an effect of correctness, and moderate evidence in favour of an effect of session, suggesting that the variance in correct and error response times differed, and that the variance also changed over sessions. For the skew of response times, there was decisive evidence in favour of an effect of correctness, and strong evidence in favour of the effects of feedback delay and the interaction between feedback delay and correctness, suggesting that the skew in correct and error response times differed, and that both the skew, and the difference between correct and error skew, differed between groups. There was no evidence in favour of any main effects or interactions on choice accuracy or the proportion of responses faster than 1 second.

Table 1: Bayes factors in favour of the inclusion of each experimental factor and their interactions (rows) for each of the 5 summary statistics assessed in the raw behavioural data (columns). Bayes factors greater than 1 indicate evidence in favour on an effect, whereas Bayes factors less than 1 indicate evidence in favour of no effect (null hypothesis). Within the table, D = Delay Group, S = Session and C = Correctness (i.e., correct/error responses), and $p(\text{resp.}) < 1\text{s}$ represents the proportion of responses faster than 1 second.

Factor	Mean RT	Variance RT	Skew RT	Accuracy	$p(\text{resp.}) < 1\text{s}$
D	0.307	0.297	52.8	0.408	0.671
S	2	5.96	0.261	0.05	0.018
C	$1.61 \cdot 10^{15}$	40.2	$1.64 \cdot 10^9$	N/A	N/A
D*S	0.167	0.306	0.473	0.008	0.003
D*C	0.313	0.16	14.3	N/A	N/A
S*C	0.032	0.048	0.031	N/A	N/A
D*S*C	$2.07 \cdot 10^{-4}$	$2.98 \cdot 10^{-4}$	0.005	N/A	N/A

The qualitative trends for the 5 summary statistics are shown in Figure 3. For mean response times, correct responses were quicker than errors (as is often the case; Ratcliff, 1978), and there was a tendency toward quicker responses over sessions. Error responses also had greater variance than correct responses, and this variability appeared to decrease over sessions. For the skew in response times, the feedback delay group showed less skew than the no delay group, and error response times showed less skew than correct response

times, with most of the difference between corrects and errors attributable to the large difference for the feedback delay group. There also appeared to be a decrease in skew over sessions for the feedback delay group (especially for errors), though the Bayesian ANOVA suggested there was weak evidence against the interaction between feedback delay group and session for skew. Accuracy and the proportion of responses faster than one second appeared to slightly increase across sessions, though there was strong evidence against both of these effects. Importantly, a large proportion of responses occurred faster than 1 second in the feedback delay group (more than 50%, on average, in most sessions), suggesting that participants were not withholding their responses until the 1 second feedback delay deadline had elapsed.

Figure 3 also shows the predictions of the fixed thresholds (orange) and collapsing thresholds (blue) models for the 5 summary statistics, overlaid on data (black). In general, the fixed thresholds model appears to provide a better account of most trends in most sessions for the no delay group, and in particular outperforming the collapsing thresholds model in the variance in response time. The collapsing thresholds model appears to provide a better account of most trends in most sessions for the feedback delay group, in particular outperforming the fixed thresholds model in the variance and skew in response times. The reason for the increasing advantage of the collapsing thresholds model in the feedback delay group over sessions is clear from these trends: the collapsing thresholds model accurately predicts the decreasing skew in response times, especially in errors, and the fixed thresholds model incorrectly predicts larger and constant skew over sessions. In contrast, in the no delay group it is less clear why the collapsing thresholds model is increasingly favoured in later sessions. It is possible that small advantages across different multiple summary statistics ‘add up’ to provide net evidence in favour of the collapsing thresholds model, or it could be that the collapsing thresholds model provides a better account of performance in

some difficulty levels at later sessions (cf. Figure 2); the difficulty factor is not represented in the summary statistics shown in Figure 3. Nevertheless, the reason for the collapsing thresholds advantage is not as clear as for the feedback delay group.

Discussion

We aimed to identify the reason behind previously reported differences in the decision strategies used by humans and non-human primates. Previous studies had generally suggested that humans' urgency to make a decision remains unchanged, regardless of the duration of the decision, represented as a fixed thresholds diffusion model. In contrast, non-human primates' urgency to make a decision had been found to increase as the time spent making the decision increased, represented as a collapsing thresholds diffusion model, at least in some contexts. We outlined two hypotheses to explain these observed cross-species differences: that there are genuine differences in decision strategy between humans and non-human primates, or that the differences found in previous studies were the result of different methodologies that unintentionally provided participants with different decision environments. We tested these hypotheses in an experimental paradigm where human participants made perceptual decisions while two key factors identified in previous research were manipulated: extended task practice, and the imposition of a minimum duration between stimulus onset and provision of feedback, which we refer to as the feedback delay.

Our investigation of these factors has clear implications for understanding human decision-making. Firstly, our findings indicate that introducing a delay between stimulus onset and the earliest time at which feedback will be provided, ostensibly to discourage very rapid responding, causes people to increase their urgency as decisions take longer. In the first few sessions of the task, participants who experienced a feedback delay adopted collapsing thresholds, and those who experienced no feedback delay had a level of urgency that

remained constant over decision time. These inferences were also supported by model-free, behavioural analyses: response time distributions were less skewed in the feedback delay group, which could not be accounted for by the fixed thresholds model which overpredicted the skew. However, we again note that we are not claiming that the collapsing thresholds diffusion model is the most appropriate model for experimental tasks with feedback delays. Specifically, the use of a feedback delay could result in participants adopting any number of potential strategies that are inconsistent with the assumptions of evidence accumulation models, which would make collapsing thresholds diffusion model a poor model of the process at an absolute level. For example, participants could potentially withhold their responses until after the feedback delay has passed, despite the level of evidence for an alternative having already reached a decision threshold, since responding prior to the feedback delay has no benefit. Importantly, this strategy would be inconsistent with the most basic assumption of evidence accumulation models, and therefore inconsistent with the collapsing thresholds diffusion model, but would likely be better accounted for by a collapsing threshold than a fixed threshold. Rather than attempting to claim which model/s are most appropriate for experimental tasks with feedback delays, we only wish to claim that implementing a feedback delay in an experimental task leads to evidence in favour of collapsing thresholds over fixed thresholds, which has key implications for previous findings in the collapsing thresholds literature.

Secondly, extended task practice appeared to have an impact on decision strategy. The increase in urgency of those in the feedback delay group became more pronounced with practice, and those in the no delay group shifted from a strategy of constant urgency to one of slightly increasing urgency. Importantly, the latter effect did not emerge until the 7th session – over 2500 decisions worth of practice – which is very many decisions per participant in the study of human decision-making, but very few in the study of non-human primates.

However, the effect of practice was not clearly supported by the model-free, behavioural analyses. Specifically, although there appears to be a decreasing trend in skew over sessions in the feedback delay group, especially for errors, Bayes factors found weak evidence against this effect. When considering the model predictions for the summary statistics, there does not appear to be a single, clear reason why the collapsing thresholds model was preferred by DIC in the later sessions for no delay group; it could be a combination of smaller effects. However, for the feedback delay group, the collapsing thresholds model clearly predicts the change in skew across all sessions, whereas the predictions of the fixed thresholds model became further and further from the data in the later sessions. Taken together, it appears that the effect of practice is weaker than the effect of feedback delay. Nevertheless, there appears to be clear evidence that participants behave more consistently with a collapsing thresholds strategy with increased time on task, especially when considering the skew in the feedback delay group.

Our results support the notion that previous cross-species differences in decision strategy were due to differences in experimental procedure. Our findings that implementing a feedback delay, or providing extended practice on the task, pushed human participants toward a decision strategy of increasing urgency suggests that a large proportion of the differences in decision strategy previously observed between humans and non-human primates can be attributed to the experimental procedure. Importantly, all non-human primate studies that have shown evidence for an increasing urgency strategy have used an experimental procedure with extended practice on the task (Hawkins, Forstmann, et al., 2015), and one of the most commonly used datasets to provide evidence for increasing urgency implemented both extended practice and a feedback delay (Roitman & Shadlen, 2002). This result is consistent with previous research that drew a similar conclusion in humans, though using a data set where the duration of task practice differed across par-

ticipants (cf. Palmer et al., 2005 data in Figure 5 of Hawkins, Forstmann, et al., 2015). We wish to be clear that we are not claiming that humans and non-human primates use identical decision strategies. What we do argue is that it appears the specific cross-species differences in decision strategy, and more generally the findings of increasing urgency strategies adopted by organisms, which have both been recently reported in high-profile journals, can be attributed to the experimental procedure.

We believe our findings have three important implications. Firstly, and most importantly, previously reported differences in decision strategy between humans and non-human primates are more likely the result of the experimental procedure than innate differences. Secondly, our findings suggest that feedback delays may have unintended effects on decision processes. One possibility is that feedback delays encourage participants to avoid responding too long after the end of the feedback delay period, as this becomes “wasted time”. This is conceptually similar to tasks with an explicit response deadline, where collapsing thresholds are the most efficient strategy (Frazier & Yu, 2007) and are implemented by humans (Evans, Hawkins, & Brown, Submitted). Although feedback delays are not universally applied in non-human primate research, data sets that have used this particular experimental manipulation (Roitman & Shadlen, 2002) have been over-represented in the literature providing evidence in favour of increasing urgency strategies of decision making (Ditterich, 2006a, 2006b; Drugowitsch et al., 2012; Hawkins, Forstmann, et al., 2015; Hawkins, Wagenmakers, et al., 2015; Huang & Rao, 2013). This is most likely due to Roitman and Shadlen’s commendable efforts in making their data freely available, which is not the case for many research studies on non-human primate decision-making. Therefore, we recommend that 1) the use of a feedback delay be carefully considered by researchers in future, as it may unintentionally alter participants’ decision strategy, and 2) non-human primate researchers aim to make their data freely available where possible, to prevent over-

representation of a subset of data sets in the literature, which has the potential to skew theoretical consensus on important issues. Thirdly, our findings suggest that care must be taken when assessing, and drawing inferences from, data where participants (regardless of species) are highly practiced. Importantly, even in the standard decision-making task we studied here, we found that human participants shifted from the standardly-assumed constant urgency strategy to the more efficient increasing urgency strategy following considerable practice. In future research it will be important to avoid drawing conclusions about general processes from a subset of the data (before or after extended practice).

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, *19*(6), 716–723.
- 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, & Psychophysics*, *73*(2), 640–657.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. *Cognitive psychology*, *57*(3), 153–178.
- Chandrasekaran, C., Peixoto, D., Newsome, W. T., & Shenoy, K. V. (2017). Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nature communications*, *8*(1), 614.
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature neuroscience*, *11*(6), 693–702.
- Ditterich, J. (2006a). Evidence for time-variant decision making. *European Journal of Neuroscience*, *24*(12), 3628–3641.
- Ditterich, J. (2006b). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Networks*, *19*(8), 981–1012.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*, *32*(11), 3612–3628.
- Dutilh, G., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E.-J. (2009). A diffusion model decomposition of the practice effect. *Psychonomic Bulletin & Review*, *16*(6), 1026–1036.
- Evans, N. J. (Submitted). A method, framework, and tutorial for efficiently simulating models of decision-making.
- Evans, N. J., Bennett, A. J., & Brown, S. D. (2018). Optimal or not; depends on the task. *Psychonomic bulletin & review*, 1–8.
- Evans, N. J., & Brown, S. D. (2017). People adopt optimal policies in simple decision-making, after practice and guidance. *Psychonomic Bulletin & Review*, *24*(2), 597–606.

- Evans, N. J., Hawkins, G. E., Boehm, U., Wagenmakers, E.-J., & Brown, S. D. (2017). The computations that support simple decision-making: A comparison between the diffusion and urgency-gating models. *Scientific reports*, *7*, 16433.
- Evans, N. J., Hawkins, G. E., & Brown, S. D. (Submitted). The role of passing time in decision-making.
- Evans, N. J., Howard, Z. L., Heathcote, A., & Brown, S. D. (2017). Model flexibility analysis does not measure the persuasiveness of a fit. *Psychological review*, *124*(3), 339.
- Evans, N. J., Rae, B., Bushmakin, M., Rubin, M., & Brown, S. D. (2017). Need for closure is associated with urgency in perceptual decision-making. *Memory & Cognition*, 1–13.
- Frazier, P. I., & Yu, A. J. (2007). Sequential hypothesis testing under stochastic deadlines. In *Nips* (pp. 465–472).
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *The Journal of Neuroscience*, *35*(6), 2476–2484.
- Hawkins, G. E., Marley, A., Heathcote, A., Flynn, T. N., Louviere, J. J., & Brown, S. D. (2014). Integrating cognitive process and descriptive models of attitudes and preferences. *Cognitive science*, *38*(4), 701–735.
- Hawkins, G. E., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Discriminating evidence accumulation from urgency signals in speeded decision making. *Journal of neurophysiology*, *114*(1), 40–47.
- Ho, T. C., Yang, G., Wu, J., Cassey, P., Brown, S. D., Hoang, N., . . . others (2014). Functional connectivity of negative emotional processing in adolescent depression. *Journal of affective disorders*, *155*, 65–74.
- Holmes, W. R. (2015). A practical guide to the probability density approximation (pda) with improved implementation and error characterization. *Journal of Mathematical Psychology*, *68*, 13–24.
- Huang, Y., & Rao, R. P. (2013). Reward optimization in the primate brain: a probabilistic model of decision making under uncertainty. *PloS one*, *8*(1), e53344.
- JASP Team. (2018). *JASP (Version 0.9)[Computer software]*. Retrieved from <https://jasp>

- stats.org/
- Juan, C.-H., Shorter-Jacobi, S. M., & Schall, J. D. (2004). Dissociation of spatial attention and saccade preparation. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(43), 15541–15544.
- Khodadadi, A., Fakhari, P., & Busemeyer, J. R. (2017). Learning to allocate limited time to decisions with different expected outcomes. *Cognitive psychology*, *95*, 17–49.
- Leimbach, F., Georgiev, D., Litvak, V., Antoniadis, C., Limousin, P., Jahanshahi, M., & Bogacz, R. (2018). Deep brain stimulation of the subthalamic nucleus does not affect the decrease of decision threshold during the choice process when there is no conflict, time pressure, or reward. *Journal of cognitive neuroscience*, *30*(6), 876–884.
- Malhotra, G., Leslie, D. S., Ludwig, C. J., & Bogacz, R. (2017). Overcoming indecision by changing the decision boundary. *Journal of Experimental Psychology: General*, *146*(6), 776.
- Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature communications*, *7*, 13526.
- Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology*, *44*(1), 190–204.
- O’Connell, R. G., Shadlen, M. N., Wong-Lin, K., & Kelly, S. P. (2018). Bridging neural and computational viewpoints on perceptual decision-making. *Trends in Neurosciences*.
- Palestro, J. J., Weichart, E., Sederberg, P. B., & Turner, B. M. (2018). Some task demands induce collapsing bounds: Evidence from a behavioral analysis. *Psychonomic Bulletin & Review*, 1–24.
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of vision*, *5*(5), 1–1.
- Pilly, P. K., & Seitz, A. R. (2009). What a difference a parameter makes: A psychophysical comparison of random dot motion algorithms. *Vision Research*, *49*(13), 1599–1612.
- Purcell, B. A., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2012). From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *Journal of Neuroscience*, *32*(10), 3433–3446.
- Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *The Behavioral*

- and brain sciences*, 1–107.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, *85*(2), 59.
- Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of neurophysiology*, *90*(3), 1392–1407.
- Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of neurophysiology*, *97*(2), 1756–1774.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, *20*(4), 260–281.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic bulletin & review*, *9*(3), 438–481.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological review*, *106*(2), 261.
- Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of neuroscience*, *22*(21), 9475–9489.
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, *56*(5), 356–374.
- Schall, J. D. (1999). Weighing the evidence: how the brain makes a decision. *nature neuroscience*, *2*(2), 108.
- Schall, J. D., & Hanes, D. P. (1993). Neural basis of saccade target selection in frontal eye field during visual search. *Nature*, *366*(6454), 467.
- Schwarz, G., et al. (1978). Estimating the dimension of a model. *The annals of statistics*, *6*(2), 461–464.
- Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, *80*(3),

- 791–806.
- Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the national academy of sciences*, *93*(2), 628–633.
- Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical bayesian methods. *Cognitive Science*, *32*(8), 1248–1284.
- Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, *44*, 408–463.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(4), 583–639.
- Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and aging*, *25*(2), 377.
- Starns, J. J., & Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychonomic bulletin & review*, *19*(1), 139–145.
- Stone, M. (1960). Models for choice–reaction time. *Psychometrika*, *25*, 251–260.
- 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.
- Thura, D., & Cisek, P. (2016). Modulation of premotor and primary motor cortical activity during volitional adjustments of speed-accuracy trade-offs. *Journal of Neuroscience*, *36*(3), 938–956.
- Turner, B. M., Sederberg, P. B., Brown, S. D., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological methods*, *18*(3), 368.
- van Maanen, L., Fontanesi, L., Hawkins, G. E., & Forstmann, B. U. (2016). Striatal activation reflects urgency in perceptual decision making. *NeuroImage*, *139*, 294–303.
- Voskuilen, C., Ratcliff, R., & Smith, P. L. (2016). Comparing fixed and collapsing boundary versions of the diffusion model. *Journal of Mathematical Psychology*, *73*, 59–79.
- Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods*, *39*(4), 767–775.

Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 326–339.

Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, 447(7148), 1075.

Supplementary Material

All data for this article can be found at <https://osf.io/2vnam/>.

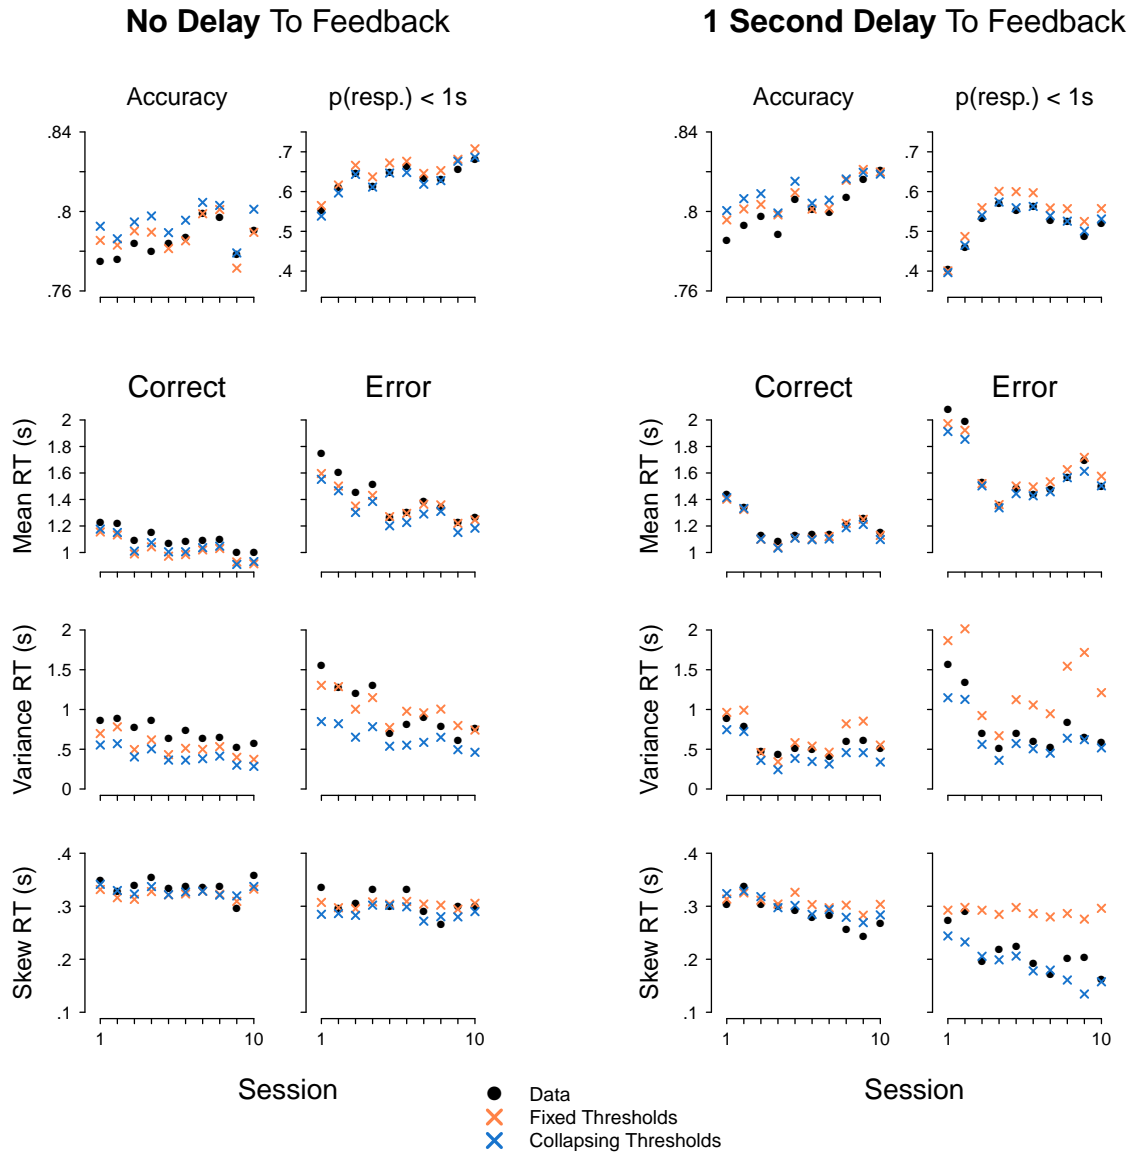


Figure 3. Summary statistics for the empirical data (black) and the fixed and collapsing thresholds models (orange and blue crosses, respectively), separately for the no delay to feedback and 1 second delay feedback conditions (two leftmost and two rightmost columns, respectively). Accuracy and the proportion of responses faster than 1 second ($p(\text{resp.}) < 1\text{s}$) are shown in the upper row. Response time mean, variance and skew are shown in rows 2-4, separately for correct and error responses. The x-axes show sessions 1-10.