A cross-species framework for investigating perceptual evidence accumulation

Sucheta Chakravarty[1]*, Cristina Delgado-Sallent[2]*, Gary A. Kane[2], Hongjie Xia[3], Quan H. Do[3], Ryan A. Senne[4], and Benjamin B. Scott[1,2,5]

[1]Psychological & Brain Sciences, [2]Center for Systems Neuroscience, [3]Graduate Program in Neuroscience, [4]Department of Biomedical Engineering, Boston University
Boston, MA, USA.

Author Note

*These authors contributed equally.

[5]Corresponding author: Benjamin B. Scott (bbs@bu.edu), Department of Psychological and Brain Sciences, Boston University, Boston, MA, 02215, USA, Tel: +1.617.353.2580.

Highlights

1. Development of a free response evidence accumulation task for rats and mice
2. Synchronized video game allows direct comparisons with humans
3. Rat, mouse and human behavior are well fit by the same decision models
4. Model parameters reveal species-specific priorities in accumulation strategy

Summary

Cross-species studies are important for a comprehensive understanding of brain functions. However, direct quantitative comparison of behaviors across species presents a significant challenge. To enable such comparisons in perceptual decision-making, we developed a synchronized evidence accumulation task for rodents and humans, by aligning mechanics, stimuli, and training. Rats, mice and humans readily learned the task and exhibited qualitatively similar performance. Quantitative model comparison revealed that all three species employed an evidence accumulation strategy, but differed in speed, accuracy, and key decision parameters. Human performance prioritized accuracy, whereas rodent performance was limited by internal time-pressure. Rats optimized reward rate, while mice appeared to switch between evidence accumulation and other strategies trial-to-trial. Together, these results reveal striking similarities and species-specific priorities in decision-making. Furthermore, the synchronized behavioral framework we present may facilitate future studies involving cross-species comparisons, such as evaluating the face validity of animal models of neuropsychiatric disorders.

Keywords: perception; evidence accumulation; cross-species; computational modeling

Abbreviations: DDM: drift diffusion model; GAMM: generalized additive mixture model; GLM: generalized linear model; GLM-HMM: generalized linear model-hidden markov model; ITI: inter-trial interval; RT: response time

A cross-species framework for investigating perceptual evidence accumulation

Introduction

Quantitative comparisons of behavior across species are crucial to understand complex cognitive processes (Badre et al., 2015; Esteves et al., 2021; Schmack et al., 2021; Young, 2023). However, such comparisons are challenging for several reasons (Barron et al., 2020; Redish et al., 2021). Research in individual laboratories commonly focuses on single species, leading to marked differences in experimental techniques across species (Badre et al., 2015). These experiments often exploit the innate abilities of individual species. For example, humans are commonly provided with verbal task instructions, which is not possible with other animals. The disconnect between human and animal research can be costly, since animal models serve an important role in preclinical research of human neuropsychiatric disorders (Markou et al., 2009; Nestler & Hyman, 2010).

One domain where significant progress in cross species analysis has been made is perceptual decision making (Brunton et al., 2013; Nguyen & Reinagel, 2022; Pedrosa et al., 2023; Schmack et al., 2021; Shevinsky & Reinagel, 2019). Perceptual decision making is a core behavior observed in many species, and alterations have been observed across a range of human neuropsychological disorders (Horga & Abi-Dargham, 2019; Kavcic et al., 2011; Rizzo & Nawrot, 1998; Robertson & Baron-Cohen, 2017; van den Boogert et al., 2022). Previous research has investigated perceptual decisions in humans, non-human primates, and rodents, using similar stimuli (Carandini & Churchland, 2013; Hanks & Summerfield, 2017; Shadlen & Kiani, 2013). However, direct comparisons of cross-species behavior based on previous studies is difficult, due to marked differences in the paradigms, sensory statistics, and training protocols.

Here we present a behavioral framework for investigating perceptual decision making in mice, rats and humans. We implemented a free response version of the pulse-based evidence accumulation task, previously used in rats (Brunton et al., 2013; Gupta et al., 2024; Scott et al., 2015) and mice (Odoemene et al., 2018; Pinto et al., 2018). Briefly, the task uses sequences of brief visual pulses (flashes), presented to the left and right sides while the subject must choose the side with the higher pulse probability to obtain the reward. This task offers a speed-accuracy trade off as longer response times allow the observer to sample more pulses, which enables better estimates of the relative pulse probability. We trained rats and mice to perform this task using computer controlled 3-port light chambers, and an automated training facility. In parallel, we developed a video game based on the same task to gather data from human participants online.

We synchronized task parameters and training protocols between the rodent task and the human video game; we used non-verbal, reward feedback-driven training in all three species. Further, the high-throughput rodent training facility, and the online game facilitated large-scale data collection.

All three species learned to perform the task and relied on evidence accumulation as a choice strategy. However, there were cross-species differences in several key aspects of the accumulation process. Overall, humans were slower and more accurate than rodents, particularly mice. Drift diffusion models (DDM) revealed highest decision thresholds in humans, and lowest thresholds in mice. Further, collapsing boundary model fits indicated that rodent behavior may be limited by internal time-pressures. Interestingly, rats appeared to optimize response times to maximize reward rate, whereas humans opted for better accuracy. Mice showed high animal-to-animal and trial-to-trial variability in choice behavior and model fits suggested they alternated between evidence accumulation- and other strategies across trials. Together these data reveal important similarities and species-specific priorities in decision making across rodents and humans, and highlight the value of synchronized behavioral frameworks for cross-species studies.

Results

A synchronized framework for cross-species studies of perceptual decision making

We developed a free-response version of the pulse-based evidence accumulation task (Brunton et al., 2013; Scott et al., 2015) that was synchronized across rodents and humans (Figure 1A; see methods for details). In this task, sensory information was presented as a sequence of randomly-timed pulses of light from two sources, one to the left and one to the right. Pulses were brief ($10\ ms$) and binned into $100\ ms$ bins—the odds were set so that the probability (p) of a flash on one side in a bin was complementary to the flash probability on the other side for the same bin ($1-p$). The pulses continued until the subject selected one of the light sources. Identification of the light source with the greater probability of pulsing yielded a correct response.

Rodents performed the task in a three port operant chamber; they initiated each trial with a nose poke at the center port, followed by a cue period during which sequences of brief light flashes were presented in the left and right light ports, the flashes continued until the animal poked on one of the side ports. Rodents were rewarded with a drop of sugar water for poking on the side with greater flash probability (Figure 1A, C). For human experiments, we designed an online video game that preserved the same mechanics and stimulus statistics (flash duration, flash rate, and generative flash probability) from the rodent task (Figure 1A). In the video game, participants were

presented with a field of asteroids. Each trial started when they mouse-clicked on an asteroid. During the cue period, sequences of brief flashes in the form of an alien spaceship were presented bilaterally. The flashes continued until the participant clicked on one of the sides. If they chose the side with the greater flash probability, the initially selected asteroid would be destroyed (Figure 1A-B). Importantly, both the rodent task and the human video game used the same stimulus statistics and both offered a speed-accuracy tradeoff. Longer response times facilitated greater probability of choosing the correct side, due to greater opportunity for accumulating evidence from the flashes.

Both the rodent task and the human video game used a non-verbal, feedback-based training pipeline (Do et al., 2023; Scott et al., 2015). This pipeline consisted of progressive stages with increasing trial difficulty, which was manipulated by altering the generative flash probability of the two sides ($p = 100\%$, $p = 90\%$, and $p = 80\%$). Correct choices were rewarded with positive feedback—a drop of sugar water for the rodents and point bonuses for the humans, for each correct choice. Importantly, humans did not receive verbal instructions about the game rule. Rodents underwent multiple sessions of training over the course of many days, while humans completed two sessions. (Figure 1D-F). Training took $4 - 5$ weeks for mice and $1 - 3$ weeks for rats, and several minutes for humans.

Rodents and humans solved the task using an evidence accumulation strategy

Using the procedures describe above, we trained 21 rats, 95 mice and 18 adolescent humans on the task. All three species performed well on difficult trials ($p = 80\%$) and were above criteria for chance performance (upper bound of the 99% binomial confidence interval, see methods). There was substantial heterogeneity in performance across individuals of a species, revealed by their overall accuracy, response time (RT), bias and reward rate (Figure 2A-D). On average, humans were slower and more accurate, mice were the fastest and least accurate (One-way ANOVA for accuracy: $p < 0.0005, F = 271.84$; for RT: $p < 0.0005, F = 236.39$, see Figure 2A-B, E). Interestingly, there was overlap in performance across species (e.g. Figure 2A-2C). There were no detectable differences between males and females in any species (Supplementary Figure 1).

Importantly, two pieces of evidence suggested that all three species used a similar evidence accumulation strategy. First, across species, trials with longer RTs yielded increased accuracy (Figure 2F; Pearson correlations between accuracy and RT, see methods; Mouse: $p < 0.017, R = 0.80$, Rat:$p < 0.011, R = 0.83$, Human: $p < 0.033, R = 0.85$). Second, choice and

RT data from individuals of each species were well fit by the drift diffusion model (DDM; Bogacz et al., 2006; Ratcliff, 1978; Ratcliff & McKoon, 2008), which showed relatively high decision bounds (Figure 3A-C).

Species differ in key model parameters revealing differences in the accumulation process

While all three species were well fit by the DDM, our initial investigation showed cross-species differences in RT, suggesting possible differences in the accumulation process across species. To further explore cross-species differences, we examined how DDM parameters varied across species. The DDM assumes that the accumulation process is governed by four main parameters: **boundary separation** (a)**, drift rate** (v), **non-decision time** (ndt), and **starting point bias** (x_0); these were estimated for individual subjects (Figure 3A). Comparison of the DDM parameters showed that humans were characterized by the largest drift rate (v) and boundary separation (a), whereas mice had the lowest drift rate and boundary separation (Drift rate: $p < 0.0005, F = 77.72$; Boundary separation: $p < 0.0005, F = 88.29$; Figure 3D-E). Rats showed the greatest non-decision time (ndt), while mice had the shortest non-decision time ($p < 0.0005, F = 347.26$; Figure 3F). There was also a trend effect for greater starting point bias (x_0) in rodents ($p = 0.053, F = 2.99$; Figure 3G).

Thus, the DDM parameters suggested key differences in the accumulation process across species. The higher boundary separation in humans suggested a higher threshold for perceptual evidence before commitment to a decision, leading to more accurate but slower responses. The greater non-decision time and lesser boundary separation in rats suggested that rats accumulated less evidence, but spent more time in processes unrelated to decision making (for example, motor movements). Finally, mice had the least boundary separation indicating less evidence accumulation, also less time spent in non-decision processes. Accordingly, mice were the fastest and least accurate of the three species.

Rodents' choices were limited by time-pressures

Given the faster RT in mice, and model-based evidence for faster decision times in rats, we asked if rodents' choices were limited by internal time-pressures. The classic version of the DDM discussed above, does not readily account for this possibility, for it assumes that with time in a trial, the amount of evidence required to make a decision remains the same. This assumption is implemented in the model through a fixed decision boundary (a). To test if rodents experienced time-pressures, we considered two additional variants of the DDM (see methods and

Supplementary Figures 2, 3 and 4), a collapsing boundary model (Bowman et al., 2012; Shevinsky & Reinagel, 2019), and an urgency model (Ditterich, 2006).

The collapsing boundary DDM assumes that with time in a trial, the decision boundary collapses, and thus, less evidence is required to make a decision. It included three additional parameters: **initial boundary separation** (a_0), **semi-saturation constant** ($t_{0.5}$), and **amount of collapse** (k) of the boundary (see methods and Figure 4A). The urgency model incorporates an urgency signal that increases with time, raising the probability of a choice. The collapsing boundary- and urgency models are mathematically similar, but the collapsing boundary model can better explain faster RTs on more difficult trials. The urgency DDM also included three additional parameters: **slope** (u_{slope}), **magnitude** (u_{mag}), and **onset** (τ) of the urgency signal (Figure 4B).

Model comparison using the BIC revealed that the collapsing boundary DDM fit rodent data significantly better than the classic- or the urgency DDM (Repeated measures ANOVA; Mice: $p < 0.0005, F = 22.12$; Rats: $p = 0.0014, F = 11.12$; Figure 4C). However, for humans, the collapsing boundary- or the urgency model was not better than the classic DDM (One-way ANOVA; $p = 0.13, F = 2.31$; Figure 4C). These results suggest that rodents' behavior was limited by internal time-pressures. Further, the collapsing boundary DDM showed the lowest semi-saturation constant ($t_{0.5}$) for mice, which suggests that the boundary collapse was fastest for mice ($p = 0.004, F = 5.69$; Figure 4E), producing the fastest responses.

Overall, these results supported the hypothesis that internal time-pressures limited rodents' performance in the current task.

Rats balance speed and accuracy near optimally to maximize the reward rate.

Summary performance measures suggested cross-species differences in accuracy and RT, and the DDMs helped explain the differences in the underlying decision parameters. However, it remains unclear how each species settled on their specific balance of speed vs accuracy (Figure 5A). One possibility is that agents chose an RT that maximized their reward rate (Figure 5B-C and Supplementary Figure 5; see also Bogacz et al. 2006).

To investigate this possibility, we simulated a perfect accumulator for each species, accounting for cross-species differences in the inter-trial interval (Supplementary Figure 5G-I, also see methods). Interestingly rats were close to the simulated optimal RT, while humans responded

slower than the optimal RT and mice were faster (Figure 5D-F). These observations suggested that rats may balance speed and accuracy to maximize reward rate.

To explore whether rats flexibly shift their decision strategies to obtain higher reward rates, we designed a new version of the task in which reward volume increased progressively for longer RTs (see methods, and Figure 6A-B, F). We trained 6 rats in this version of the task and examined how speed and accuracy shifted as reward contingencies changed. As expected, as reward volume increases for longer RTs, rats slowed their responses. Simulations with an accumulator model indicated that the increased RT observed in rats was near the optimal for reward rate maximization (Figure 6B-C). Interestingly however, rats' accuracies also improved, suggesting that they continued to accumulate evidence during trials with extended RTs (Figure 6C-E and 6G-H). Together, these results suggest that rats can flexibly alter the speed-accuracy balance to maximize reward rate.

Mice were more influenced by alternate strategies across latent states

Previous research suggests existence of alternate decision strategies in evidence accumulation tasks (Ashwood et al., 2020; Carandini & Churchland, 2013; Roy et al., 2021). To investigate if subjects in our task displayed shifts in strategy, we used a generalized linear model (GLM) that estimated the binomial probability of choosing the right- over the left side, as a linear function of perceptual evidence and the following alternate strategies: **win-stay- lose-switch** (determined by the choice and reward), and **previous choice** (determined by choice independent of reward). Perceptual evidence was the flash difference, divided by total flashes ($\frac{\delta flash}{\Sigma flash}$). The intercept term represented side-bias (Figure 7A, also see methods). The GLM suggested that rats and humans relied more on perceptual evidence, but mice were influenced by alternate strategies ($p < 0.0005, F = 156.28$; Figure 7D-E). The GLM predicted the psychometric curve for rats and humans well, but for mice the fit was poorer ($p < 0.0005, F = 38.67$; Figure 7B-C).

The influence of multiple strategies on decisions may suggest that subjects dynamically switch behavioral strategies on a trial-by-trial basis (Ashwood et al., 2022). To test this hypothesis, we used a modeling framework previously proposed by Ashwood et al., (2022) that combined a hidden Markov model (HMM) with a GLM. Using this GLM-HMM, we investigated the number of latent states underlying choice behavior in each species. We considered the original GLM, a GLM with lapse states, and multiple GLM-HMMs with increasing latent states—2, 3, 4, and 5 states,

respectively. Model comparison showed that the original GLM fit the data from rats and humans better (Figure 7F). However, mice data were better fit by a 3-state GLM-HMM (Figure 7F); the 3 states were evidence accumulation and left-side and right-side biases (Supplementary Figure 6). Thus, mice appeared to alternate between different latent states that prioritized evidence accumulation or alternate decision strategies.

In line with the above results, we also found interesting, time-varying trends in raw behavior (Supplementary Figure 7). Rats were remarkably stable in maintaining a constant RT and accuracy across their 120-minute session. Rats performed similar numbers of trials across time and had a stable reward rate across the session. In contrast, in mice, RT, accuracy and reward rate changed over the course of a session. Thus, mouse behavior in this task was more variable trial-to-trial when compared to rats or humans.

Discussion

We present a behavioral framework, synchronized across rodents and humans, for investigations of perceptual decision making. Inspired by pulse-based evidence accumulation tasks, previously used in rats, we designed a free response task that offered a speed-accuracy trade-off. In parallel, we developed a video game based on the same task to test humans online. Together we used this task and game to assess species-specific priorities in decision making. Task parameters and training procedures were deliberately synchronized across species, facilitating direct, quantitative comparisons of behavior. All three species learned to perform the task, their behavior was well fit by evidence accumulation models, and we found significant species-specific differences in key parameters. Unlike humans, rodents were best fit by a collapsing boundary DDM suggesting that their performance was limited by internal time-pressures. Rat's response times were close to optimal to maximize reward rate, even as task conditions changed. Mice alternated between evidence accumulation- and other strategies, and exhibited more variable behavior. Overall, this study highlights the importance of a synchronized behavioral framework for meaningful, quantitative comparisons of cross-species perceptual decision making.

Our framework has several key attributes to facilitate quantitative comparisons of behavior across species. First, to minimize training differences, we used a similar non-verbal feedback-based training strategy in which difficulty was slowly increased to facilitate learning. Second, in humans we used a online video game-based approach. Video games can be advantageous over traditional psychological experiments; games are engaging and can be used to collect objective

behavioral measures with well-controlled stimuli. Also, online games are helpful for collecting larger datasets (see also, Do et al., 2023). Finally, the short, self-paced nature of the video game enabled data collection within a relatively brief period of time (mean completion time = 11.06 minutes +/- 2.26 S.D., while the task could be implemented in a high-throughput, semi-automated rodent training facility allowing collection of behavioral data from many animals. Together, this framework is advantageous for collecting larger datasets across different species within a relatively short amount of time.

An interesting result revealed by this approach is that rodent performance was best fit by a collapsing bound model, suggesting their decisions were limited by internal time-pressures. This was revealed by a comparison of the classic- (Ratcliff, 1978) and the collapsing boundary (Bowman et al., 2012; Shevinsky & Reinagel, 2019) version of the DDM. The classic DDM, a highly successful model for explaining perceptual decision making behavior (Ratcliff, 2014; Ratcliff et al., 1999; Ratcliff & Rouder, 1998), assumes that the amount of evidence required to make a decision remains constant. Others have suggested situations in which less evidence is required with time (Busemeyer & Rapoport, 1988; Drugowitsch et al., 2012; Thura et al., 2012, Ditterich, 2006). However, experimental evidence for the collapsing boundary model has been inconclusive (Hawkins et al., 2015; Voskuilen et al., 2016). Notably, while previous studies used these models to explain human and non-human primate data, here we show that rodents decisions are best fit by a collapsing bound model.

While humans and rats were best fit by a single perceptual accumulation strategy across all trials, we found that mice were best fit by a model that alternated between evidence-accumulation and other strategies across different latent states. This finding is consistent with previous studies that have examined mouse behavior in other perceptual decision tasks (Ashwood et al., 2022; Carandini & Churchland, 2013; Gupta et al., 2024; Hanks & Summerfield, 2017).

The framework we provide here provides a complement to cross species studies using ethological approaches. Recent research has highlighted the importance of focusing on ethological behaviors in different species (Juavinett et al., 2018). Indeed, focus on ethological behaviors can provide unique insights into cross-species comparisons. However, this rarely allows for direct, quantitative cross-species comparisons as ethologically-relevant behaviors are often species-specific, reflecting adaptations to that organism's particular niche. Here, we followed an approach that is

complementary to the ethological focus, by designing a task that can be performed in similar ways by multiple species, in order to better appreciate the unique and common aspects of behavior.

Limitations

Species specific differences we report here could reflect uncontrolled environmental variables rather than genetic differences. Despite a deliberate attempt to synchronize most aspects of our task and training across species, there were several differences in the human and rodent setups that could have contributed to the cross-species behavioral differences. First, rewards were different for rodents and humans; drops of sugar water for rodents vs point bonuses for humans. Moreover, while mice were water scheduled to increase motivation and rats were food scheduled. Second, rodents and humans performed the task for different durations, which could have impacted their motivational states differentially. Third, humans participated in the task from their home environment, with the experimenter only present remotely via Zoom. In contrast, the rodent underwent requisite animal handling procedures which were conducted by trained researchers.

Conclusion

In sum, we present a synchronized behavioral framework for high-throughput data, and direct, quantitative comparisons of perceptual decision-making across species. Beyond basic neuroscience investigations, this framework could be used in translational research on neuropsychological conditions impacting perceptual abilities (van den Boogert et al., 2022), such as autism spectrum disorder (Robertson & Baron-Cohen, 2017), Schizophrenia (Horga & Abi-Dargham, 2019), and Alzheimer's disease (Kavcic et al., 2011; Rizzo & Nawrot, 1998). Indeed, our video game-based approach could be used to assess quantitative behavior traits in humans with neuropsychiatric disorder and these quantitative traits could then be used as an objective behavioral benchmark, or biomarker, to aid in better selection and assessment of genetic animal models. Moreover, as this task is amenable to neurophysiological investigations, circuit mechanisms of behavioral traits can be followed up in the preclinical animal model. Finally, we point out that the non-verbal, feedback-driven training pipeline for humans may allow inclusion of participants with communication difficulties, such as non-verbal autistic individuals, improving the reach of this research to underrepresented groups (McKinney et al., 2021).

Acknowledgements

We thank Christa Rose, Arielle Rubel, Julie Gomez, Sanaa Ahmed, Anosha Khawaja-Lopez, Benjamin Lee and Sina Analoui for help with collection of rodent data, and Vanessa Torres-

Lacarra and Phoebe Lesk for collection of human data. We thank Helen Tager-Flusberg for feedback on human data collection, Joseph McGuire for feedback on the analyses, and Josh Sanders for feedback on the behavioral training software. This work was supported by NIMH award number R56MH132732 and SFARI award number 874568 to BBS and a Center for Systems Neuroscience distinguished fellowship award to CDS and GAK.

Contributions

SC, CDS, GAK, HX and QHD designed the experiments. SC and CDS analyzed the data. SC, CDS and BBS wrote the manuscript. All authors contributed to the conceptualization and execution of the study and edited the manuscript.

Materials and methods

Rodents

All experiments and procedures were performed in accordance with protocols approved by the Boston University Animal Care and Use Committee. Long Evans adult rats (N=21; aged 3 months to 2 years) were purchased from Taconic or bred in house. C57BL/6NJ adult mice (N=95; aged 2 to 8 months) were purchased from Jackson Laboratory. Both male and female rats or mice were utilized and trained simultaneously in the same room, albeit in separate operant chambers. Rats were food restricted to 80-100% of their body weight and fed once per day (typically 3-4 pellets of food per day). Mice were water deprived to 80-90% of their body weight and got a minimum of 1mL of water a day. Rats received 0.03 mL reward, whereas mice received 0.005 mL reward based on 10% sucrose (100g/L). Reward volume was consistent within a session. Both rats and mice were housed on a 14:10 ON:OFF light schedule with the ON phase corresponding to daylight hours in Boston, MA USA.

Rodent behavioral control system

Behavioral control system was inspired by Dhawale et al., 2017 and Poddar et al., 2013. Rodents were trained in custom acrylic chambers with three nose ports. Nose ports were 3D printed (Sanworks or custom made) and equipped with a visible LED for stimulus delivery (Sanworks), peristaltic pump for reward delivery, and an IR LED and photodetector as a beam break (Sanworks). Behavioral control software to implement the task and control individual boxes was written in MATLAB. Boxes were controlled through a Teensy-based microcontroller system (Bpod Sanworks). A custom written python application was used to control multiple Bpod instances from a single control computer, requiring an edited version of the Bpod MATLAB software library

(edited Bpod library: https://github.com/RatAcad/Bpod_Gen2; Custom python application: https://github.com/RatAcad/BpodAcademy). Further details about the software implementation can be found in the respective code repositories. The floor of the chambers contained bedding.

Rodent daily training

Mice were housed in groups of 2-4 mice per cage whereas rats were housed in pairs in an animal facility. Both were moved to the training room in the laboratory each morning for training. Mice ran for 1 hour shifts and rats ran for 2 hour shifts, 5 days per week in the late morning (10am-noon) or early afternoon (12pm-3pm). Rat feeding was conducted in the late afternoon after behavioral training (1-4pm).

Rodent training pipeline

Rats and mice were rewarded with 10% sucrose at all stages of training and testing. Training took 1-4 weeks depending on the rodent and progressed through 3 stages. In the first stage (1-3 days) rodents were rewarded for inserting their nose into a side port with an illuminated LED (400 trials). In the second stage rodents received reward for inserting their nose into the center port and then the side port with an illuminated LED (400 trials). In the third stage, they received a reward for inserting their nose into the center port and then the side port with a flashing LED. Once rodents reached criterion 90 -100% correct over 400 trials, the probability of flashes on the incorrect side increased in the following way 100:0 -> 90:10 -> 80:20. Progression through this third stage took 1-5 days per condition. After completion of all stages and conditions rodents performed the task with a flash probability of 80:20.

Rodent behavioral task

Rodents were free to move within their cage during the task. At the start of a trial, the light in the center port turned on indicating that the rodent could initiate the trial at its leisure. Once the rodent nose poked in the center port, a single flash would occur simultaneously in both the left and right ports. After this one light flash on both sides, the rodent would see a series of light flashes according to a Bernoulli process – every 100 ms, the rodent would see a flash on either the right or the left side with a probability of 80:20 that the flash would occur on the correct side vs. the incorrect side. The correct side was drawn randomly on each trial. To record a response, the rodent nose poked in either the left or right port and the series of light flashes would terminate as soon as this decision was recorded. If the rodent responded correctly, the light in the correct side port turned on for 3 s and a 30uL-rat or 5uL-mice sucrose water reward was delivered

immediately. If the rodent responded incorrectly, all lights turned off for a 5 s time-out and no reward was delivered, and new trials would start after a 2s delay. Trials had an external response deadline of 8 seconds, after which it would be counted as an omission.

Rats reward inflection time contingency task

Rats underwent progressive training on a new version of the free-response evidence accumulation task, where longer response times yielded greater rewards. Training involved seven different reward inflection paradigms, each based on a sigmoidal curve (Figure 6F). The sigmoidal curve for each reward inflection time was determined using the following equation:

$$R(t) = R_{min} + \frac{R_{max} - R_{min}}{1 + e^{\frac{-(f(t)-b)}{2}}}$$

In this equation, R(t) represents the reward obtained at time t, while Rmin and Rmax denote the minimum and maximum rewards received by the rats, respectively. The function f(t) represents the vector of the number of flashes presented to the rat over time, with b indicating the flash bin where the inflection point occurs. For each paradigm, the inflection point marks the flash bin where the rat receives the baseline reward from the task (30 µl). The slope of the sigmoidal curve remains constant at 2, while b progressively shifts the curve by 5 flash bins (0.5 seconds) for each reward inflection paradigm. This adjustment allows for the delay of the optimal response for each reward inflection paradigm. Rats completed a minimum of 1500 trials for each paradigm. Rats were trained in the reward inflection paradigm at 75:25 task difficulty.

Human participants

Human participants $(N = 18)$ were adolescents (age: $13 - 18\ years$), who were invited to play the game online. All participants were typically-developing siblings to autistic individuals, and were recruited from one of the following databases: the SPARK database sponsored by the Simons Foundation, and the Center for Autism Research Excellence (CARE) registry maintained at Boston University. Informed consent and assent were obtained from participants and parents; parents also provided demographic information, and completed the following autism-relevant surveys on their behalf: Vineland Adaptive Behavior Scales (VABS) 3, Social Responsiveness Scale (SRS) 2, Adolescent Adult Sensory Profile (AASP), and Behavioral Inflexibility Scale (BIS). All participants were required to have normal or corrected-to-normal vision. Participants with a seizure disorder were not recruited. All procedures were approved by an institutional review board at Boston University. Participants first played another game, also based on the pulse-based evidence task, but with a fixed response design—the response could only be made after the cue

period. Results from the fixed-response game, and that for the autistic participants are part of another investigation. Participants played the current, free-response game typically months after the fixed-response game. For each successful session, participants received an amazon gift card worth USD 25.

Human video-game and experimental procedure

The video game was developed using a JavaScript application called PixiJS, and run online using the cloud platform Heroku; data was saved online using the no-SQL database application MongoDB, in JSON format. Demographic and survey data were collected separately using Boston University's REDCap platform. Each online game session was supervised by a member of the lab, trained in human behavioral testing procedures. This experimenter was present on Zoom and provided guidance on the game setup, game rule was not disclosed; the Zoom sessions were recorded. The fixed-response and the free-response games had different visual narratives. In the fixed response game, participants collected gems. In the free-response game, they destroyed asteroids. Both games included a total of 200 trials, a successful session meant that the participant completed all trials. In the free-response game, the first 10 trials served as learning trials, equivalent to Stage 4 in rodents (90:10 generative flash probability), where participants needed to achieve above 90% accuracy to be included in the study. The remaining trials were task trials, equivalent to rodents Stage 5 (80:20 generative flash probability). In each trial, participants selected the asteroid they wanted to destroy. Subsequently, a series of alien-shaped stimuli were presented according to a Poisson process, with an 80% probability of appearing on the correct side and a 20% probability of appearing on the incorrect side, or not at all, every 100 ms. The correct side was randomly determined for each trial. Participants responded by selecting the alien on the correct side, concluding the stimuli presentation. A correct response resulted in the destruction of the asteroid, awarded points, and a wait time of approximately 1 second before the next trial. An incorrect response led to the asteroid remaining intact, no points awarded, and a wait time of approximately 3 seconds before the next trial. If participants failed to respond within 8 seconds, the trial was counted as an omission.

Synchronization of the free response pulse-based evidence accumulation task across species

We developed a free-response version of the Poisson clicks task, wherein two competing streams of visual pulses were presented on the left and right sides of the subject's visual field. Subjects were rewarded to select the side with more pulses to receive a reward. Pulses were generated

using a Poisson process for humans and a Bernoulli process for rodents, with the pulse probability manipulated across task stages. The task maintained consistent stimulus statistics, including flash duration, flash rate, and generative probability, for both rodents and humans—specifically, a 10 Hz flash rate and an 80% flash probability on the rewarded side versus a 20% flash probability on the other side in the final stage (Stage 5).

The task structure remained uniform across species, comprising four stages: initiation, free-response cue period, choice, and reward/time-out (Figure 1A). Human participants engaged in a single session (consisting of 200 trials) of the video game online under the supervision of a researcher via video call (Figure 1B). Multiple rodents performed the task in a 3-port chamber on a daily basis, with data transferred on the same day using Datajoint (Figure 1C).

Rodents underwent training through a series of stages until they reached Stage 5 (with an 80-20% probability on the correct side), whereas humans completed 10-15 training trials at Stage 4 (90-10%) before advancing to Stage 5 (Figures 1D-F).

Data analyses

Data analysis was conducted using Python version 3.7.1 and R version 4.3.1. Multiple sessions were introduced per rodent, resulting in an average of 3605 ± 231 total trials per mouse and 4107 ± 549 total trials per rat. Chance performance of the task was determined by calculating the upper bound of the 99% binomial confidence interval for each species performing at chance, yielding an estimated chance performance of 52% for rodents and 60% for humans.

Accuracy was calculated by dividing the number of correct trials by the total number of trials and then multiplying the result by 100. Response time is defined as the average duration between the initiation of a stimulus and the subsequent response. Percentage bias was calculated by extracting the absolute difference between right and left trials and dividing by the total number of trials, the result was multiplied by 100. The average accuracy, response time, and reward rate during a session and across sessions were computed using a rolling average of 5 minutes during the 60 (mouse) or 120 (rat) minutes of the session. Trials across the session were calculated by counting the trials per minute and averaging them across sessions and subjects (Supplementary Figure 7). Standard statistical tests (e.g., One-way ANOVAs, Pearson and Spearman correlations, multiple comparisons and t-students) were conducted using the pingouin package (Vallat, 2018).

To model the accuracy versus response time relationship shown in Figure 2F, a generalized additive mixed model (GAMM) was utilized, employing the mgcv package (Wood et al., 2017). The data points on the plot represent the average accuracy of all trials at each response time, calculated using 0.2-second bins over a span of 3 seconds; only bins that presented more than 50 trials were selected. Pearson correlations were performed using the bins within the first 1.5 seconds for each species. All trials for each species were collectively fitted within the model framework and the data.

Drift diffusion models

Drift diffusion models (DDMs) are widely utilized to characterize decision-making in two-choice paradigms. The DDM assumes that decision-makers accumulate noisy evidence for each choice option over time, ultimately committing to a response when a predefined decision threshold is reached. In our study, we employed three variants of DDMs: the classic DDM (Bogacz et al., 2006; Ratcliff, 1978; Ratcliff & McKoon, 2008), the collapsing boundary DDM (Bowman et al., 2012; Hawkins et al., 2015; Shevinsky & Reinagel, 2019; Voskuilen et al., 2016) and the urgency DDM (Ditterich, 2006; Hawkins et al., 2015). The classic DDM assumes that the amount of evidence required to reach a decision is fixed over time, and implements it with fixed decision boundaries (a). By convention, we assume the lower boundary to be 0, which sets a total boundary separation of $a - 0 = a$. In addition, it includes the following parameters:

1) The drift rate (v), which is the rate at which information is sampled; evidence (x) grows linearly with time (t): the change in evidence (dx), is updated by the drift rate (v), and a gaussian noise (W) scaled with a constant (c):

   $$dx = v_i * dt + cdW$$

   Here, i represents a trial.

2) The non-decision time (ndt), which is the time taken up by processes not-related to evidence accumulation.
3) The starting point bias (z), which determines the location on the distance between the two boundaries from which the evidence accumulation process starts, it reflects a bias to choose one option over the other one. z is relative to a.

Thus, evidence is accumulated over time using the drift rate embedded with noise. Evidence is accumulated until the decision threshold is reached. We fit this model to the choice-RT data of individuals from each species. For each model we also assumed that drift rate, non-decision time

and starting point bias can vary from trial-to-trial following normal, half-normal and uniform distributions, respectively:

Drift rate: $v_i \sim N(v, theta_v)$;

Starting point: $x0_i \sim Unif(x0 - theta_x0, x0 + theta_x0)$;

Non-decision time: $ndt_i \sim Unif(ndt - theta_ndt, ndt + theta_ndt)$.

Unlike the classic DDM, both collapsing boundary model and the urgency model assume that the amount of evidence required to reach a decision is not fixed over time. The collapsing boundary model included three additional parameters— 1) initial boundary separation (a_0), 2) the semi-saturation constant ($t_{0.5}$), or the time when the boundary had collapsed by half, and 3) the total amount of collapse (k). The urgency model also included three additional parameters— 1) the slope of a (linear) urgency signal (u_{slope}), 2) total amount of urgency (u_{mag}), and 3) the onset delay of urgency (τ).

The basic idea behind fitting a DDM to experimental data is to find a set of parameters which can simulate trials with the same accuracy and RT distributions as the experimental data. This is done iteratively, by perturbing the parameters, starting from an initial set of values, and comparing the fit between the simulated and experimental trials. Here, we used the χ^2 method to compare between the RT histograms of the simulated- and experimental data, because the method is relatively fast and less prone to outlier RTs (Ratcliff & Tuerlinckx, 2002). All models were fitted using a custom R package (rddm; https://github.com/gkane26/rddm).

Simulation analyses

For all simulations with the perfect accumulator, we included 10,000 trials for each choice of RT; RT typically varied between 100 ms and 4 s, with an increment of 200 ms. For each choice of RT, we calculated the accuracy and/or reward rate of the accumulator, considering the 10,000 trials. Consistent with the current task, each trial had a generative flash probability (p) of 80%. For simulations of the reward-RT contingency training experiment conducted in rats, this was 75%. For each trial, the correct side (left or right) was pre-determined with a coin flip. Individual flashes were generated with a Bernoulli process for the simulations for rodents and a Poisson process for the simulations for humans. Sequence of flashes continued only up to the choice of RT. The perfect accumulator always chose the side with greater total number of flashes. If the trial ended with the same flash counts on each side, the perfect accumulator randomly picked a side. We also considered a noisy accumulator, which was informed by the accuracy of each species at

specific RTs; this was obtained from the accuracy predictions of the GAMM (Supplementary Figure 6).

For simulations of the reward-rate, in addition to RT, we considered species-specific choices of inter-trial intervals (ITI). As mentioned in the results, ITI was the interval between the response and the next stimulus onset; it included the time to initiate a trial and the time spent in post-response events, such as rewards, timeouts etc. The average ITI was different across species, there was also a difference in the average ITI for rewarded and unrewarded trials. These differences were accounted for in the simulations. In the current experiment, each correct response produced the same amount of rewards, assuming each reward to be of $1\ unit$, a total accuracy of N would produce a total of N rewards. For the reward-RT contingency training experiment conducted in rats, rewards were quantified as the proportion of this baseline reward (30 µl or 1 unit). Accordingly, depending on the RT condition, a total accuracy of N would produce a total of $s * N$ rewards, where s depends on the relative reward proportion. For each choice of RT, reward-rate was calculated for the 10,000 trials as:

$$Reward\ rate\ =\ Total\ rewards \div (RT + ITI)$$

The optimal RT was the RT corresponding to the peak reward-rate in the RT versus reward-rate plot. Note that these simulations assumed that RT included the decision- or evidence-accumulation process only, whereas other models have also incorporated periods of non-accumulation (Simen et al., 2009).

GLM and GLM-HMM

The generalized linear model (GLM) was utilized to elucidate choice probability, specifically the tendency to select the right option over the left option. This model incorporated three main predictors as a linear function:

$$\log\left({}^{\pi}/_{1-\pi}\right) = \beta_{bias} + \beta_{{}^{R-L}/_{R+L}} + \beta_{WSLS} + \beta_{PC}$$

1. Pulse or flash difference: Calculated as the disparity between flashes on the two sides, normalized by the total number of flashes on both sides.
2. Win-stay or lose-switch strategy: Reflecting the tendency to stick with a chosen side after a win or switch sides after a loss.

3. Previous choice strategy: Indicating the inclination to follow the previous choice made.

The intercept term in the model represented the inherent bias towards one side over the other. Notably, this analysis was restricted to data from the 80% versus 20% flash stage.

To compare model fit across species with varying trial numbers, we employed the log-likelihood per trial in bits, as outlined by Ashwood et al. (2022). This is calculated as follows:

$$LL_{Bits\ per\ trial} = \frac{LL_{data}}{n_{test} log(2)}$$

where LL_{data} is the log-likelihood of the entire data per species. n_{test} is the number of trials in the data. Predictive accuracy, following the methodology of Ashwood et al. (2022), was also utilized to gauge the goodness of fit.

Afterward, we compared these fits with a 2-state classic lapse model, which posits that subjects transition between two latent states. In one state, termed the engaged state, decisions are influenced by the stimuli, leading to dependent probabilities associated with the left and right stimuli. In the other state, referred to as the disengaged or lapse stage, decisions are entirely independent of the stimuli. In this scenario, subjects essentially flip a biased coin on each trial, with a fixed probability (γr+γl), and subsequently adopt one of two strategies based on the outcome: a strategy contingent on the stimulus or a strategy that disregards the stimulus entirely.

$$p(y_t = 1|x_t) = \left\{ \begin{array}{ll} \frac{1}{1+e^{-w*x}}, & z_t = "engaged" \\ \frac{y_r}{y_r+y_l}, & z_t = "lapse", \end{array} \right\}$$

Furthermore, to elucidate potential transitions through different task stages, we applied a modeling framework incorporating hidden Markov models (HMMs). This framework, known as GLM-HMM, assumes distinct latent states that animals may inhabit during task performance, each potentially influencing decision strategies differently. This model comprises multiple independent GLMs, each corresponding to a latent state and characterizing choice outcomes as a Bernoulli process. Transition between states is governed by transition probabilities, and weight vectors within each state capture the relative influence of decision strategies. We fitted the GLM-HMM to cross-species data, varying the number of latent states and employing hyperparameters σ = 2

and α = 2 to govern the prior. A detailed description of the model is presented in Ashwood et al., (2022),

References

Ashwood, Z. C., Roy, N. A., Bak, J. H., & Pillow, J. W. (2020). Inferring learning rules from animal decision-making. *Advances in Neural Information Processing Systems*, *33*, 3442–3453.

Ashwood, Z. C., Roy, N. A., Stone, I. R., Urai, A. E., Churchland, A. K., Pouget, A., & Pillow, J. W. (2022). Mice alternate between discrete strategies during perceptual decision-making. *Nature Neuroscience*, *25*(2), Article 2. https://doi.org/10.1038/s41593-021-01007-z

Badre, D., Frank, M. J., & Moore, C. I. (2015). Interactionist Neuroscience. *Neuron*, *88*(5), 855–860. https://doi.org/10.1016/j.neuron.2015.10.021

Barron, H. C., Reeve, H. M., Koolschijn, R. S., Perestenko, P. V., Shpektor, A., Nili, H., Rothaermel, R., Campo-Urriza, N., O'Reilly, J. X., Bannerman, D. M., Behrens, T. E. J., & Dupret, D. (2020). Neuronal Computation Underlying Inferential Reasoning in Humans and Mice. *Cell*, *183*(1), 228-243.e21. https://doi.org/10.1016/j.cell.2020.08.035

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. https://doi.org/10.1037/0033-295X.113.4.700

Bowman, N. E., Kording, K. P., & Gottfried, J. A. (2012). Temporal Integration of Olfactory Perceptual Evidence in Human Orbitofrontal Cortex. *Neuron*, *75*(5), 916. https://doi.org/10.1016/j.neuron.2012.06.035

Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science (New York, N.Y.)*, *340*(6128), 95–98. https://doi.org/10.1126/science.1233912

Busemeyer, J. R., & Rapoport, A. (1988). Psychological models of deferred decision making. *Journal of Mathematical Psychology*, *32*(2), 91–134. https://doi.org/10.1016/0022-2496(88)90042-9

Carandini, M., & Churchland, A. K. (2013). Probing perceptual decisions in rodents. *Nature Neuroscience*, *16*(7), Article 7. https://doi.org/10.1038/nn.3410

Dhawale, A. K., Poddar, R., Wolff, S. B., Normand, V. A., Kopelowitz, E., & Ölveczky, B. P. (2017). Automated long-term recording and analysis of neural activity in behaving animals. *eLife*, *6*, e27702. https://doi.org/10.7554/eLife.27702

Ditterich, J. (2006). Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks: The Official Journal of the International Neural Network Society*, *19*(8), 981–1012. https://doi.org/10.1016/j.neunet.2006.05.042

Do, Q., Li, Y., Kane, G. A., McGuire, J. T., & Scott, B. B. (2023). Assessing evidence accumulation and rule learning in humans with an online game. *Journal of Neurophysiology*, *129*(1), 131–143. https://doi.org/10.1152/jn.00124.2022

Drugowitsch, J., Moreno-Bote, R., 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

Esteves, M., Moreira, P. S., Sousa, N., & Leite-Almeida, H. (2021). Assessing Impulsivity in Humans and Rodents: Taking the Translational Road. *Frontiers in Behavioral Neuroscience*, *15*. https://doi.org/10.3389/fnbeh.2021.647922

Gupta, D., DePasquale, B., Kopec, C. D., & Brody, C. D. (2024). Trial-history biases in evidence accumulation can give rise to apparent lapses in decision-making. *Nature Communications*, *15*(1), 662. https://doi.org/10.1038/s41467-024-44880-5

Hanks, T. D., & Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron*, *93*(1), 15–31. https://doi.org/10.1016/j.neuron.2016.12.003

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. https://doi.org/10.1523/JNEUROSCI.2410-14.2015

Horga, G., & Abi-Dargham, A. (2019). An integrative framework for perceptual disturbances in psychosis. *Nature Reviews Neuroscience*, *20*(12), 763–778. https://doi.org/10.1038/s41583-019-0234-1

Juavinett, A. L., Erlich, J. C., & Churchland, A. K. (2018). Decision-making behaviors: Weighing ethology, complexity, and sensorimotor compatibility. *Current Opinion in Neurobiology*, *49*, 42–50. https://doi.org/10.1016/j.conb.2017.11.001

Kavcic, V., Vaughn, W., & Duffy, C. J. (2011). Distinct Visual Motion Processing Impairments In Aging and Alzheimer's Disease. *Vision Research*, *51*(3), 386–395. https://doi.org/10.1016/j.visres.2010.12.004

Markou, A., Chiamulera, C., Geyer, M. A., Tricklebank, M., & Steckler, T. (2009). Removing Obstacles in Neuroscience Drug Discovery: The Future Path for Animal Models. *Neuropsychopharmacology*, *34*(1), 74–89. https://doi.org/10.1038/npp.2008.173

McKinney, A., Weisblatt, E. J., Hotson, K. L., Bilal Ahmed, Z., Dias, C., BenShalom, D., Foster, J., Murphy, S., Villar, S. S., & Belmonte, M. K. (2021). Overcoming hurdles to intervention studies with autistic children with profound communication difficulties and their families. *Autism*, *25*(6), 1627–1639. https://doi.org/10.1177/1362361321998916

Nestler, E. J., & Hyman, S. E. (2010). Animal models of neuropsychiatric disorders. *Nature Neuroscience*, *13*(10), 1161–1169. https://doi.org/10.1038/nn.2647

Nguyen, Q. N., & Reinagel, P. (2022). Different Forms of Variability Could Explain a Difference Between Human and Rat Decision Making. *Frontiers in Neuroscience*, *16*. https://doi.org/10.3389/fnins.2022.794681

Odoemene, O., Pisupati, S., Nguyen, H., & Churchland, A. K. (2018). Visual Evidence Accumulation Guides Decision-Making in Unrestrained Mice. *The Journal of*

Neuroscience, *38*(47), 10143–10155. https://doi.org/10.1523/JNEUROSCI.3478-17.2018

Pedrosa, V., Menichini, E., Pajot-Moric, Q., Vincent, P., Zhou, L., Teachen, L., Latham, P., & Akrami, A. (2023). *Humans, rats and mice show species-specific adaptations to sensory statistics in categorisation behaviour* (p. 2023.01.30.526119). bioRxiv. https://doi.org/10.1101/2023.01.30.526119

Pinto, L., Koay, S. A., Engelhard, B., Yoon, A. M., Deverett, B., Thiberge, S. Y., Witten, I. B., Tank, D. W., & Brody, C. D. (2018). An Accumulation-of-Evidence Task Using Visual Pulses for Mice Navigating in Virtual Reality. *Frontiers in Behavioral Neuroscience*, *12*. https://www.frontiersin.org/articles/10.3389/fnbeh.2018.00036

Poddar, R., Kawai, R., & Ölveczky, B. P. (2013). A Fully Automated High-Throughput Training System for Rodents. *PLOS ONE*, *8*(12), e83171. https://doi.org/10.1371/journal.pone.0083171

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. https://doi.org/10.1037/0033-295X.85.2.59

Ratcliff, R. (2014). Measuring psychometric functions with the diffusion model. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(2), 870–888. https://doi.org/10.1037/a0034954

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, *9*(5), 347–356. https://doi.org/10.1111/1467-9280.00067

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. https://doi.org/10.3758/bf03196302

Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction

time. *Psychological Review*, *106*(2), 261–300. https://doi.org/10.1037/0033-295X.106.2.261

Redish, A. D., Sweis, B. M., Abram, S., Duin, A., Kazinka, R., Kocharian, A., MacDonald, A., Schmidt, B., Schmitzer-Tobert, N., & Thomas, M. (2021). *Sunk cost sensitivity in mice, rats, and humans on the Restaurant Row and WebSurf tasks cannot be explained by attrition biases alone* (p. 2021.10.07.462802). bioRxiv. https://doi.org/10.1101/2021.10.07.462802

Rizzo, M., & Nawrot, M. (1998). Perception of movement and shape in Alzheimer's disease. *Brain: A Journal of Neurology*, *121*(12), 2259–2270. https://doi.org/10.1093/brain/121.12.2259

Robertson, C. E., & Baron-Cohen, S. (2017). Sensory perception in autism. *Nature Reviews Neuroscience*, *18*(11), 671–684. https://doi.org/10.1038/nrn.2017.112

Roy, N. A., Bak, J. H., Akrami, A., Brody, C. D., & Pillow, J. W. (2021). Extracting the dynamics of behavior in sensory decision-making experiments. *Neuron*, *109*(4), 597-610.e6. https://doi.org/10.1016/j.neuron.2020.12.004

Schmack, K., Bosc, M., Ott, T., Sturgill, J. F., & Kepecs, A. (2021). Striatal dopamine mediates hallucination-like perception in mice. *Science (New York, N.Y.)*, *372*(6537), eabf4740. https://doi.org/10.1126/science.abf4740

Scott, B. B., Constantinople, C. M., Erlich, J. C., Tank, D. W., & Brody, C. D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *eLife*, *4*, e11308. https://doi.org/10.7554/eLife.11308

Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, *80*(3), 791–806. https://doi.org/10.1016/j.neuron.2013.10.047

Shevinsky, C. A., & Reinagel, P. (2019). The Interaction Between Elapsed Time and Decision Accuracy Differs Between Humans and Rats. *Frontiers in Neuroscience*, *13*, 1211. https://doi.org/10.3389/fnins.2019.01211

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(6), 1865–1897. https://doi.org/10.1037/a0016926

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

Vallat, R. (2018). Pingouin: Statistics in Python. *Journal of Open Source Software*, *3*(31), 1026. https://doi.org/10.21105/joss.01026

van den Boogert, F., Klein, K., Spaan, P., Sizoo, B., Bouman, Y. H. A., Hoogendijk, W. J. G., & Roza, S. J. (2022). Sensory processing difficulties in psychiatric disorders: A meta-analysis. *Journal of Psychiatric Research*, *151*, 173–180. https://doi.org/10.1016/j.jpsychires.2022.04.020

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. https://doi.org/10.1016/j.jmp.2016.04.008

Wood, S. N., Li, Z., Shaddick, G., & Augustin, N. H. (2017). Generalized Additive Models for Gigadata: Modeling the U.K. Black Smoke Network Daily Data. *Journal of the American Statistical Association*, *112*(519), 1199–1210. https://doi.org/10.1080/01621459.2016.1195744

Young, J. W. (2023). Development of cross-species translational paradigms for psychiatric research in the Research Domain Criteria era. *Neuroscience and Biobehavioral Reviews*, *148*, 105119. https://doi.org/10.1016/j.neubiorev.2023.105119

Figures

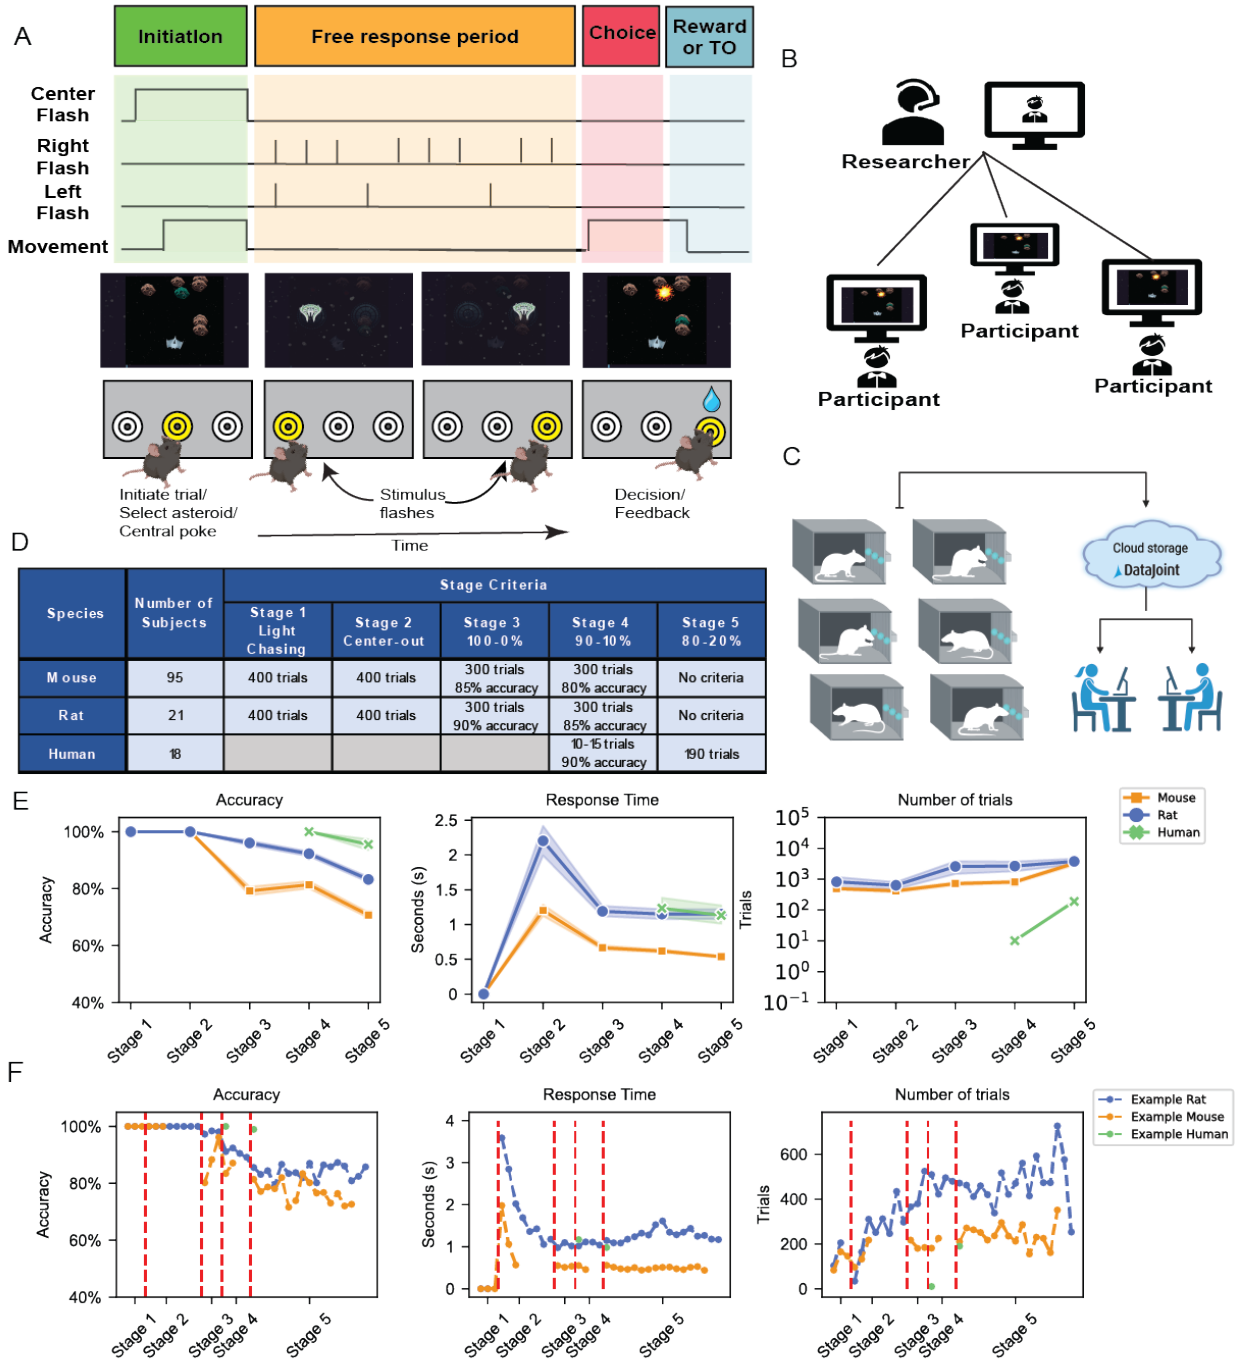

Species	Number of Subjects	Stage Criteria				
		Stage 1 Light Chasing	Stage 2 Center-out	Stage 3 100-0%	Stage 4 90-10%	Stage 5 80-20%
Mouse	95	400 trials	400 trials	300 trials 85% accuracy	300 trials 80% accuracy	No criteria
Rat	21	400 trials	400 trials	300 trials 90% accuracy	300 trials 85% accuracy	No criteria
Human	18				10-15 trials 90% accuracy	190 trials

Figure 1: Free response pulse-based evidence accumulation task for humans and rodents. A) Schematic representation of the free response evidence accumulation task for humans and rodents. The goal of the task is to identify the side with the higher probability of visual stimuli. The task structure is the same across species and consists of 4 different phases: 1) Initiation: the subject starts the task by clicking an asteroid (humans) or poking in the center (rodents), 2) Free response period or Cue period: the visual stimuli will be presented in the right or left side following

a Poisson (humans) or Bernoulli (rodents) process. 3) Choice: during the cue period, the subject decides when to respond by selecting a side, the stimuli are shown until the subject selects a side. 4) Reward or time-out (TO): the subject will receive a reward (increase of the score for humans and sugar water for the rodents) or a time-out depending if they selected the side with more stimuli probability or not. The human version consists of an online game format with the goal of destroying asteroids. In the rodent version, rodents play the task in a 3-port operant training chamber, in which the center port is used for initiation and the left and right ports are used for making choices. B) In humans, zoom video calls are used to monitor and collect the data online. C) Rodents behavioral sessions were conducted in multiple rodents at the same time and the data was collected via Datajoint. D) Table showing the number of subjects and training criteria across stages and species. Rodents learned the task in six different stages: Stage 1: poke in the side with the light; Stage 2: poke in the center to initiate and then in the side with the light. Stage 3-5: Task with different difficulties (100-0%, 90-10%, 80-20% flash probability on the correct vs incorrect side). Humans only participate in one session of 200 trials that was divided in two stages. The first 10-15 trials are learning trials (90-10% probability) in which they have to perform above 90% accuracy. The rest of the trials are done at 80-20% probability. E) Evolution of averaged accuracy, response time and number of trials per stage in mice, rats and humans. Mice averaged a lower accuracy and faster response times than rats and humans. F) Example mouse, rat and human evolution across stages and sessions. Note that the red dashed lines represent the transitions across stages.

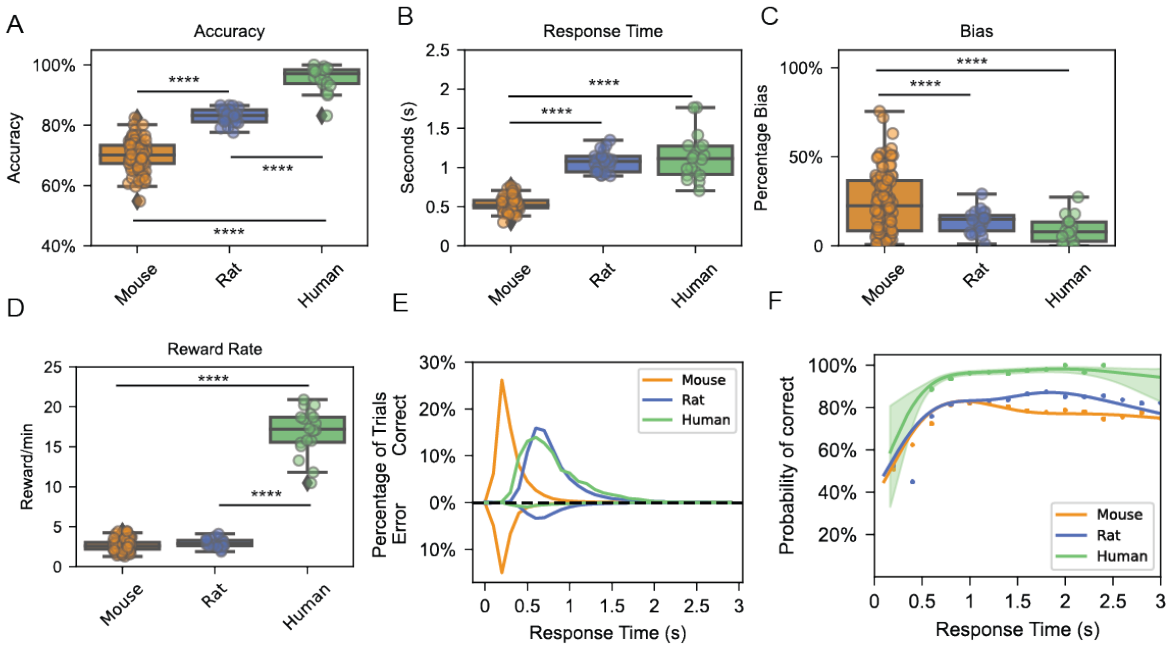

Figure 2: Descriptive statistics during the 80-20% probability phase reveal that all three species accumulate evidence over time to perform the task, albeit with distinct speed-accuracy trade-offs. A) Accuracy varies significantly across species (One-way ANOVA; $p < 0.0005$, $F = 271.84$), with mice exhibiting the lowest performance and humans demonstrating the highest. B) Mice exhibit faster response times compared to humans and rats, which display similar response times (One-way ANOVA; $p < 0.0005$, $F = 236.39$). C) Normalized bias is comparable between humans and rats but is higher in mice (One-way ANOVA; $p < 0.0005$, $F = 11.13$). D) Reward rate was similar for rodents but larger for humans (One-way ANOVA; $p < 0.0005$, $F = 1060.97$). E) Histogram of response times across species reveals that rats and humans respond at similar speeds during correct or incorrect trials, whereas mice respond faster in error trials than in correct trials (Paired t-student; $p < 0.0005$). Note that positive values represent correct trials, while negative values represent incorrect trials. F) Accuracy vs response time derived from a generalized additive mixed model on the three species in which we can observe the evidence accumulation process in each of the species. All of the species peak between 0.7-1 seconds but the humans present a higher peak than the rodents. Each dot represents the averaged accuracy per each 0.2-second bin for each species (Pearson Correlations Accuracy-RT, see methods; Mouse: $p < 0.017, R = 0.80$, Rat: $p < 0.011, R = 0.83$, Human: $p < 0.033, R = 0.85$). All comparisons are unpaired two-sample t-tests, corrected for multiple comparisons with Bonferroni correction. Notations: * = p<0.05, ** = p<0.01, *** = p<0.005, **** = p<0.001.

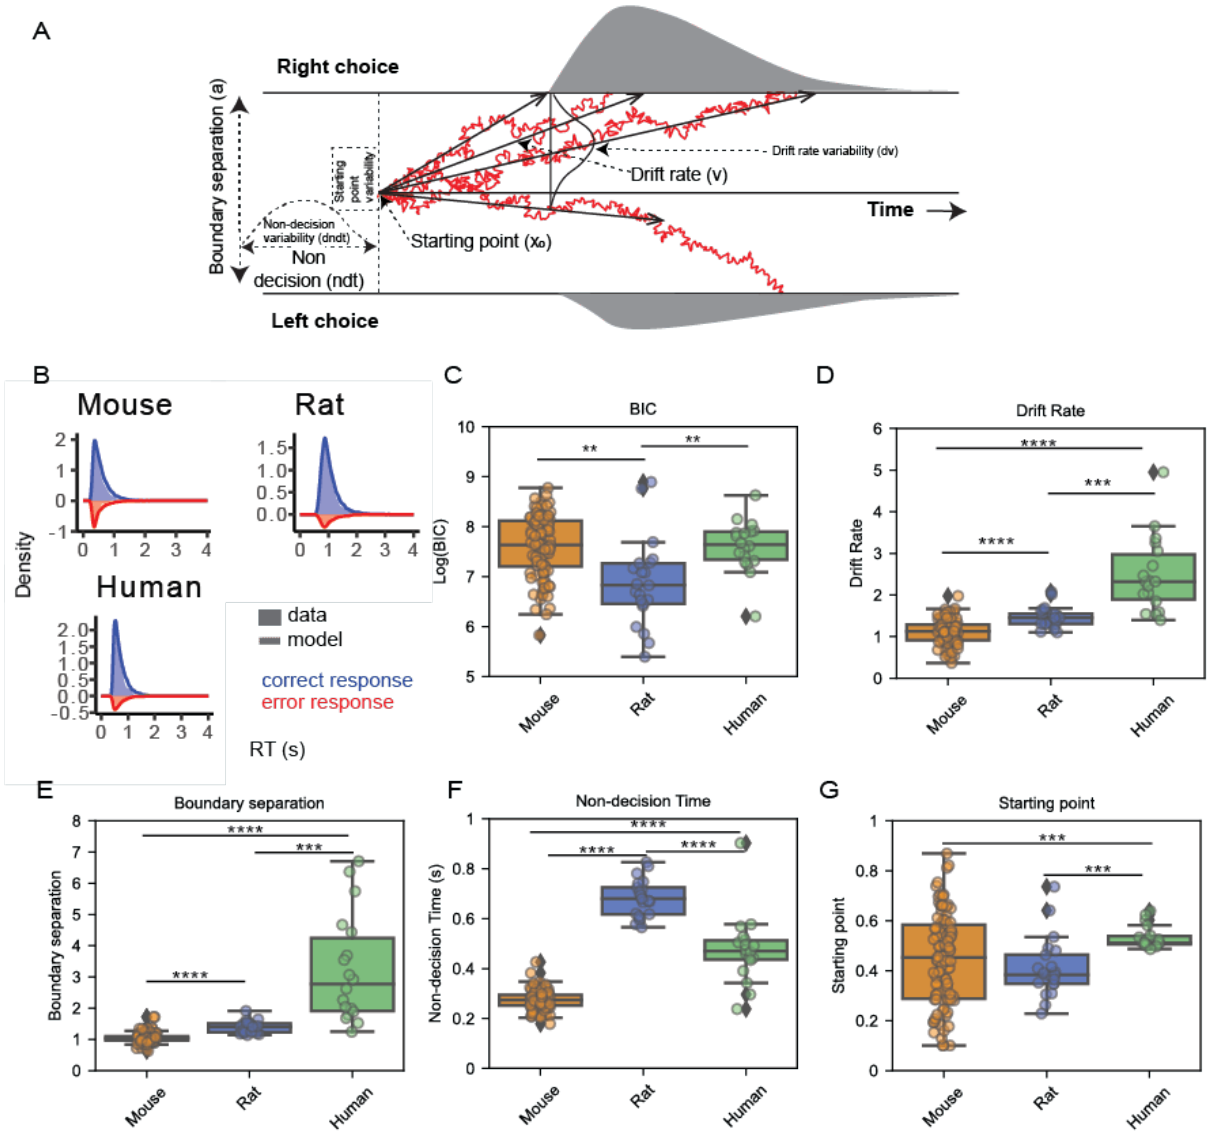

Figure 3: A drift diffusion model was able to explain the response time data for all three species. A) Illustration of the diffusion model: two boundaries indicate decision thresholds for left and right choice, greater distance between boundaries would mean more separability between the two choice decisions. Drift rate is the average rate of evidence accumulation, the model assumes drift rate varies from trial to trial, following a normal distribution. Starting point measures the bias in choosing one side over the other. Starting point is also assumed to vary from trial to trial, following a uniform distribution. Non-decision time indexes time spent in processes other than the decision process. Trial to trial variability in non-decision time follows a half-normal distribution. Evidence on each trial is accumulated until a decision boundary is reached, evidence accumulation is governed by non-decision time, starting point bias, drift rate; drift across time is

subject to a diffusion noise. B) Fit of the diffusion model to one example subject from each species, response time distributions are plotted separately for correct and error responses. C) Comparison of model performance across species. Rats are better fitted by the model than mice and humans (DDM BIC; Humans: 2772 ± 262, Rats: 1492 ± 403, Mice: 2778 ± 136; One-way ANOVA; $p < 0.0005$, $F = 10.97$). D-G) Estimates of model parameters compared across species. Rodents present lower drift (One-way ANOVA; $p < 0.0005$, $F = 77.72$) and boundaries (One-way ANOVA; $p < 0.0005$, $F = 88.29$) than humans. Rats show the slower non-decision time (One-way ANOVA; $p < 0.0005$, $F = 347.26$). All comparisons are unpaired two-sample t-tests, corrected for multiple comparisons with Bonferroni correction. Notations: * = $p<0.05$, ** = $p<0.01$, *** = $p<0.005$, **** = $p<0.001$.

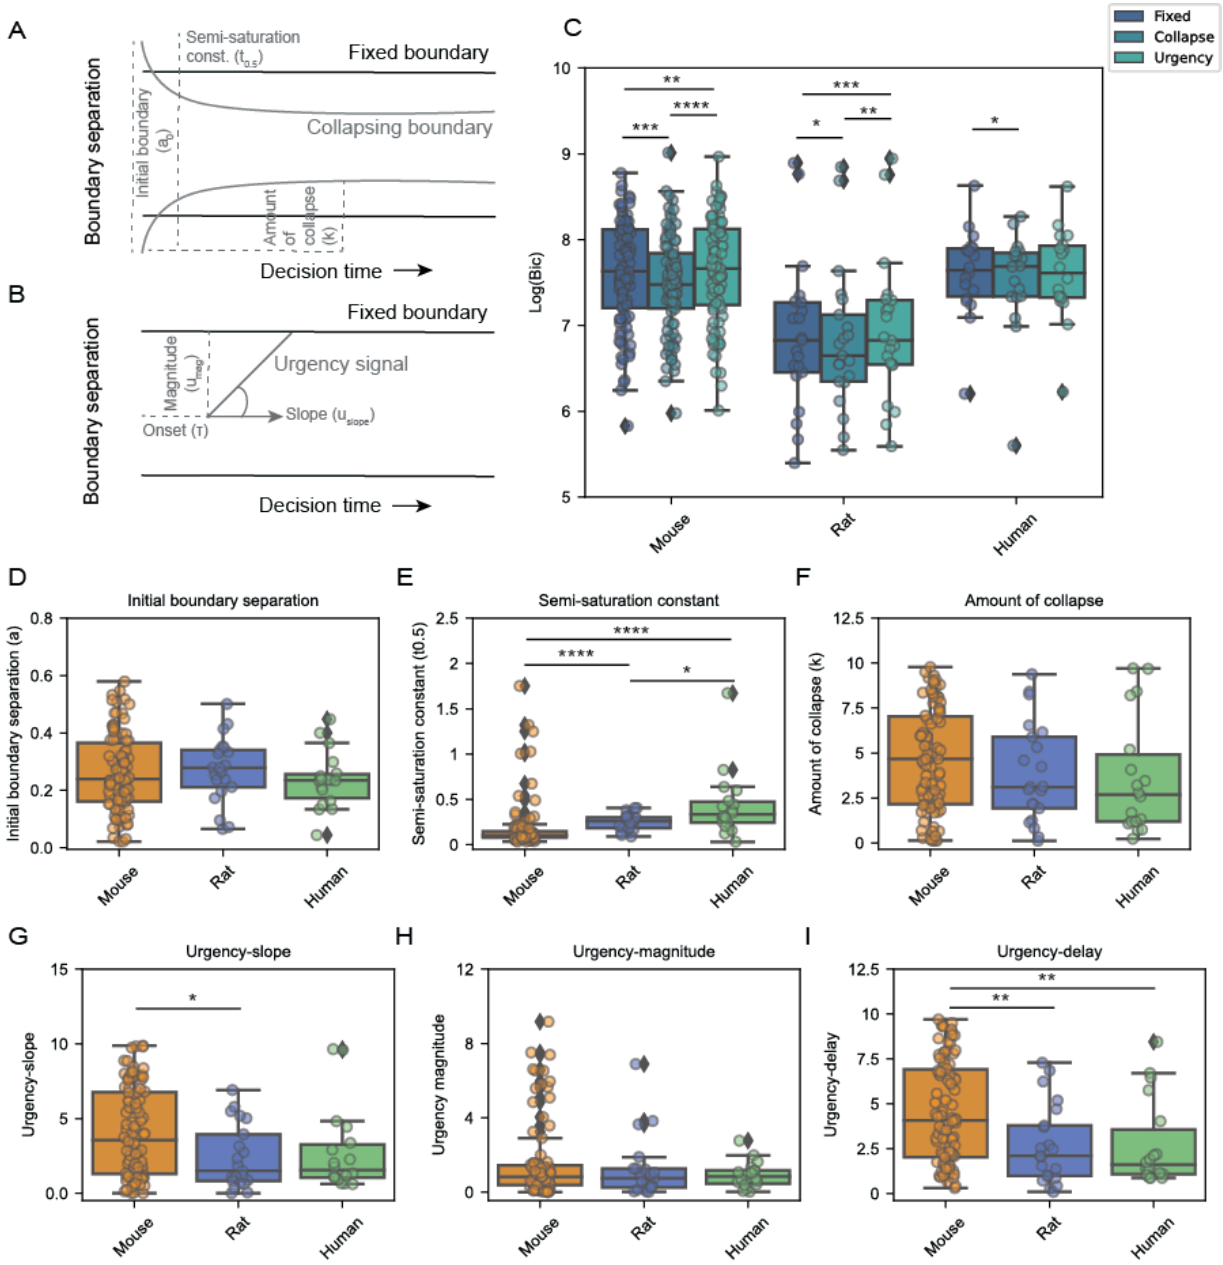

Figure 4: Response times for rats and mice may be better explained by a decision boundary that collapses over time than a fixed decision boundary. A) Illustration of the diffusion model with collapsing boundary, three (free) parameters are used to estimate the boundary: initial boundary separation, semi-saturation constant, and amount of collapse; the model is otherwise the same as the fixed boundary model. B) Illustration of the diffusion model with an urgency signal, three (free) parameters are used to estimate the (linear) urgency signal: slope of the signal, magnitude of the signal, and onset delay for the signal. C) Comparison of model performance for the three diffusion models: fixed boundary model, collapsing boundary

model, and the urgency model, for each species. Rats and mice were better fitted by the collapsing boundaries (Repeated measures ANOVA; Mice: $p < 0.0005$, $F = 22.12$; Rats: $p = 0.0014$, $F = 11.12$). D-F) Comparison of boundary related parameters of the collapsing boundary model, across species. Mice show the lowest semi-saturation constant (One-way ANOVA; $p = 0.004$, $F = 5.69$). G-I) Comparison of urgency related parameters of the urgency model across species. Mice present the highest urgency-slope (One-way ANOVA; $p = 0.027$, $F = 3.71$) and urgency delay (One-way ANOVA; $p = 0.003$, $F = 6.06$). Notations: * = $p<0.05$, ** = $p<0.01$, *** = $p<0.005$, **** = $p<0.001$.

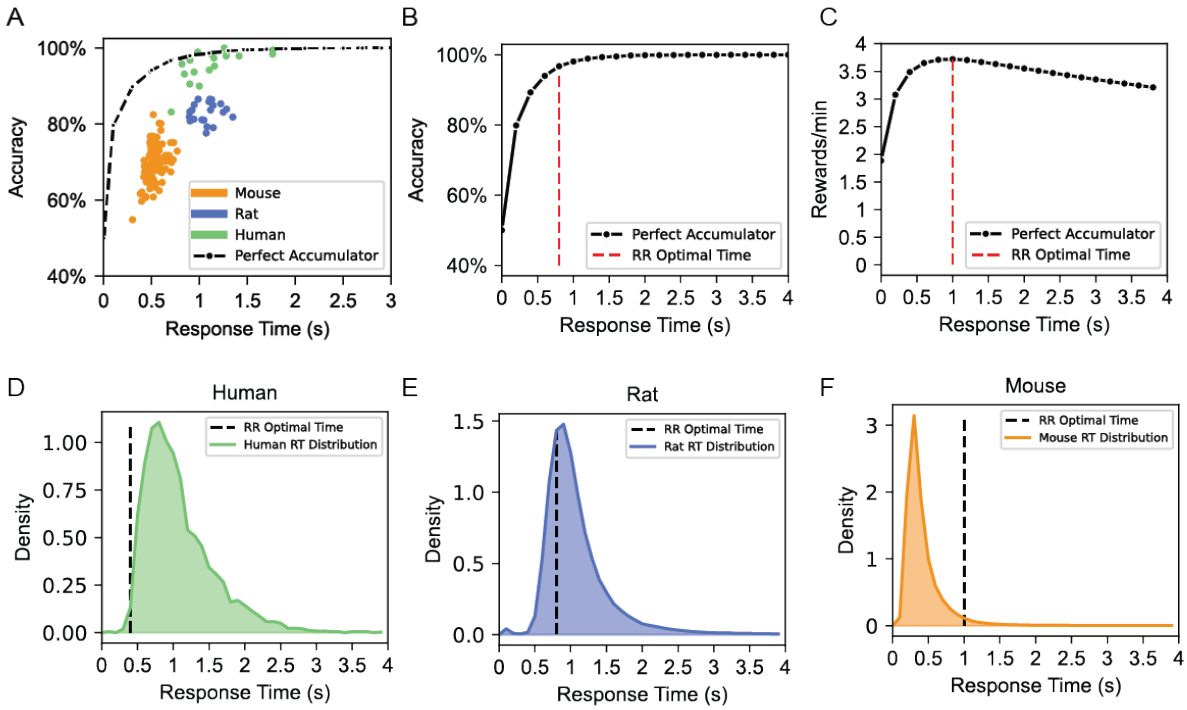

Figure 5: Rats balance speed and accuracy near optimally to maximize the reward rate. A) The speed-accuracy trade-off plot illustrates the relationship between response time and accuracy across species, each dot represents the mean accuracy and response time of an individual. The dashed line represents the accuracy as a function of different choices of RT for a perfect accumulator. B-C) Representations of accuracy (B) and reward rate (C) as a function of different choice of RT for a perfect accumulator, incorporating the average trial initiation time observed in mice, which was different for correct and error responses. In the simulations, flash ratio was kept at 80% vs 20%, and flash rate was 10 Hz. The red line marks the RT associated with the peak reward rate of the accumulator. D-F) Histograms illustrate the distribution of RT observed in humans (D), mice (E), and rats (F) in comparison to the RT associated with the peak reward rate of the perfect accumulator (black dashed lines, Supplementary Figure 5G-I).

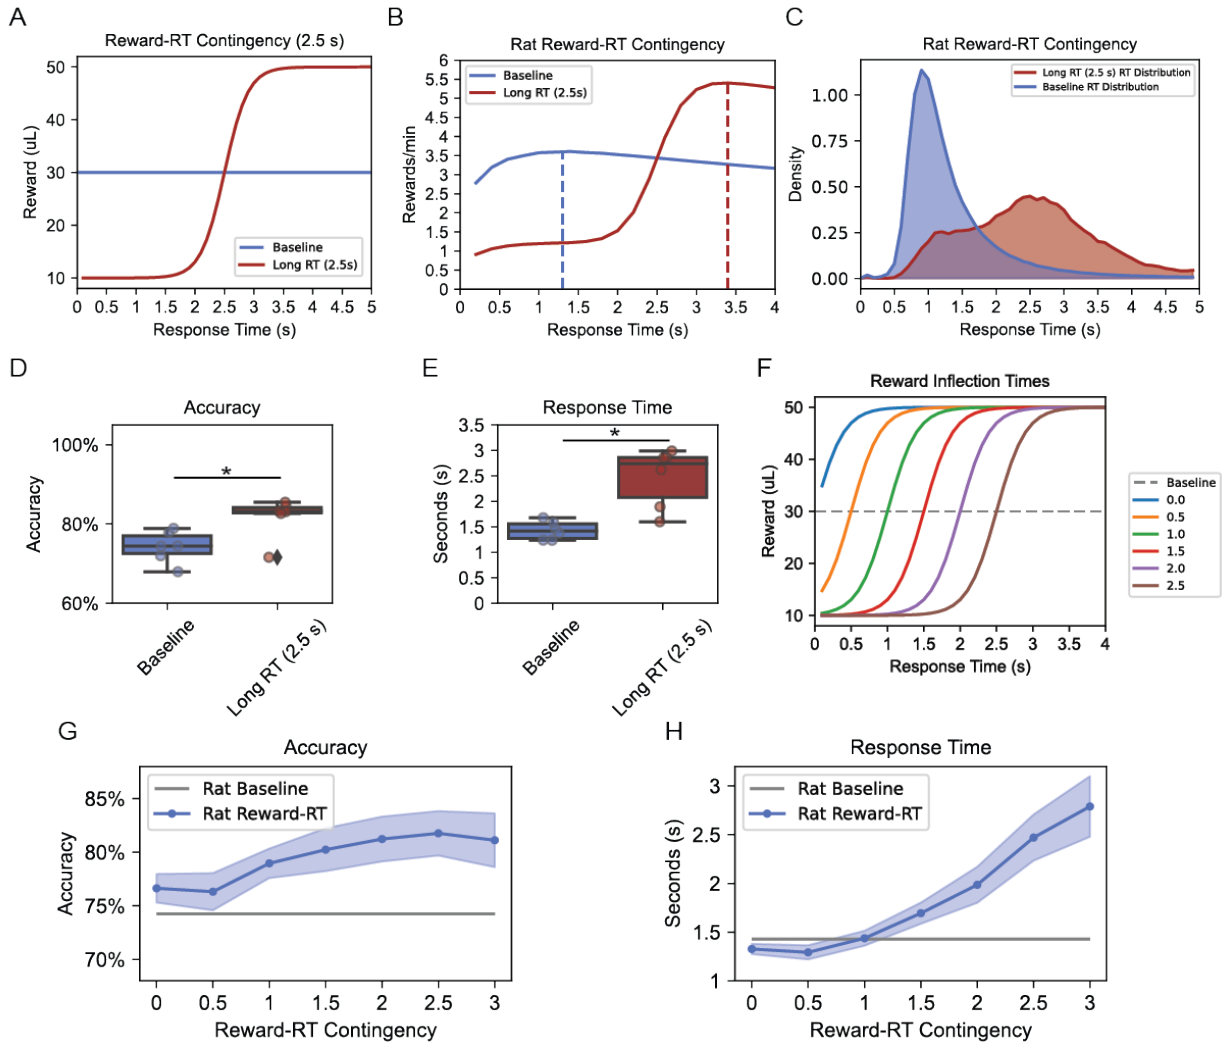

Figure 6: Rats shift their speed-accuracy trade-off to maximize reward rate. A) The Reward-RT contingency graph illustrates the variation of reward with response times, conforming to a sigmoidal curve. The response time at which the reward equals the baseline reward (30 µl) is considered as the inflection point. B) Reward rate of a perfect accumulator, for different choices of RT, and separately for the baseline condition (constant reward), and a long RT condition (inflection point = 2.5 s), these simulations were done considering the ITI of rats. In the simulations, flash ratio was kept at 75% vs 25%. The colored dashed lines represent the response times corresponding to the peak reward rates for each condition. C) Histograms displaying rat response times in the baseline and long RT conditions (n = 6 rats). D-E) Differences in accuracy (Paired t-student; p = 0.025; D) and RT (Wilcoxon test; p = 0.031; of rats between the baseline and a long RT (inflection point = 2.5s) condition. F) Schematic of the seven long RT conditions with different inflection points introduced progressively in the rat experiments. The inflection point is systematically delayed by 0.5 seconds across conditions (see methods for details). G) Accuracy

showed a progressive increase with rise of the inflection point, plateauing at 80% accuracy. H) Response time showed a similar, progressive increase. Notations: * = $p<0.05$, ** = $p<0.01$, *** = $p<0.005$, **** = $p<0.001$.

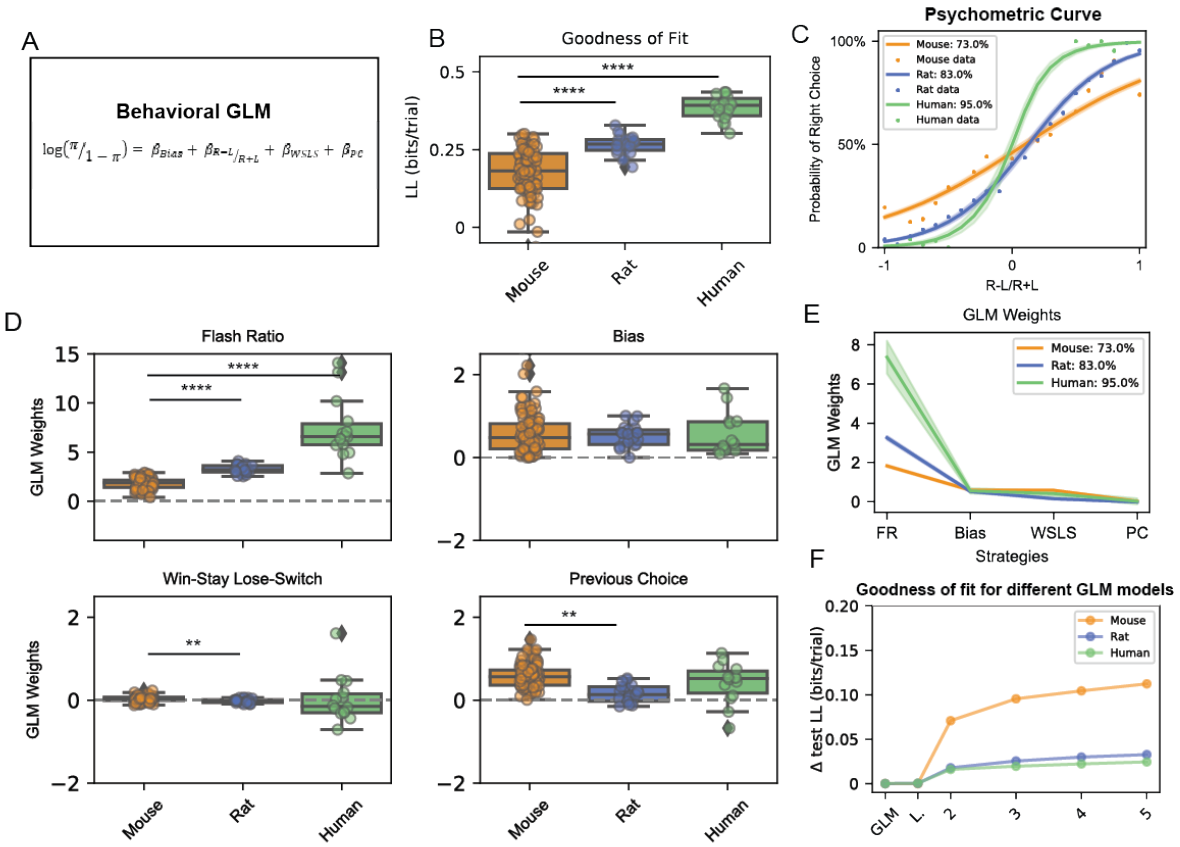

Figure 7: Generalized linear model showed that mice alternate between different strategies while performing the task. A) Generalized linear model with 4 parameters: Bias (Intercept), Flash Ratio (Right flashes - Left Flashes) / Total flashes), Win-stay Lose-switch, and Previous choice. B) Goodness of fit of the GLM presents a better fit for the human data compared to mice or rats. Test log-likelihood is in units of bits per trial, and is relative to a 'null' Bernoulli coin-flip model (One-way ANOVA; $p < 0.0005$, $F = 38.67$). C) Psychometric curve fits, and mean predictive accuracy across species. D) Estimated GLM weights showed that humans relied on the flash ratio more than the rodents, followed by the rats (One-way ANOVA; $p < 0.0005$, $F = 156.28$), whereas mice tended to use a previous choice strategy (One-way ANOVA; $p < 0.0005$, $F = 18.2$). E) Mean estimated GLM weights across species. F) Model comparisons among the one-state GLM, the classic lapse model ('L'), and multiple GLM-HMMs with different numbers of latent states. To compare across species, we measured the change in test log-likelihood, as a function of the number of states included in the model, relative to a one-state GLM, separately for each species. Mice showed greater increase in LL for a GLM-HMM with 3 latent states, suggesting that mice transitioned through multiple latent states while performing the task. All comparisons are based

on unpaired two-sample t-tests, corrected for multiple comparisons with Bonferroni correction. Notations: * = $p<0.05$, ** = $p<0.01$, *** = $p<0.005$, **** = $p<0.001$.

Supplementary Figures

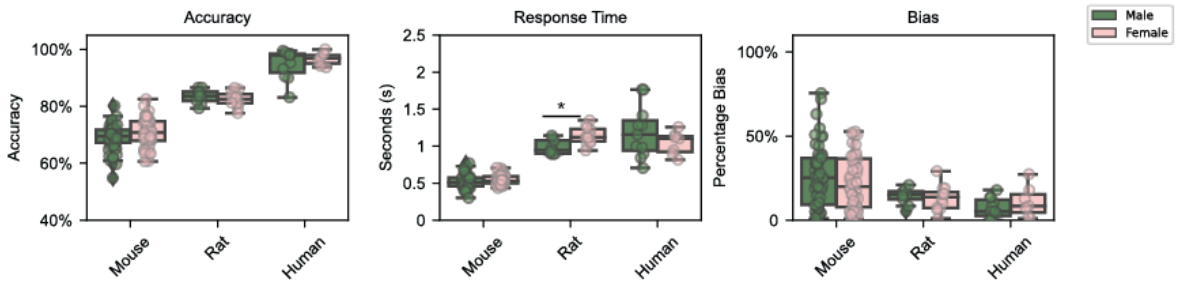

Supplementary Figure 1: Comparison of accuracy, response time and bias across species, separately for female and male subjects. Notations: * = p<0.05, ** = p<0.01, *** = p<0.005, **** = p<0.001.

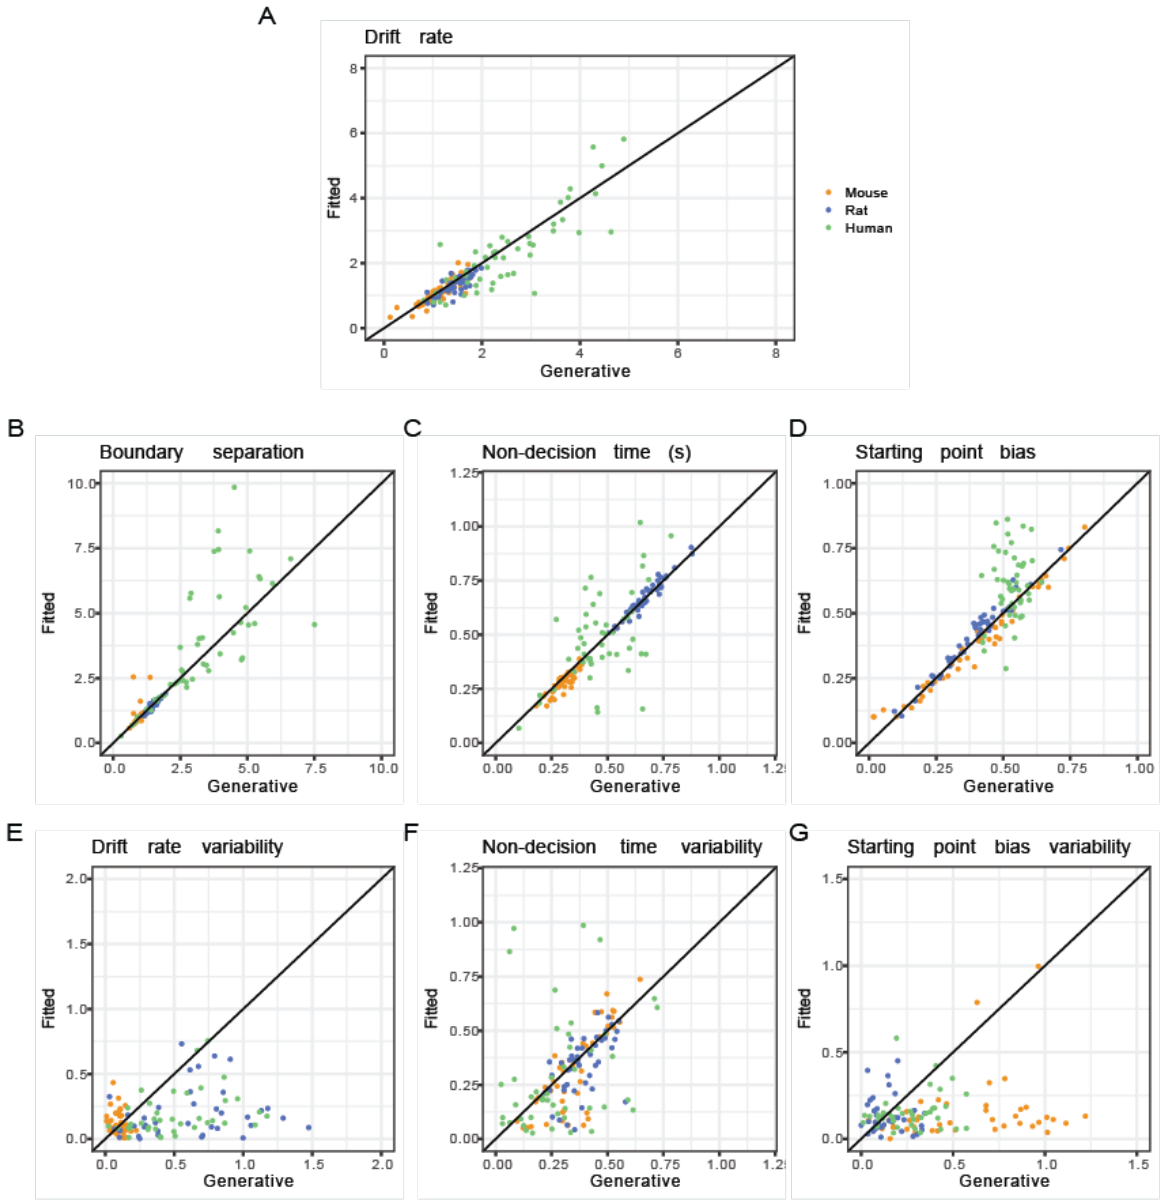

Supplementary Figure 2: Parameter recovery for the DDM. Based on the parameters observed with the empirical data (see Figure 3-4), we simulated 50 datasets per species with 600 trials each. After fitting the DDM to each of these simulated datasets, we compared the generative and fitted parameters.

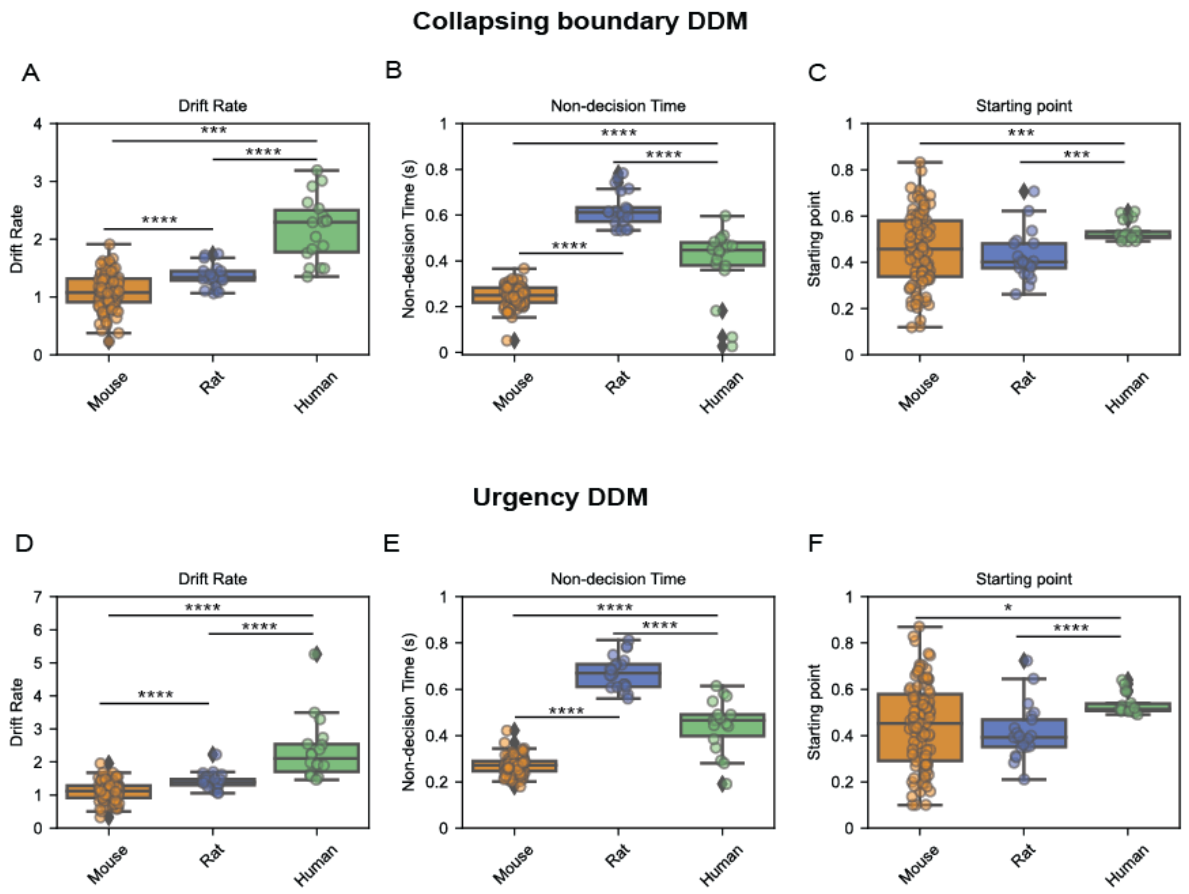

Supplementary Figure 3: Comparison of drift rate, non-decision time and starting point bias across species, separately for the collapsing boundary DDM and the urgency DDM. Both models present similar parameter estimates and cross-species trends as the base DDM. Notations: * = p<0.05, ** = p<0.01, *** = p<0.005, **** = p<0.001.

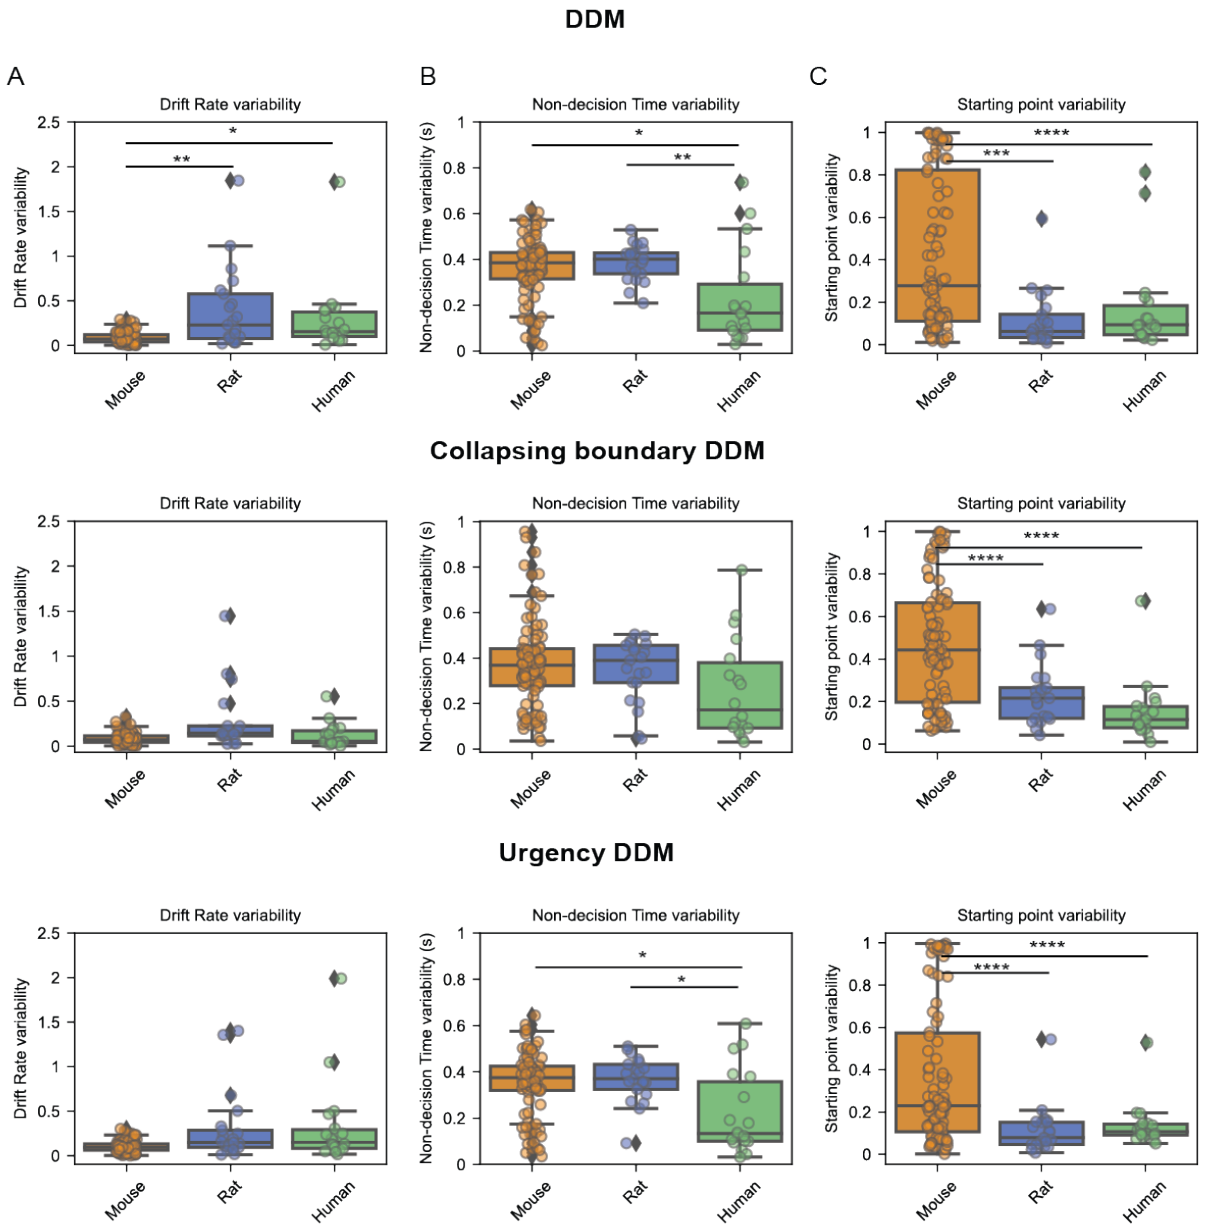

Supplementary Figure 4: Comparison of DDM parameters that control the trial-to-trial variability of its main parameters, such as drift rate, non-decision time and starting point bias. We assumed variability of drift rate from trial to trial follows a normal distribution, that for non-decision time follows a half-normal distribution, and for starting point bias, trial-to-trial variability follows a uniform distribution. A–C) corresponds to the original DDM, D–F) corresponds to the collapsing boundary model, G–I) corresponds to the urgency DDM. Notations: * = p<0.05, ** = p<0.01, *** = p<0.005, **** = p<0.001.

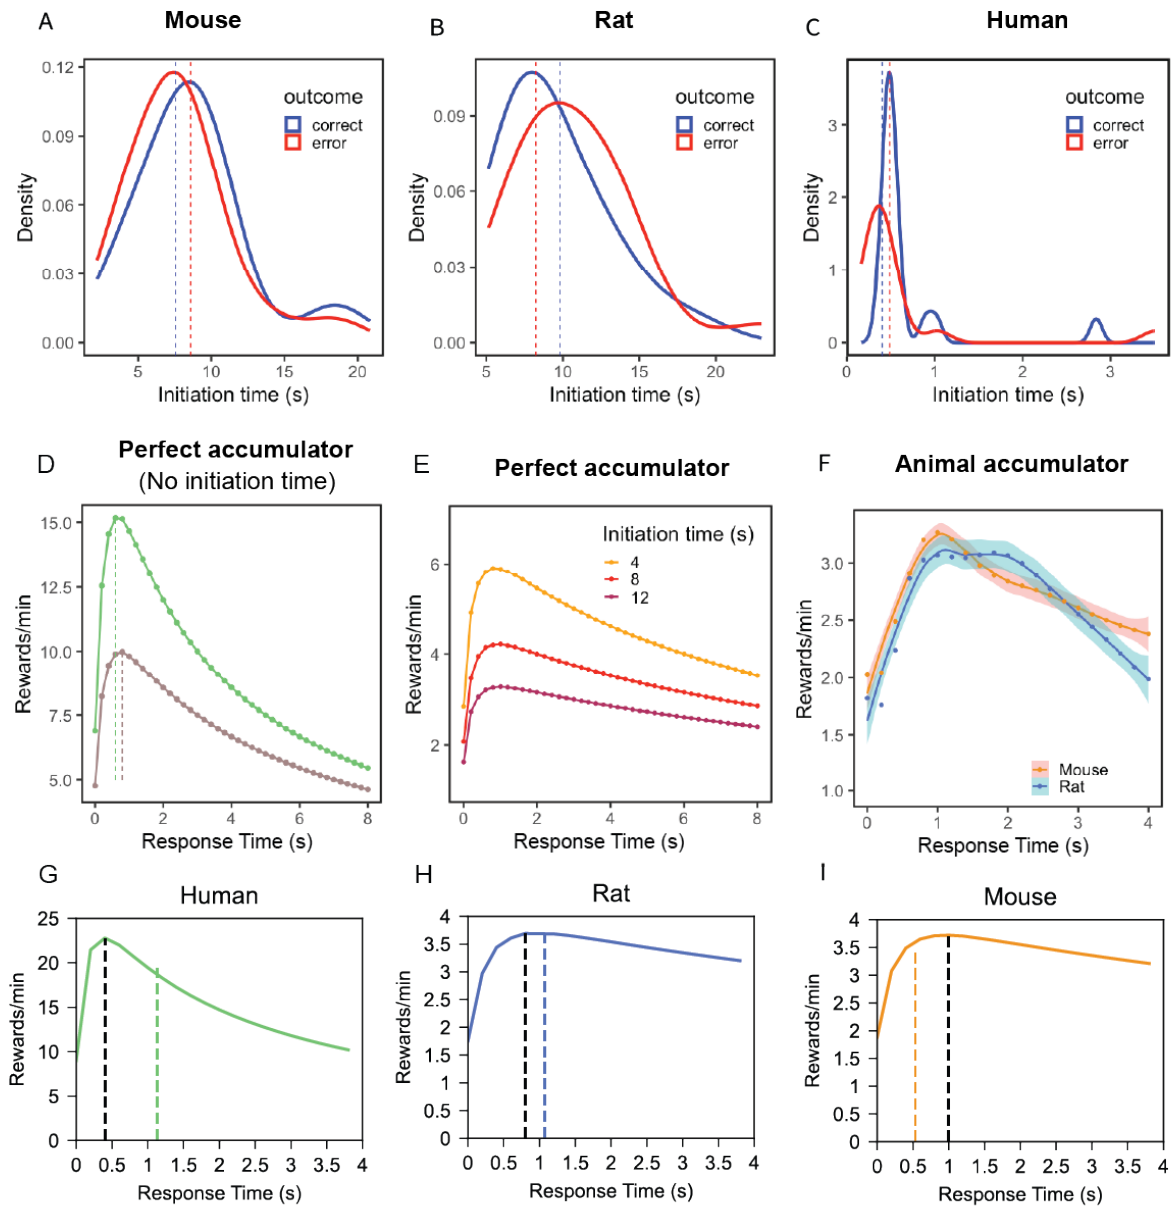

Supplementary Figure 5: Differences in intertrial intervals across species and its effects on the reward rate simulations. A-C) Distributions of initiation times across all correct and erroneous responses, separately for mice, rats and humans. Initiation time was the period between the start of a trial and the center poke, after the center poke, flash stimuli were presented; for humans this was the time to choose an asteroid. D) Reward rate (rewards/min) as a function of different choices of RT for a perfect accumulator, simulated separately for humans (green line) and rodents (purple line). RT corresponding to the peak reward rate is noted. E) Reward rate of a perfect accumulator taking into account different initiation times. F) Reward rate as a function of different choices of RT for noisy (animal) accumulators. Separate simulations were run to

account for average trial initiation times observed in rats and in mice, for correct and error responses. In the simulations, flash ratio was kept at 80% vs 20%, and flash rate was 10 Hz, same as that for the actual experiments. In addition to initiation time, choice accuracy for at different RT levels was predicted by the GAMM, previously fit to animal choice/RT data (see Figure 2F). G-I) Reward rate as a function of different choices of RT for humans (G), mice (H) and rats (I) based on a perfect accumulator, incorporating the average trial initiation time observed in each species, for correct and error responses. The black dashed line represents the response time where the reward rate of the accumulator peaks and the colored line is the average response time of each species.

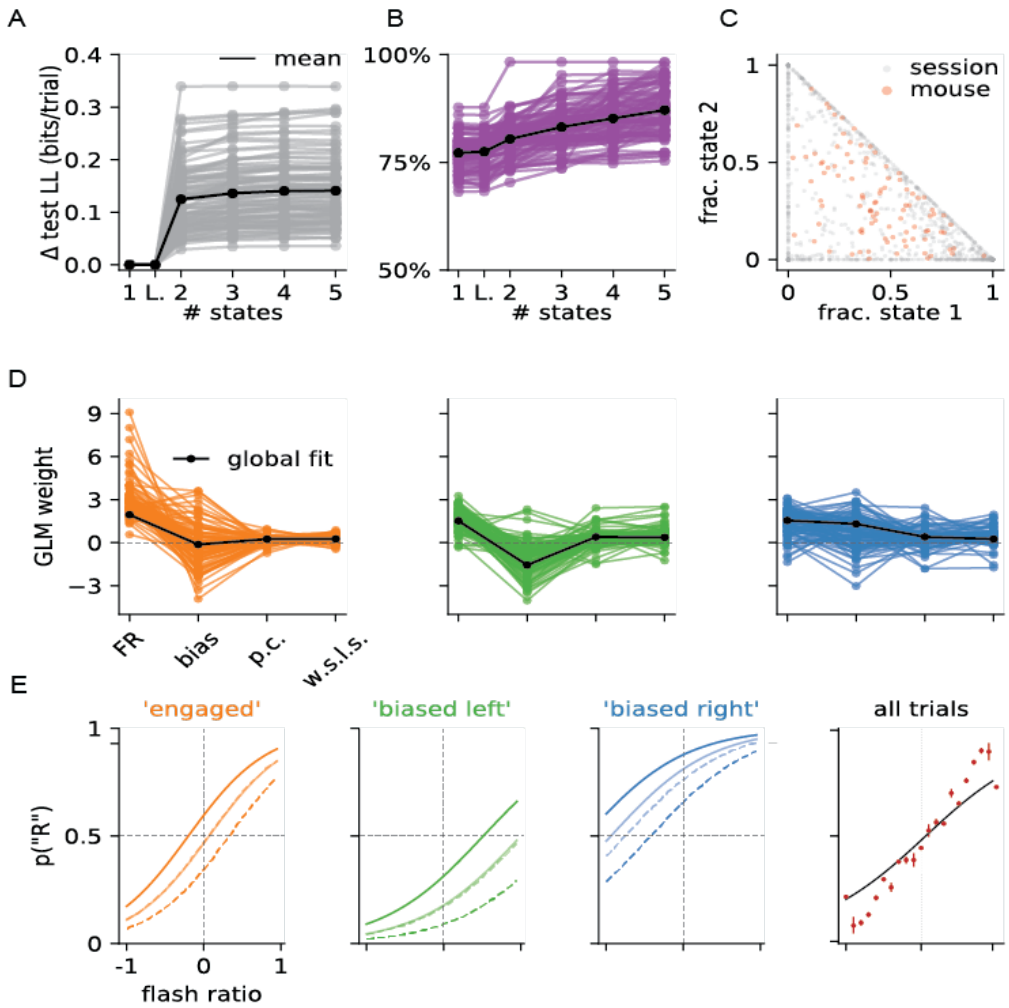

Supplementary Figure 6: Mouse data is better fit with a hidden Markov model with 3 latent states. A) Change in test log-likelihood as a function of the number of states included in the GLM-HMM, relative to a (one-state) GLM, modeled separately for each mouse. The classic lapse model, a restricted form of the two-state model, is labeled as 'L'. Each trace represents a single mouse and the solid black line indicates the mean across animals. B) Change in predictive accuracy relative to a one-state GLM for each mouse, indicating the percentage improvement in predicting choice. C) Gray dots correspond to individual sessions across all mice, indicating the fraction of trials spent in state 1 (engaged) and state 2 (biased left). Points at the vertices (1,0), (0,1) or (0,0) indicate sessions with no state changes, whereas points along the sides of the triangle indicate sessions that involve only two of the three states. Red dots correspond to the same fractional occupancies for each of the mice, revealing that the engaged state predominated but that all mice spent time in all three states. D) Inferred GLM weights for each mouse, for each

of the three states in the three-state model. The solid black curve represents a global fit using pooled data from all mice. E) Psychometric curve for each state, conditioned on previous reward (solid line if right and dashed line if left) and previous choice (darker color if rewarded and dashed line if not rewarded) and all the states combined.

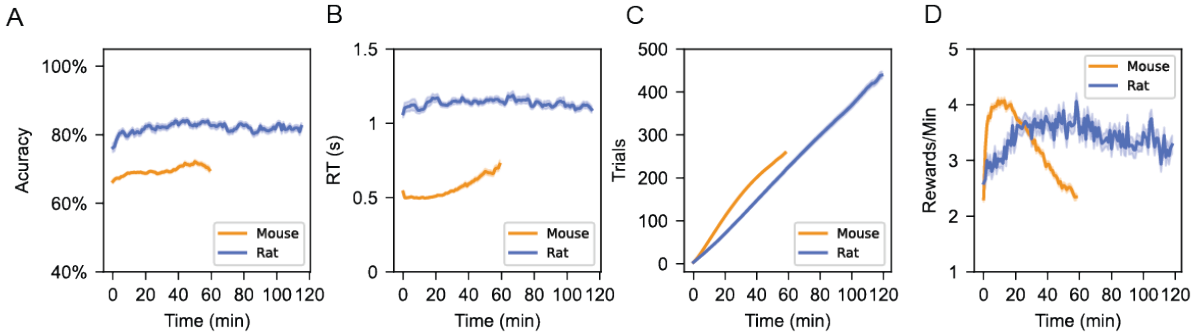

Supplementary Figure 7: Rats exhibited stability throughout the entire session, whereas mice displayed greater variability. A-D) Averaged accuracy (A), response time (B), number of trials (C), and reward rate (D) across a session. It is important to note that mouse sessions lasted for 60 minutes and rat sessions for 120 minutes.