A qualitative difference in decision-making of rats vs. humans explained by quantitative differences in behavioral variability

Quynh Nhu Nguyen[1], Pamela Reinagel[1]*

[1] Section of Neurobiology, Division of Biological Sciences, University of California San Diego, La Jolla, California, United States of America.

* Corresponding author

Email: preinagel@ucsd.edu (PR)

Abstract

When observers make rapid, difficult sensory decisions, their response time is highly variable from trial to trial. We previously compared humans and rats performing the same visual motion discrimination task. Their response time distributions were similar, but for humans accuracy was negatively correlated with response time, whereas for rats it was positively correlated. This is of interest because different mathematical theories of decision-making differ in their predictions regarding the correlation of accuracy with response time. On the premise that sensory decision-making mechanisms are likely to be conserved in mammals, our objective is to reconcile these results within a common theoretical framework. A bounded drift diffusion model (DDM) with stochastic parameters is a strong candidate, because it is known to be able to produce either late errors like humans, or early errors like rats. We consider here such a model with seven free parameters: the evidence accumulator's starting point z, drift rate v, non-decision-time t, threshold separation a, and three noise terms σ_z, σ_v and σ_t. We fit the model parameters to data from both rats and humans. Trial data simulated by the model recapitulate quantitative details of the relationship between accuracy and response time in both species. On this model, the species difference can be explained by greater variability in the starting point of the diffusion process (σ_z) in rats, and greater variability in the drift rate (σ_v) in humans.

Introduction

One might expect decision-making by humans to be quite different from that of rats. In decisions with wide-reaching long-term consequences, we expect (or at least wish) humans would avail themselves of abstract conceptual thought, logical reasoning, and culturally accumulated knowledge that would be unavailable to a rat. Yet all organisms face a continuous challenge of selecting among alternative available actions in order to pursue goals. In order to select an action, sensory information, internal knowledge, and goals are combined to assess and evaluate the likely outcomes of possible actions relative to survival needs. Often there is not enough time to acquire the evidence necessary to determine with certainty the optimal course of action, so an action must be selected despite unresolved or unresolvable uncertainty. Some mechanism is needed to ensure timely commitment and to optimize outcome on average, and this must adapt flexibly to prevailing sensory context, shifting goal priorities, the urgency of action, and the severity of consequences of errors. When it comes to the continuous sensory guidance of moment-by-moment actions, decisions about sensory evidence are made in a fraction of a second. We speculate that in this case, mechanisms are largely conserved across mammals.

A now-classic series of studies in humans and non-human primates introduced the use of a stochastic visual motion task to study decision making [1-12]. In each trial a visual stimulus provides

information regarding which of two available actions is associated with reward and which is associated with non-reward or penalty. Stimulus strength is modulated by the motion coherence, which is defined as the fraction of the dots in the display that are "signal" (moving toward the rewarded response side). The remaining dots are "noise" (moving in random directions). As stimulus strength increases, accuracy increases and response time decreases for both monkeys [10] and humans [9]. This is parsimoniously explained by drift diffusion models, which postulate that noisy sensory evidence is integrated over time until the accumulated evidence reaches a decision threshold [5, 6, 9, 13-20]. Although this class of model is highly successful, more data are needed to test model predictions and differentiate among competing versions of the model or alternative model classes [14, 15, 21-31].

When monkeys or humans perform this task, among trials of the same stimulus strength the interleaved trials with longer response times are more likely to be errors [9, 10]. In its simplest form the drift diffusion model does not explain this result; therefore the observation has been an important constraint for recent theoretical efforts. The result can be explained if the decision bound is not constant but instead decays as a function of time [32, 33]. A collapsing decision bound can be rationalized as an optimal strategy under some task constraints [22, 30, 34, 35] though this argument has been challenged by others [36, 37]. There are alternative ways to explain the data within the drift diffusion model framework without positing an urgency signal or decaying bound [14, 38-40].

In the interest of developing a rodent model for perceptual decision-making, we trained rats to perform the random dot motion task [41]. Rats differed from humans and other primates in one key respect: for rats, later decisions are more likely to be accurate [41, 42]. The same has also been reported for image discriminations in rats [43], for visual orientation decisions in mice [44], and is found in humans in some other tasks [45-48]. This result is not readily explained by some models suggested to explain the late errors of primates (reviewed in [15, 49, 50]). The purpose of this study is to test one variant of the drift-diffusion model [14, 16] for its ability to explain data from both species performing a random dot motion visual decision task.

In a standard drift diffusion model (DDM; also called the Decision Diffusion Model), the relative sensory evidence in favor of a decision (e.g. "motion is rightward" vs. "motion is leftward") is accumulated by an internal decision variable, resulting in a biased random walk, i.e., diffusion with drift (Fig 1A). The strength of the drift is determined by the sensory signal to noise ratio (e.g. the strength of a visual motion signal). When the decision variable reaches either decision threshold, the agent commits to a choice. The time at which the decision variable crosses a threshold (response time), and the identity of the decision threshold that is crossed (correct vs. incorrect), vary from trial to trial; the distribution of response times and the probability of correct response depend on the drift rate and the starting point of the accumulator. The asymmetric shape of the response time distribution of this model (Fig 1B,D) resembles that observed in experimental data.

In this simple form, the drift diffusion model depends on only four parameters: the threshold separation $\boldsymbol{a}$, non-decision-time $\boldsymbol{t}$, accumulator starting point $\boldsymbol{z}$, and drift rate $\boldsymbol{v}$ (Fig 1F). The experimental data that constrain these parameters are the accuracy (percent of decisions that are correct) and the shape of the distribution of response times, for error trials and correct trials at each stimulus strength tested.

An interesting feature of this model is that for any given stimulus strength and decision threshold, errors and correct responses have the same response time distributions (Fig 1B-E, red vs.

green). Errors are on average the same speed as correct responses, even if the signal is so strong that errors are very rare.

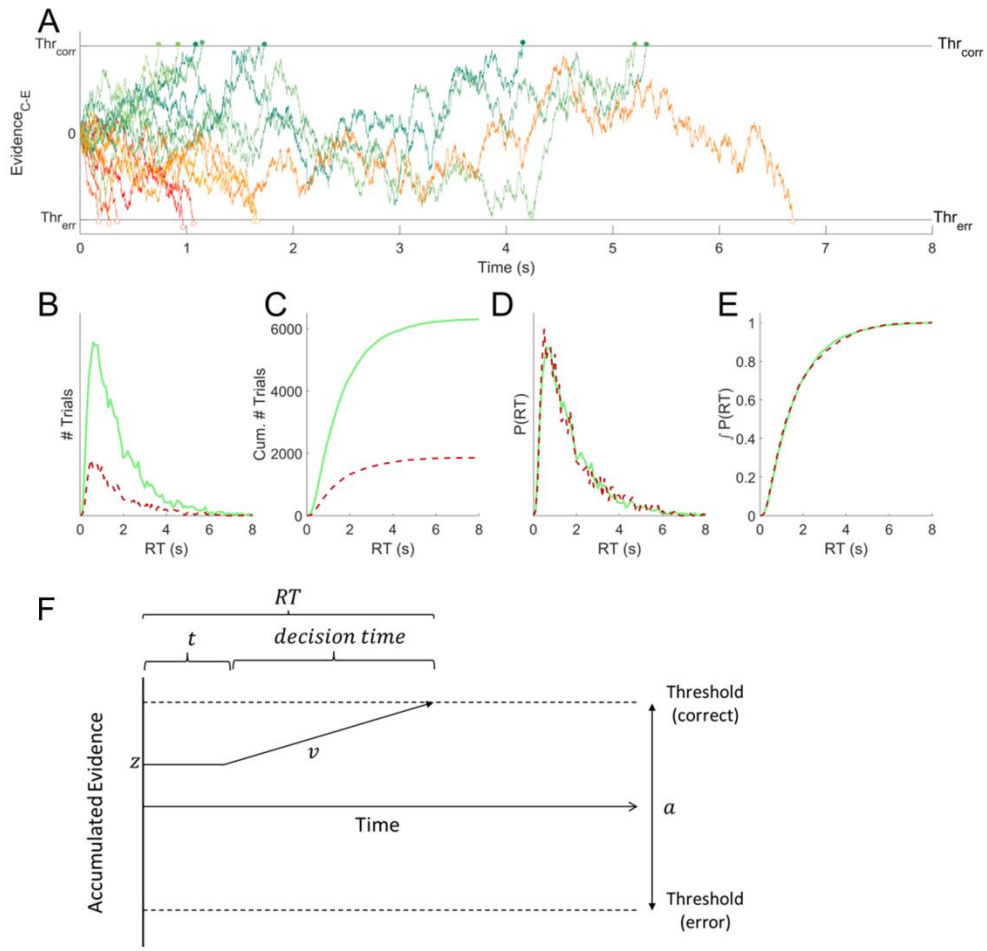

Fig 1. Illustration of a standard drift diffusion model. (A) Noisy sensory evidence is accumulated over time, resulting in a random walk in each trial; a decision occurs when the accumulated evidence reaches one of the decision thresholds. The traces show eight randomly chosen example trials that terminated in correct decisions (green shades, solid symbols) and eight that terminated in error decisions (red shades, open symbols). Although an equal number of traces are shown, in the condition illustrated 77% of the traces terminated by crossing the correct boundary. (B) Response time distributions for errors (red, dashed) vs. correct trials (green, solid) of trials simulated as in A. (C) Cumulatives of the distributions shown in B. (D,E) Distributions in B and C normalized to the number of trials, to better compare shapes. (F) Definition of parameters for a standard drift diffusion model. The parameter a is the distance between the error and correct thresholds. The starting point z is expressed as a fraction of a, such that z=0.5 is an unbiased or neutral starting point. The drift rate, v, depends on the sensory signal strength. These parameters are constant during a trial and across all trials of the same stimulus strength. The observable outcomes are response time (RT) and decision (correct or error). The non-decision-time t reflects both sensory latency and motor latency, but is drawn at time t=0 for graphical simplicity. Decision time is defined as RT-t. For the simulations illustrated in A-E, t=0.

Because the ratio of the two curves (Fig 1C) is fixed, the accuracy of the decision is independent of response time (RT). We note that this does not, but may at first appear to, contradict two other facts. First, responses to stronger stimuli are both more accurate and faster, which in this model is explained

by an increased drift rate of the random walks. In this sense response time is negatively correlated with accuracy – but only when comparing trials of differing stimulus strengths. Second, conservative subjects take more time to respond and are more accurate, which in this model is explained by an increased separation between the decision thresholds. In this sense response time is positively correlated with accuracy – but only when comparing trials in blocks of trials with different overall degrees of caution. Both of these facts are consistent with the claim that within a block of fixed overall caution, comparing among the trials of the same stimulus strength, response time and accuracy are not correlated.

Both humans and rats deviate systematically from the prediction that correct and error trials have the same probability distribution, however. In the random dot motion discrimination task, for example, correct trials of rat subjects are more likely to have longer response times compared to errors (Fig 2A). If we compare the response time of each error trial to the temporally closest correct trial of the same stimulus strength [42], we see that on average, correct trials have longer response times than neighboring error trials of the same stimulus strength (Fig 2B). This trend is found in most individual experiments on rats (Fig 2C). If one analyzes accuracy as a function of response time for stimulus strengths near the psychophysical threshold, the slope is positive (Fig 2D), indicating that accuracy is increasing with RT.

Humans also violate the model prediction, but in the opposite way. For humans, errors are more likely to have longer response times (Fig 2E). On average correct trials have shorter response times than adjacent error trials of the same stimulus strength (Fig 2F). This effect is found in most individual experiments on humans (Fig 2G). The slope of accuracy vs. response time is negative (Fig 2H); in other words, accuracy declines with RT. This result has been widely reported in studies of both humans and monkeys [9, 10].

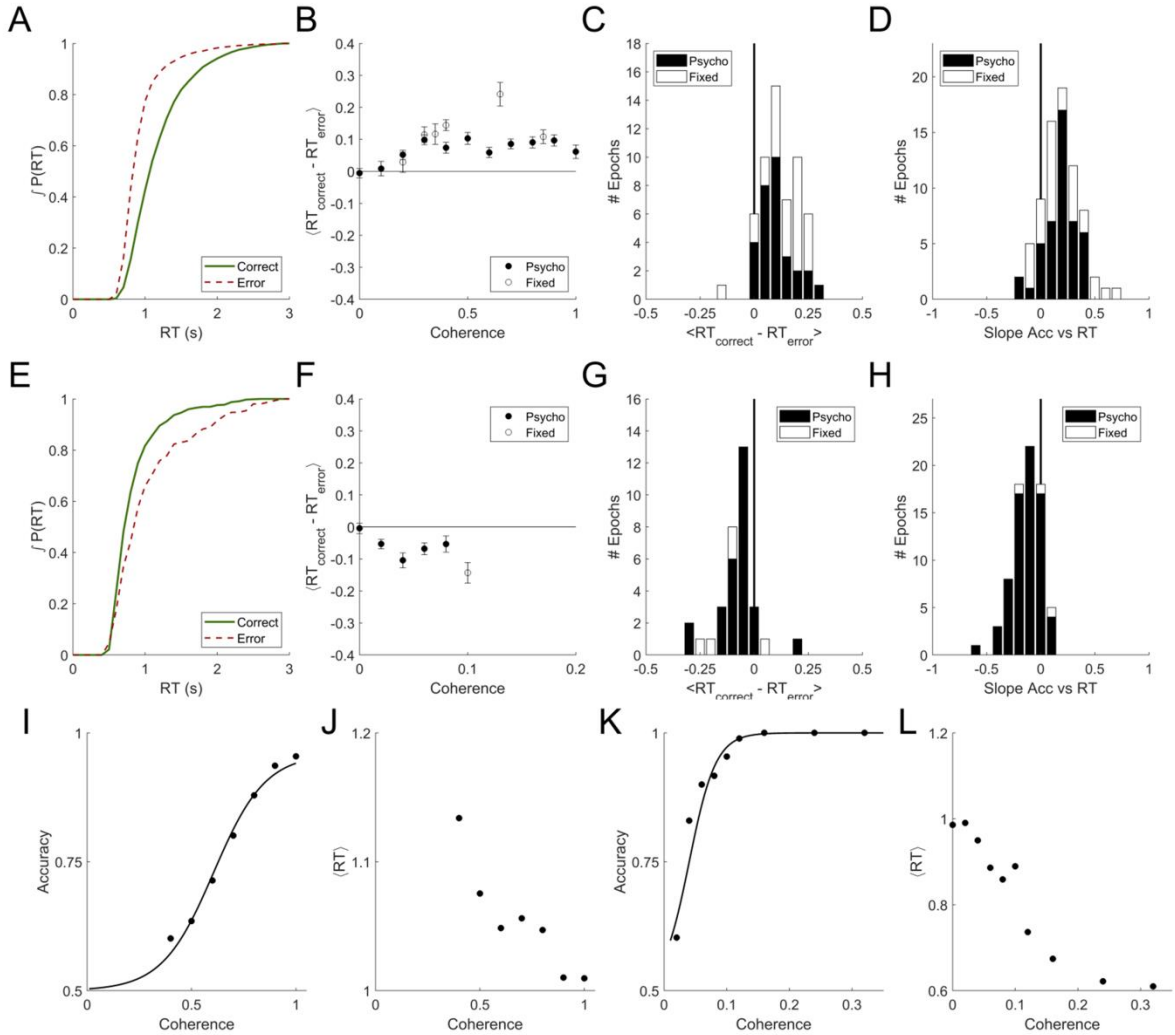

Fig 2. Summary of experimental data from rats and humans. (A) Cumulative probability function of response time for errors (red, dashed) and correct trials (green, solid) for one example experiment in a rat, a fixed-coherence epoch with an easy motion stimulus for which the rat was 82% correct. The red dashed curve is to the left, indicating that early responses are more likely to be errors. (B) The average difference in response time between a correct trial and a temporally nearby error trial of the same motion strength; positive values indicate that correct trials had longer response times. This was computed within experimental epoch and then averaged over epochs (minimum 2 epochs). Coherences were either randomly interleaved during the experiment (closed symbols, N=46) or fixed within the epoch (open symbols, N=43). Error bars show standard error of the mean (SEM) over epochs contributing data at that coherence. (C) Distribution over epochs of the average difference between correct and error RT, for coherences ≥0.4. Psychometric (mixed coherence) epochs (black, N=30) and fixed-coherence epochs (white, N=26). (D) Distribution over epochs of the slope of accuracy vs. RT for coherences with accuracy between 70-90%. Trials were divided into 6 quantiles of RT, and accuracy and mean RT computed within quantile. A line was fit to the resulting points for which mean RT was ≤1.5s (minimum 3 points), and the slope of this line is reported. Psychometric epochs (black, N=46) and fixed-coherence epochs (white, N=43). (E) Like A, for one example experiment in a human subject, a fixed-coherence epoch with an easy coherence for which the subject was 83% correct. (F) Like B, for human subjects (N=75 psychometric and N=5 fixed-coherence epochs). Coherences up to 0.35 were tested in humans, but for coherence >0.2 there were not enough errors for this analysis. (G) Like C, but for human subjects, using coherences ≥0.04 (N=64 mixed coherence and N=4 fixed-coherence epochs). (H) Like D, but for human subjects (N=75 mixed coherence and N=5 fixed-coherence epochs). (I) Accuracy vs. coherence at tested coherences for one example rat epoch, and psychometric curve fit. (J) Mean response times at each coherence, for the same epoch analyzed in I. (K-L) Like I-J, for one example human psychometric epoch. See Statistical Methods for additional details. For analysis excluding epochs with lapse or bias see S5 Appendix. Re-analysis of data from [42, 51].

Results

Standard Drift Diffusion Model

We fit the standard drift diffusion model (Fig 1F) to the behavioral data from both rats and humans. We used a hierarchical Bayesian estimation method to estimate epoch-level values for each parameter: non-decision-time $\boldsymbol{t}$, drift rate parameters ($\boldsymbol{v_1 \ldots v_n}$) one for each coherence, and threshold separation $\boldsymbol{a}$. In the experiments, the correct response side was left or right with equal probability independently each trial. Although subjects could certainly be biased to the left or right response, it is impossible to be systematically biased towards or away from the *correct* response threshold. Therefore we fixed the starting point $\boldsymbol{z}$ at 0.5 (no bias) with respect to the correct answer (but see S4 Appendix).

The distributions of parameters estimated for each population are described in Fig 3. It should be noted, however, that each epoch is characterized by a specific combination of parameters, not by random sampling from these distributions.

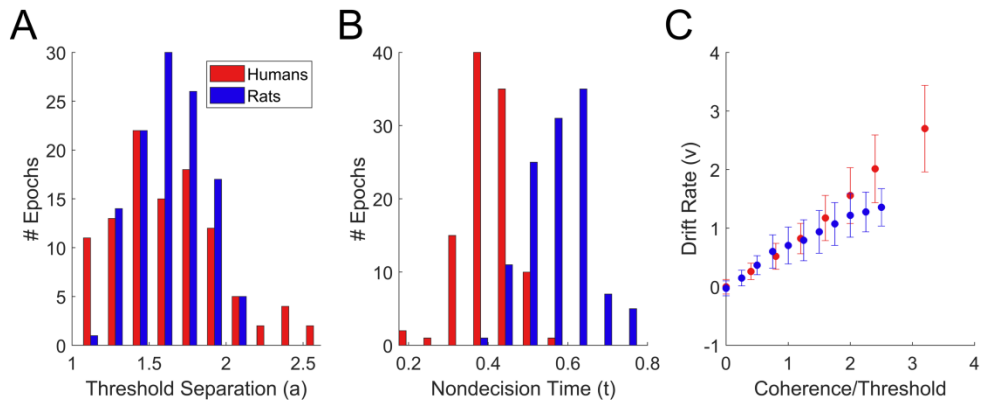

Fig 3. Estimated parameters of a standard drift diffusion model fit to data from humans or rats. All trend-stationary epochs of all subjects were included (N=104 epochs from humans, red; N=115 epochs from rats, blue). The starting point was constrained to be neutral with respect to the correct response (z=0.5). (A) Distributions of threshold separation parameter a: a_{human} = 1.62±0.35, N=104 (mean±SD over the population); a_{rat} = 1.62±0.22, N=115. (B) Distributions of non-decision-time t: t_{human} = 0.40±0.06 s, t_{rat} = 0.57±0.07 s. (C) Drift rate parameter v as a function of stimulus strength, expressed as motion coherence divided by the median psychophysical coherence threshold of the species: 0.05 (humans) or 0.40 (rats). Psychophysical threshold is defined as the motion coherence at which accuracy is 70%, based on a psychometric fit of accuracy vs. coherence from the same experiments [42]. Symbols indicate population mean, error bars indicate population SD. Humans and rats were tested with different coherences so drift rates (v) cannot be directly compared. See Statistical Methods, S1 Appendix and S2 Data for additional details.

We used the parameters fit to each experimental epoch to simulate trials from the model, and subjected the simulated population data to the identical analysis as the real data. The model successfully approximated the psychometric curves and the average response time distributions of individual subjects (Fig 4I-L), but failed to replicate the difference between error and correct trials within condition for either humans or rats (Fig 4A-H). This result is expected, and serves as a baseline to which other models may be compared.

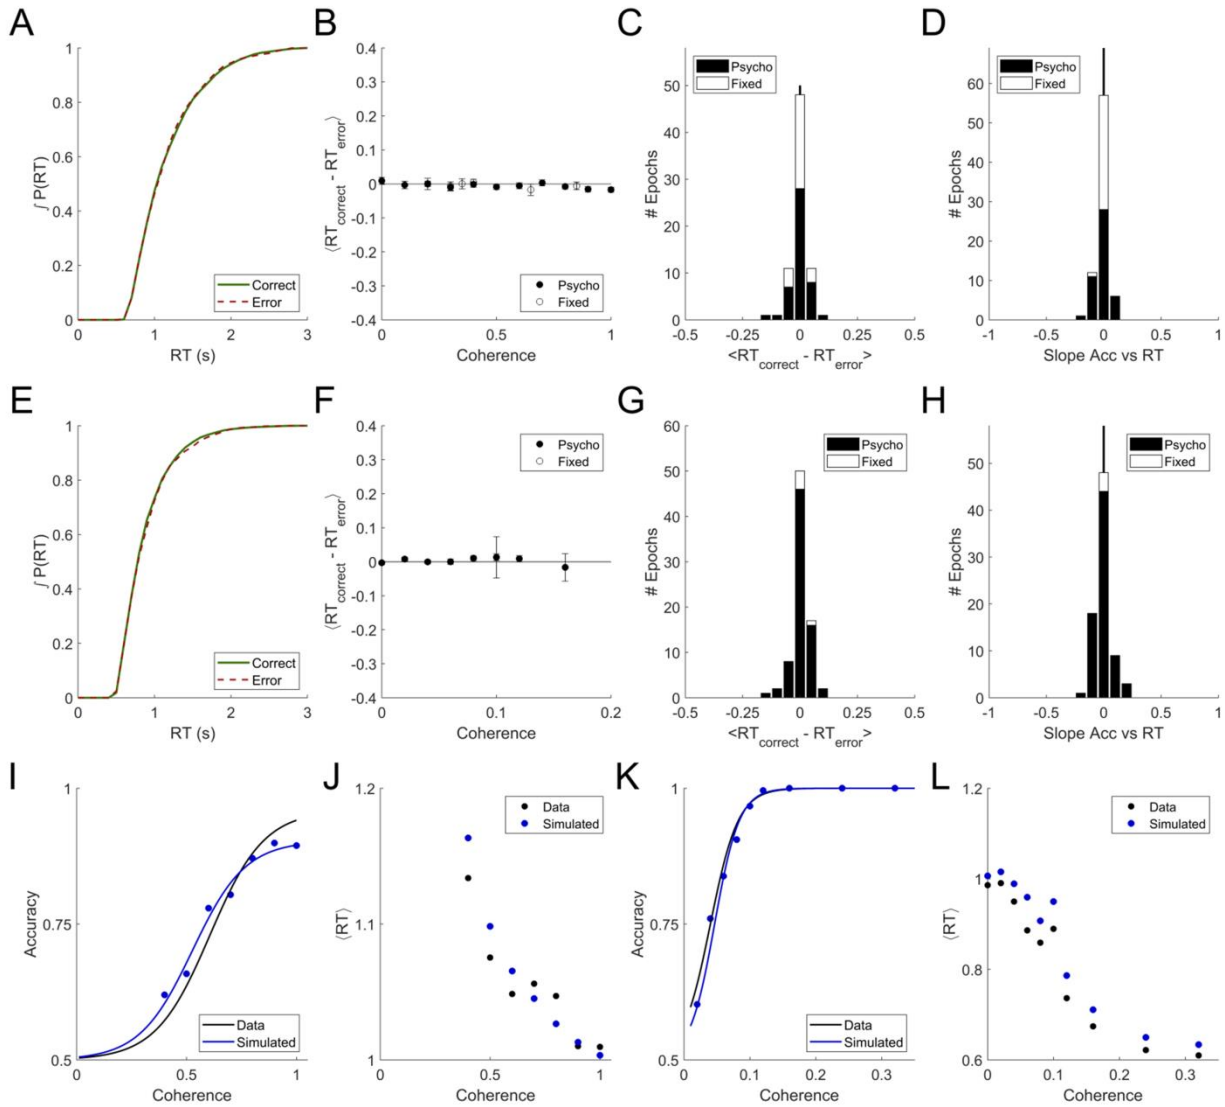

Fig 4. Standard DDM fails to replicate the interaction between elapsed time and accuracy in either species. (A) Cumulative probability distributions of error vs. correct trials in simulated data for the same rat epoch analyzed in Fig 2A. (B) Average difference in RT between nearby correct and error trials as a function of motion strength in simulated rat epochs, cf. Fig 2B. (C) Average RT difference in each simulated rat epoch, cf. Fig 2C. (D) Slope of accuracy vs. RT for simulated trials, cf. Fig 2E. (E-H) Analysis of simulated trials from models fit to human data, cf. Fig 2E-H. (I) Accuracy vs. coherence and psychometric curve fit (blue) for the simulated trials from the model fit to the example rat epoch analyzed in Fig 2I. Psychometric curve from the data (black) is shown for comparison. (J) Chronometric function from the simulated trials (blue), compared with the data (black, c.f. Fig 2J). (K-L) Like I-J for the example human epoch of Fig 2K-L. See Statistical Methods and S5 Appendix for additional details.

Drift Diffusion Model with variable parameters

It was previously shown that adding noise to the parameters of a bounded drift diffusion model can differentially affect the error and correct response time distributions [14, 16]. Depending on parameter values, this can produce either an increase or a decrease in accuracy as a function of response time. Therefore this model could account for both the human and rat results and provide a basis for understanding the difference between them. The stochastic parameter model we

implemented has three additional parameters: variability in starting point σ_z, variability in non-decision-time σ_t, and variability in drift rate σ_v (Fig 5).

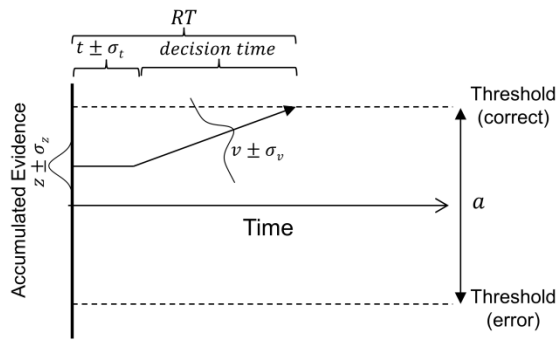

Fig 5. Addition of variability to the parameters of the drift diffusion model. Parameters *t*, *z*, and *v* in the drift diffusion model remain constant within a trial, but vary from trial to trial according to the variability parameters σ_t, σ_z, and σ_v respectively. After (Ratcliff and McKoon, 2008).

In a preliminary analysis (S4 Appendix), we systematically explored the parameter space of the model illustrated in Fig 5. We assessed the simulated models for the effect of interest by comparing the mean response time of correct vs. error trials. On the basis of these exploratory simulations we hypothesized that the result typicaly found in rats ($RT_{correct} > RT_{error}$) could be explained by a high σ_z/a ratio, and/or low σ_v, whereas the result usually found in humans ($RT_{correct} < RT_{error}$) could be explained by a low σ_z/a ratio, and/or high σ_v. We also demonstrated that a starting point bias towards one side (left or right) is sufficient for the standard DDM to produce the sign of the effect in rats ($RT_{correct} > RT_{error}$). We don't think bias fully explains the effect in rats, however, because the sign of the effect is the same in unbiased rats. Therefore we expected that the model term σ_z might capture effects of bias as well as other sources of starting point variability. In other simulations (not shown) the parameter σ_t had no effect on whether accuracy increased or decreased with response time. Therefore we did not expect σ_t to explain the dependence of accuracy on elapsed time in either species, and had no other basis for predicting its value.

When we fit the model shown in Fig 5 to the data set described in Fig 2, we found that adding the three parameters σ_t, σ_z, and σ_v improved the model fits of both species, even after penalizing for the larger number of free parameters (deviance information criterion, DIC, see Statistical Methods). The enhanced model also fit better as assessed by the residual error of RT distributions on cross-validation (discussed below). The population distributions of the epoch-level parameter estimates are shown in Fig 6.

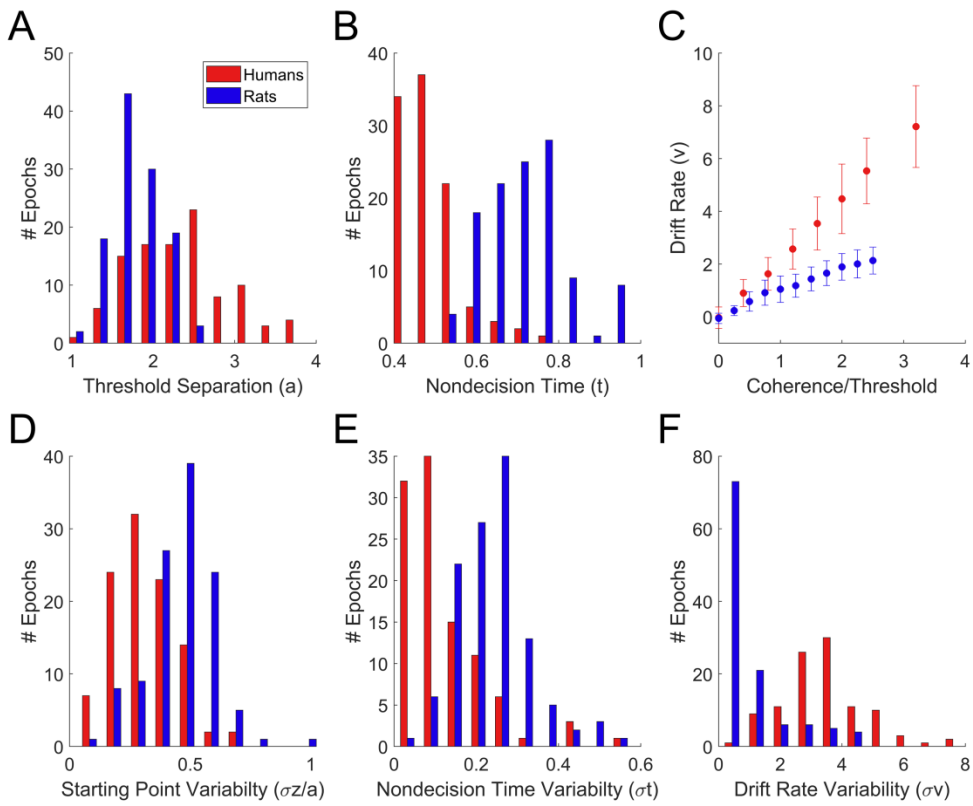

Fig 6. Estimated parameters of Drift Diffusion Model with noise terms fit to data from humans or rats. Data sets as in Fig 3. Parameters estimated for individual human (red) and rat (blue) experimental epochs. (A) Distributions of threshold separation a: a_{human} = 2.33±0.60, a_{rat} = 1.79±0.32. (B) Distributions of nondecision time t: t_{human} = 0.48±0.07 s, t_{rat} = 0.72±0.10 s. (C) Drift rate v for each stimulus strength (mean±SD over subjects) where strength is given by coherence divided by species median coherence threshold, 0.05 (humans) or 0.40 (rats). (D) Distributions of starting point variability parameter σ_z, expressed as a fraction of the threshold separation a: Human σ_z/a = 0.31±0.12, rat σ_z/a = 0.46±0.14. The raw parameter σ_z (not shown) in the human data set was σ_z = 0.68±0.18, in the rat data set σ_z = 0.81±0.22. (E) Distributions of nondecision time variability σ_t: human σ_t = 0.12±0.10 s, rat σ_t = 0.24±0.09 s. (F) Distributions of drift rate variability parameter σ_v: human σ_v = 3.42±1.37, rat σ_v = 0.93±1.13. See Statistical Methods, S1 Appendix and S2 Data for additional details.

These parameter fits were consistent with our preliminary simulations (S4 Appendix) showing that a high σ_z/a ratio with low σ_v can produce the early-error pattern seen in rats, whereas a low σ_z/a ratio with high σ_v can produce the late-error pattern seen in humans. The model was not explicitly optimized for its ability to account for rats' increase in accuracy with elapsed time or humans' decline in accuracy with elapsed time, however. To test whether it did so, we simulated trials for each epoch based on their individual parameters in the variable parameter model (Fig 6), and performed the identical analysis as for real data (Fig 2) and the fixed-parameter model (Fig 4). The model with noise parameters succeeded in capturing the sign of the interaction between elapsed time and accuracy in both species (Fig 7).

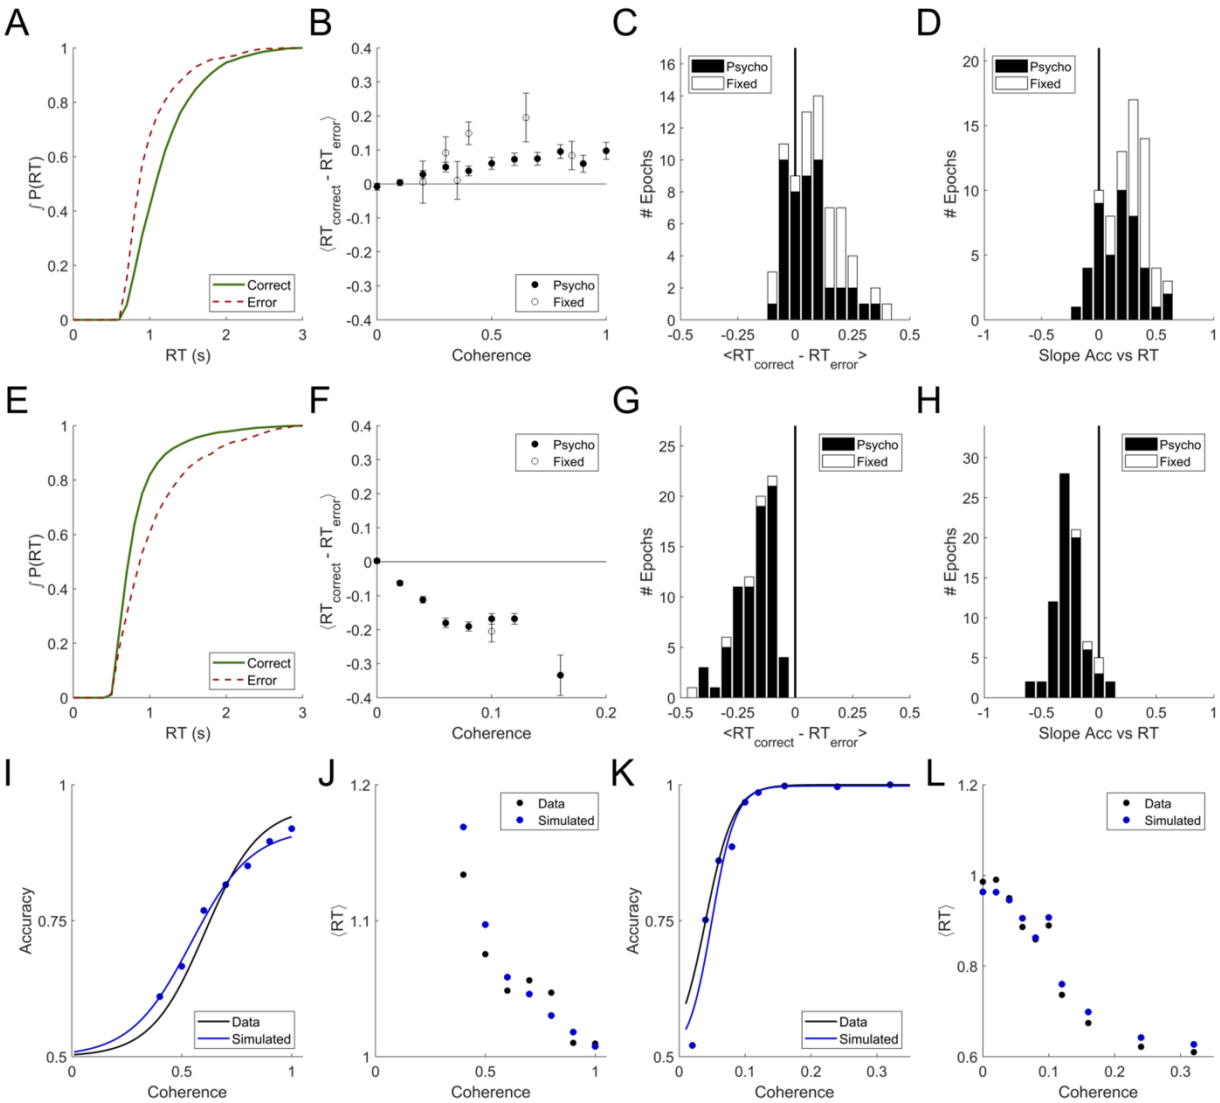

Fig 7. Simulated data from DDM with noise parameters. (A-D) simulations based on parameters fit to data from rats, **(**E-H) simulations based on parameters fit to data from humans. (A,E) Cumulative response time distributions of simulated error trials (red, dashed) and correct trials (green, solid) for the same example subjects shown in Fig 2A,E and 4A,E. (B,F) Average response time difference between correct and error trials, computed on simulated trials within epoch, and averaged across epochs as described in Fig 2B. (C,G) Distribution across epochs of mean response time difference within epoch, as described in Fig 2C,G. (D,H) Distribution over epochs of the slope of accuracy vs. response time, as described in Fig 2D,H. (I,J) Psychometric curve and chronometric function from the simulated trials for the rat epoch analyzed in Fig 2I,J and 4I,J (blue) compared to the data (black). (K,L) Psychometric curve and chronometric function from the simulated trials for the human epoch analyzed in Fig 2K,L and 4K,L (blue) compared to the data (black). See Statistical Methods and S5 Appendix for additional details.

Lesions of the model

To evaluate the contributions of the different noise parameters, we performed lesions within the model. The overall model fit was measured by the residual error between the real and predicted response time distributions using cross-validation data. This was computed separately for each coherence and summed over coherences. Here we can see that for individual subjects (Fig 8A,E) and

over the population (Fig 8B,F) the model with all three noise parameters ('+++', cf. Fig 5) fit the data better than the model without any ('- - -', cf. Fig 1F). We note that because cross-validation data were used this improvement is not attributable to over-fitting.

In single-parameter "lesions" of the model, we re-simulated trials using the parameters fit to the full model (Fig 6) after setting one of the noise parameters (σ_t, σ_z, or σ_v) to zero for all subjects. For rats, the average fit of the model was much worse if we ablated σ_z compared to the ablation of either σ_t or σ_v (Fig 8A,B). For humans, the error was much worse if we ablated σ_v compared to either σ_z or σ_t (Fig 8E,F). Double-lesion results were consistent with this: for rats the double-lesion model that retained only σ_z had the best fit; for humans the double-lesion keeping only σ_v had the best fit. But both σ_z and σ_v are involved in both species, as the double lesion led to greater error than either lesion alone (Fig 8B,F).

To assess the interaction between elapsed time and accuracy, we computed the difference between the mean response times of correct and error trials of the same coherence, $\langle RT_{corrrect} - RT_{error} \rangle$. The model with all three noise parameters (+++) replicated the fact that for rats, errors had shorter response times (Fig 8C-D). Ablating σ_v from the model resulted in this effect being overestimated; ablating σ_z abolished this property and uncovered a weak opposite effect; ablating σ_t was inconsequential. We infer that starting point variability is the model parameter that specifically accounts for rats' increasing accuracy with response time, but this coexists with and is partly offset by a weaker effect in the opposite direction due to drift rate variability.

This model also replicated the trend that human errors have longer response times, but overshot the magnitude of the real effect (Fig 8G-H). Ablating σ_z caused the trend to be further overestimated; ablating σ_v caused the model to predict the opposite effect instead; ablating σ_t was inconsequential. We infer that drift rate variability is the model parameter responsible for reproducing humans' declining accuracy with response time, but this coexists with a weaker effect in the opposite direction due to starting point variability.

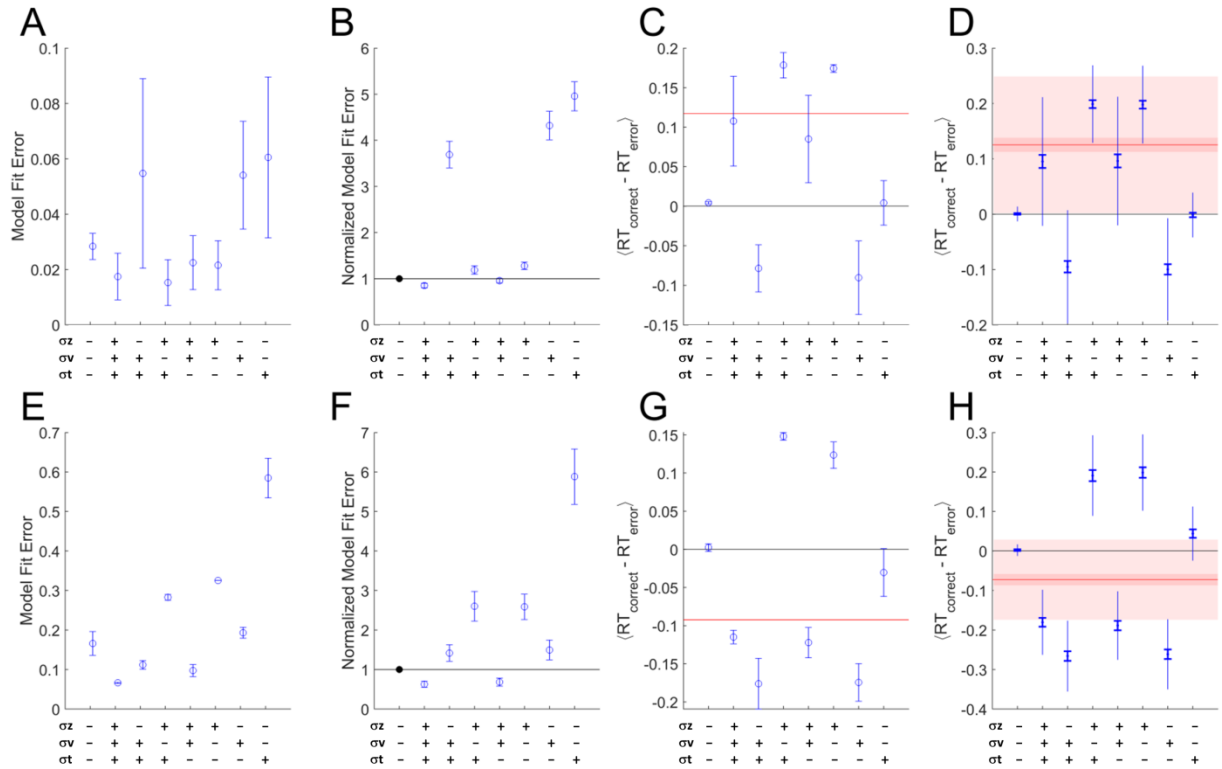

Fig 8. Effects of noise parameter ablations in the model. Results are shown for rats (top row) and humans (bottom row). The leftmost column in each plot (---) shows results for the model that was fit without noise parameters (c.f. Fig. 1,3,4). The second column (+++) represents the model that was fit with three noise parameters σ_t, σ_z, and σ_v (c.f. Fig. 5-7). The remaining columns are results of simulations using parameters fit in the latter model, after selective ablation of noise parameters as indicated below the panels. (A,E) The cross-validation error between simulated and real data for one example subject of each species. Error bars indicate SEM over the two data splits, averaged within split over two independent model fitting runs. (B,F) The cross-validation error (blue) normalized within epoch to the error of the fit of the first model (black) for N=98 rat epochs and N=63 human epochs. Error bars represent SEM over population. (C,G) The difference between the mean response time of correct trials and of errors, for the same example subject shown in (A,E). Difference was calculated within coherence, for all coherences for which subject's accuracy was between 70-90% correct, and then averaged over coherences. Error bars indicate SEM over the two data splits, averaged over two independent runs. Red line indicates the result from the data. (D,H) Difference between correct and error response times computed as in (C,G) within subject, summarized over population for N=97 rat epochs and N=54 human epochs. Red line, dark and light shading represent the mean, SEM and SD over the population for the original data. Blue points indicate population mean from simulations, horizontal hatches show population SEM; whiskers show population SD. All epochs in the data set are included. See Statistical Methods for additional details.

Discussion

The impetus for this study was an observed difference between human and rat decision-making: for humans correct decisions are on average faster than errors, whereas for rats correct decisions are on average slower than errors. Both observations violate the predictions of the standard drift diffusion model. It must be stressed that the sensory task was matched such that rats were just as accurate and just as fast as humans in the task; we are not seeking to account for any difference in performance. Nor is there any evidence that rats' behavior was any more variable or unreliable than that of humans. We do not assume that the difference in response time of correct vs. error trials is functionally significant for either species; the difference is small and accounts for a small fraction of the variance in response time.

The main reason we are interested in the effect (in either species) is because it is revealing about underlying decision-making algorithms. And the main reason we are interested in the difference between species is that a unified theoretical model for mammalian decision-making should be able to account for both. This study showed that a drift diffusion model with noise parameters [14, 16] can parsimoniously account for the interaction between accuracy and response time in both species. This provides a unified way to understand the behavior in both species and a common framework to compare them.

Decreasing accuracy with response time, as found in humans in this task, was reproduced by a dominance of drift rate variability in the model, without explicitly positing an urgency signal or collapsing bound. Similar results have been reported by others [52]. This behavioral pattern in humans and other primates has been widely reported [9, 10] and can be alternatively explained by several other competing models [14, 22, 30, 34, 35, 38-40].

Increasing accuracy with response time, as found in rats in this task, was reproduced by a dominance of starting point variability in the model. The fact that rats exhibit this behavioral pattern was only recently appreciated [41-44], and models developed to explain primate data have yet to be tested on rodent data. One plausible alternative explanation would be reliance on an internal deadline. Although we have not simulated this, it stands to reason that if decisions were triggered by a (noisy) deadline, later responses would be based on more accumulated evidence and be more accurate. That model would not easily account for why response time depends on stimulus strength, but a hybrid model might reconcile these [52]. Alternatively one might be able to attribute the effects to a trade-off between habitual vs. goal-directed strategies [53].

Non-decision-time variability was not crucial to explaining the effect of interest in either species. The data set analyzed here is not suited to test whether non-decision-time variability differs between species, because the response modalities differed and the average non-decision-time was different. In the future, it would be instructive to compare non-decision-time variability in a task for which absolute response time has been matched, and in which confidence is explicitly measured [38, 54].

In summary, this model implicates an important role for behavioral variability in explaining the effects in each species as well as the difference between species. Importantly, we did not find that one species' behavior is more variable than the other, but rather that different types of variability are found in different proportions. Alternative models have not been ruled out, but an appealing aspect of this one is the parsimonious unification of human and rat data.

According to this model, both species had some variability in both starting point and drift rate, contributing opposing effects in different proportions that balanced to produce the observed net effect in each species. The parameter distributions in rat and human populations overlapped considerably. Thus the qualitative difference between humans and rats illustrated in Fig 2 may belie an underlying similarity in their perceptual decision-making mechanisms.

Can context account for variability?

Although variability in starting point and drift rate were modeled as random noise, we do not assume variability to be random. Anything that caused starting point or drift rate to be different in different trials of the same motion strength could produce the same effects. An important future

direction will be to test for and incorporate such effects into the model.For example, if the starting point of evidence accumulation were systematically biased toward the left decision threshold, this would result in a starting point *z* that is closer to the correct decision threshold in left-motion trials, but further from it in the right-motion trials. This would manifest as starting point variability in this model (Fig S4.2). Therefore response bias could account for some of the starting point bias in both species. Response bias is more prevalent in rats than in humans [42]. We don't think this is sufficient to explain the greater starting point variability in rats, however, because correct trials have longer RT than errors even in rat epochs with no bias [42]. Moreoever our modeling results were similar when we excluded all epochs with lapse or bias (S5 Appendix).

It is known that choices in this type of task can be biased by the previous trial's stimulus, response, and outcome in mice [55, 56], rats [57], non-human primates [58], and humans [59], reviewed in [60]. If recent trial history biases the starting point for evidence accumulation, that would also be referred to the starting point variability parameter in the model. Similarly, assymetry in the drift rates of leftward vs. rightward motion could bias choices in either a static or history-dependent manner, and this would be referred to the drift rate variability parameter in the present model.

Conceptually this could be explored by fitting a separate starting point parameter $z_{j,k}$ for each target response side *j* and each recent trial history pattern *k*. Likewise, to quantify the extent to which drift rate variability can be explained by bias and history, one could fit a separate drift rate parameter $v_{i,j,k}$ for each motion strength *i*, response side *j*, and recent trial history category *k*.

Fluctuations in arousal, motivation, satiety or fatigue could conceivably modulate decision thresholds or drift rates from trial to trial independent of either response side or history, without biasing the choice to left or right. We note that fluctuations in the decision threshold parameter *a* would be referred to the starting point variability parameter in the current model [14, 16]. Therefore one might try to account for some of the starting point and/or drift rate variability by fitting separate starting point or drift rate parameters depending on other measurable factors such as eagerness (latency to trial initiation or the number or frequency of request licks), arousal (assessed by pupilometry), fatigue (the number of trials recently completed), satiety (amount of reward recently consumed), or frustration/success (fraction of recent trials penalized/rewarded).

We think a different approach to model specification and parameter fitting will be necessary to pursue these ideas, however. Due to the combinatorial explosion, fitting a different parameter for each condition combitation would result in far too many free parameters to fit, given the limited number of trials. A more promising approach would be to model starting point and drift rate parameters as analytic functions of these factors, referring any residual unexplained variability to a noise term.

Is variability a feature or a bug?

If any of the variability is due to non-random influences, as suggested above, a separate question would be whether this variability is adaptive or dysfunctional in either species. It is possible that non-sensory influences shift the decision-making computation from trial to trial in a systematic and reproducible fashion that would be functionally adaptive in the context of natural behavior, even though we have artificially broken natural spatial and temporal correlations to render it maladaptive in our laboratory task.

For example, in nature some locations may be intrinsically more reward-rich, or very recent reward yields may be informative about the expected rewards at each location. In the real world, recently experienced visual motion might be highly predictive of the direction of subsequent motion stimuli. Therefore biasing either starting point or drift rate according to location or recent stimulus or reward history may be adaptive strategies under ecological constraints, for either or both species.

The general idea that systematic biases reflect adaptive strategies could be rigorously tested by comparing any observed systematic biases to the pertinent environmental statistics. Some biases and environmental statistics could be common across mammals, others could be species-specific.

Conclusions

It has been argued that neural computations underlying sensory decisions could integrate comparative information about incoming sensory stimuli (e.g., left vs. right motion signals), internal representations of prior probability (frequency of left vs. right motion trials) and the expected values (rewards or costs) associated with correct vs. incorrect decisions, in a common currency [5, 6]. These computations are likely to be distributed throughout the brain [61-63], and mixed within brain areas and even within indivudal neurons [64, 65]. Neural mechanisms will need to be described in both rodents and primates before it is possible to engage a discussion of whether and how these mechanisms are similar or different across species.

On the premise that basic mechanisms of perceptual decision-making are likely to be conserved, fitting a single model to data from multiple species – especially where they differ – may be a powerful way to develop and distinguish among alternative computational models. If the differences between species are merely quantitative differences in a conserved underlying algorithm, it may be advantageous to study a given aspect of the computation in a species that has particularly high or low variability in a certain parameter. We have explored one model that can explain our data from both species, but other models may also be able to account for the same data in other ways, potentially making distinct testable predictions for future experiments. To enable such direct comparisons, the entire dataset analyzed in this paper is available for fitting to other models [51].

Methods

Experimental data

No new data were collected for this study. All human and rat data modeled in this manuscript were previously published in [42], where complete experimental methods as well as ethics statements regarding the use of human and animal subjects may be found. The ethics statements therein confirm that all animal and human experiments were performed in strict compliance with all local, federal and international ethical guidelines and with the approval and oversight of the University of California San Diego Institutional Animal Care and Use Committee (IACUC) and Institutional Review Board (IRB) respectively. We further confirm that all human data were de-identified for public distribution as specifically approved by the governing IRB protocol.

To summarize the experiment briefly, the task was random dot coherent motion discrimination. When subjects initiated a trial, white dots appeared at random locations on a black screen (CRT or LCD monitor) and commenced to move. Some fraction of the dots (“signal”) moved at the same speed

toward the rewarded response side. The others ("noise") moved at random velocities. Stimulus strength was modulated by the coherence (fraction of dots that were signal dots). The remaining stimulus parameters (e.g. dot size, dot density, motion speed, contrast) were chosen for each species to ensure that accuracy ranged from chance (50%) to perfect (100%) and response times ranged from ~500-2500ms for typical subjects of the species. Stimulus generation, reward delivery, and behavioral event recording were all achieved using the custom Matlab software suite Ratrix [66] and Psychtoolbox [67-69].

Healthy adult human volunteers were recruited to participate for monetary compensation. Rats were either water restricted prior to the test session, or worked in a closed-economy system with 24 hour/day access to the task as their sole source of water. Subjects voluntarily initiated each trial by a central keypress or lick respectively. This immediately triggered a random dot motion display on a computer monitor, with coherent motion either to the left or right, persisting until the subject responded. Subjects could respond at any time by a left vs. right keypress or lick. Responses on the side toward which the coherent dots were moving were rewarded with money or water. Responses on the opposite side were penalized by a brief time-out. Learning trials and initial trials of each session were excluded from analysis. Remaining trial data from each subject were broken into "epochs" during which the task was not changed and the subject's behavior (accuracy, response time) was not trending. Data selection and exclusion was described in detail by [42]. No additional selection was applied for this study.

The data set included 115 epochs from 18 individual rats, for 274,170 total rat trials, and 104 epochs from 64 individual human subjects, for 151,828 total human trials, the files designated as "allEpochs" in the data repository [51]. This data set includes all approximately stationary epochs from the experimental study, including those with lapse (less than 100% performance at the highest coherence) and/or bias (preference for left or right responses) [42]. We did not exclude the visually impaired human subject (448) during model fitting, but this subject's parameters are not shown (Fig 3 and 6), nor their residual errors computed (Fig 8).

We note that response time distributions are very broad compared to the influence of either coherence or outcome on RT. For example, in the rat epoch shown in Fig 2H-I, 3.4% of the variance in RT was attributable to coherence (P=1.6e-61), 2.1% attributable to outcome (error vs. correct) (P=1.5e-44), <0.5% attributable to an interaction between these terms (P=2.2e-05), and 94.0% attributable to noise within condition (two-factor ANOVA with replicates of 400 trials per condition). In the example human epoch shown in Fig 2J-K, these proportions were 2.6% attributable to coherence (P=4.7e-29), 0.5% to outcome (P=1.3e-07), 0.25% to interaction (P=02.3e-2), and 96.7% to within-condition noise.

Computational methods

Model fitting

Model parameters were fit with the Hierarchical Drift Diffusion Model (HDDM) package in Python [70]. The objective function for fitting is to estimate the joint posterior distribution of all parameters given the observed data, using Markov chain Monte Carlo (MCMC) sampling. The observed data are the trial-by-trial response times, labeled by coherence and distinguishing error trials from correct trials. Individual-level parameters were fit for each epoch, because distinct epochs from the

same subject differed in task parameters and/or average performance metrics. Although species-level hierarchical parameters were also fit as part of the fitting algorithm, those parameters were not interpreted or used for simulations.

The models were trained for a pre-determined 12,000 iterations (standard DDM, Fig 1,3,4) or 102,000 iterations (variable DDM, Fig 5-7). The adequacy of this training duration was assessed by examining the asymptotic behavior over training iterations of each parameter value, to confirm by visual inspection that the model parameters were no longer trending [70]. The standard DDM converged well by any criterion: parameters *a* and *t* had narrow 95% highest density intervals (HDI), Gelman-Rubin $\hat{R}$ statistics <1.01 [71, 72], low Monte Carlo Error (MC_{err}), and high run-to-run reproducibility by coefficient of variation (CV). Some parameters were less confidently estimated – notably the lower-coherence drift rates *v* and the widths of the group posteriors in both models, and the noise terms σ_t, σ_z, and σ_v in the second model [73]. These parameter values sometimes persisted in fluctuating indefinitely with training, and in this sense did not converge well as judged by the HDI or $\hat{R}$ statistic. Thus we cannot exclude that some other parameter combinations could fit the data from either species even better than the ones we report. This does not undermine our findings, however. The parameters we report were well-estimated within run (MC_{err}), reproducible across repeated fitting runs (CV), and sufficient to reproduce the phenomena of interest as well as we demonstrate (Fig 7-8). Further details on parameter estimation are provided in S1 Appendix and S2 Data.

The total number of model parameters was at most 1485 (human) or 1217 (rats), including the species-level posteriors, epoch-specific values for each parameter, including a drift rate ***v*** for every coherence and the three noise parameters (Figs 5-7). The number of data points greatly exceeded the number of parameters: 102 trials/parameter (humans) or 225 trials/parameter (rats). The rat data and human data were fit separately in parallel for speed. This also allowed the species priors to have different widths; otherwise in the HDDM software the widths of group priors are constrained to be the same across groups.

Anecdotally we found that after fitting, smoothing the curve relating coherence to drift rate *v* could improve model simulation performance in cross-validation. We conclude that imposing a smoothness constraint or a parametric form for this relationship within the model fit would be a helpful future improvement, but we did not implement this. All simulations presented in the paper were performed without such smoothing.

The drift diffusion process was simulated according to the equation $X(t) = X(t\text{-}1) \pm \delta$ with probability *p* of increasing and (1-*p*) of decreasing [74]. Here *t* is the time point of the process, with time step τ in seconds; $\delta = \sigma \cdot \sqrt{\tau}$ denotes the step size, where σ is the standard deviation of the Gaussian white noise of the diffusion; $p = 0.5 \cdot \left(1 + v \cdot \frac{\sqrt{\tau}}{\sigma}\right)$, where *v* is the drift rate. To be consistent with the simulation done by the HDDM package [70], the values for τ and σ were set at 0.1 msec and 1, respectively.

For any trial, the process starts at a starting position *z*, sampled from a uniform distribution of range σ_z, assumes a constant drift rate *v*, sampled from a normal distribution of standard deviation σ_v, and continues until $X(t)$ exceeds either threshold boundary. The non-decision-time *t*, sampled from a uniform distribution of range σ_t, is added to the elapsed time to obtain the final RT associated with that trial.

To fit the model parameters to experimental data, i.e. the trial-level response time and accuracy, the HDDM package was used. The input data was first formatted as comma-separated files with the following mandatory columns: 'subj_idx' (epoch number), 'rt' (response time), and 'response' (0 for error and 1 for correct response). The main command was either 'HDDM', if no variability parameters were included, or 'HDDMTruncated', when fitting epoch-level estimates of the variability parameters σ_z, σ_v and σ_t. The species-level parameters were initialized to informative priors [75]; the same priors were used for both human and rat data, as no rat-specific priors were available in the literature.

Aside from the default parameters *a*, *v*, and *t*, additional parameters can be fitted with the 'include' keyword argument to the main command. Otherwise, *z* is defaulted to be 0.5 (i.e. half way between the two thresholds), and the variability parameters are defaulted to zero. Additionally, any parameter can be fitted according to some specific condition(s) using the 'depends_on' keyword argument, as long as the input data file contains information about those conditions. Both models fit parameters at the epoch-level for *a*, *t*, and *v* for each coherence. In addition, both models fit species-level mean parameters for *a*, *t*, and *v* for each coherence, and species-level posterior distribution widths a_{std}, t_{std} and v_{std} (one value of v_{std} for all coherences). For the model with variability, shown in Figs 5-7, in addition to the parameters listed above, we also fit trial-by-trial variability parameters σ_t, σ_v, and σ_z at the epoch-level (one σ_v for all coherences within an epoch), and species-level posteriors (means and stds) for each of these.

For complete technical specification of models, the Python scripts used to train the models in the HDDM package are provided in S3 Code.

Cross-validation error estimation

We assessed the model fits by comparing the sum of squared errors between the response time distributions of the data vs. of trials simulated by the model (Fig 8). For this measurement we split the data into odd trials and even trials, fit models separately to each half of the data, and then compared the RT distribution of each half of the data to an equal number of trials simulated from the model fit to the other half of the data.

We re-coded the response times of error trials as negative values, in order to create a single probability distribution for each stimulus coherence in each epoch which distinguishes both the relative numbers and the RT distributions of error and correct trials. Only coherences with at least 80 trials both the training and cross-validation data sets were included in the error estimate. Some epochs were not estimated because no coherences met this criterion, i.e. the epoch did not have enough trials to split the data in half for this test.

The squared error between simulated and real RT distributions was summed over RT bins and over qualifying coherences to obtain a sum squared error for each epoch. This measure is sensitive to the bin size for discretizing RT, the number of trials of each coherence, and the number of coherences included, so comparisons between epochs or between species are not valid. But all those parameters are fixed within epoch across models, allowing models to be compared.

This measure of error is sensitive to mismatch in the psychometric curve (accuracy as a function of coherence, e.g., Fig 2,4,7 panels I,K), the chronometric function (mean RT as a function of coherence,

e.g. Fig 2,4,7 panels J,L), or the dependence of the RT distributions on the outcome within condition (error vs. correct trials, e.g. Fig 2,4,7 panels A-H). Because the sum squared error is computed on cross-validation data, models with more parameters do not have any advantage in fitting the data. Therefore no correction for model complexity is required.

In Fig 8A,E we show the mean and SEM of the sum squared error over data splits and fitting runs for one example epoch, illustrating the raw numbers that went into the population quantification in Fig 8B,F. To graphically show the population result we normalized within epoch to the mean sum squared error of the standard DDM, and report the mean and SEM over epochs of this ratio for each model.

Local response time difference analysis

In Figs 2, 4, 7, panels B,C,F, and G and Fig 8 C,D,G and H, we measure the interaction between accuracy and response time using the temporally local measure $\langle RT_{correct} - RT_{error} \rangle$ introduced in [42]. Instead of comparing the response time distributions of errors and correct trials of the same coherence estimated from the entire epoch, we compare the response time of each error trial to a temporally adjacent correct trial of the same coherence (requiring a minimum distance of 3 trials to avoid sequential effects). This method is preferred because it is robust to non-trending non-stationarities that are commonly present in both human and rat data, not detected by traditional stationarity tests, and that could confound estimation of the effect of interest. To compute an RT difference for a given coherence within an epoch, we required a minimum of 50 error trials. Therefore coherences with close to 100% accuracy had too few error trials for inclusion in the analysis. In Fig 8 some epochs were not estimated because no coherences had enough error trials for the analysis.

Statistical methods

We are interested in whether and how models can account for the phenomena observed in these example data, regardless of establishing the frequency with which the phenomena occur in the respective populations. The original data (Fig 2) came from an exploratory study that intermingled sample accumulation and analysis [42]. Moreover, individual subjects contributed multiple epochs to the data set, so the samples in the population are not independent. Therefore that study did not make claims about the statistical significance of the difference between the populations. Nevertheless, there was no intentional bias in the sampling of the populations, sample accrual was terminated for pragmatic reasons unrelated to the findings, the human data set recapitulates findings of many past human studies, and the difference between species in this data set held up in almost all proposed re-groupings of the data, including ones that included only one epoch from each subject [42]. For all these reasons we think it is likely that the data sets are representative of the respective species. It is important to disclose that the sampling procedures required by null hypothesis statistical testing were not strictly applied, but with this caveat in mind, statistical tests of population comparisons are provided below.

Statistics for Figs 2, 4, and 7

Figs 2A,E illustrate, for one example fixed-coherence epoch from each species, the cumulative RT distributions (CDFs) of correct and error trials within one fixed-coherence testing block. For the example rat (Fig 2A) the null hypothesis that the distributions are the same can be rejected with

P=2.91e-171 (N=8851,1913) by two-tailed Kolmogorov–Smirnov (KS) test. For the example human (Fig 2E), the null hypothesis can be rejected with P= 4.69e-05 (N=745,153) by KS test. In contrast to this, the error and correct CDFs are not distinguishable for trials simulated by the standard DDM model with parameters fit to the same epochs (Fig 4A,E), P=8.98e-01 (N=4100,859) for the rat example and P=4.43e-01 (N=4103,891) for the human example. The variable-parameter DDM does produce unequal error- and correct-trial CDFs (Fig 7A,E), P=1.49e-43 (N=4183,789) for the rat example and P=5.50e-34 (N=4053,881) for the human example. The sign of the difference in the simulated data (early errors for the rat, late errors for the human) was the same as in the real data.

The remaining panels of Fig 2 reproduce the previously published quantification of this effect across experimental epochs in the particular data set we modeled [42], and provide a basis for comparison to model simulations analyzed in Fig 4 and 7. It is known that the standard DDM cannot produce different error and correct response time distributions (RTDs). Therefore Fig 4 serves as a negative control for our study, showing the extent of deviation from the null hypothesis (H_0: error RTD = correct RTD) one can expect from analysis of a finite number of stochastically simulated trials by our analysis methods, when the null hypothesis is true. In Fig 4B-C and F-G, we quantify this effect by the measure $\langle RT_{correct} - RT_{error} \rangle$, which has an expected value of zero if the error and correct RTDs are the same. Panels B and F show, for descriptive purposes, the magnitude of the effect as a function of stimulus strength. Panels C and G quantify that effect in each epoch by the average over the above-threshold coherences. Compared to the null hypothesis of a random sample from a population with a median of zero, the distributions in the real data are inconsistent with that hypothesis by two-tailed Wilcoxon signed rank test (α=0.01) for both rats (Fig 2C, P=8.13e-10, N=56) and humans (Fig 2G, P=1.25e-05, N=33). When parameters of a standard DDM were fit to each epoch, the RT difference in the simulated data could not be distinguished from zero in either rats (Fig 4C, P=3.72e-01, N=73) or humans (Fig 4G, P=2.56e-01, N=80).

In panels D and H, the same effect is instead quantified by the slope of the trend between RT and accuracy, which has an expected value of zero if error and correct response time distributions are the same. Compared to the null hypothesis of a random sample from a population with a median of zero, the distributions in the real data are inconsistent with that hypothesis by two-tailed Wilcoxon signed rank test (α=0.01) for both rats (Fig 2D, P=6.49e-11, N=75) and humans (Fig 2H, P=1.14e-10, N=75). In the simulated data from the standard DDM, the slope was not distinguishable from zero for either rats (Fig 4D, P=2.68e-01, N=76) or humans (Fig 4H, P=1.41e-01, N=79). These results support our statement that the standard DDM cannot reproduce the differences between correct and error trial response time distributions that are present in this data set. It also serves to show the extent of noise in our estimates of RT difference or slope statistics from a finite number of simulated trials, when the null hypothesis is true.

For the DDM with variability terms, the median of $\langle RT_{correct} - RT_{error} \rangle$ was different from zero for both rats (Fig 7C, P=1.21e-07, N=72) and humans (Fig 7G, P=7.85e-15, N=80). The slope of accuracy vs. RT was also different from zero in both rats (Fig 7D, P=6.86e-12, N=74) and humans (Fig 7H, P=2.86e-14, N=79). Moreover the sign of the effect in the model simulation was the same as that of the data (positive for rats, negative for humans) for both the RT difference and the slope measure. These results support our claim that there exist parameters (namely, those shown in Fig 6) for which the variable DDM can produce a difference between error and correct trials with the same sign as the examples in

these human and rat data sets. Similar results were obtained in both populations after exclusion of all epochs with lapse or bias (S5 Appendix).

Statistics for Figs 3 and 6

Parameter distributions of humans vs. rats were compared using a two-tailed Wilcoxon Rank Sum test (α=0.01). For the standard DDM (Fig 3), distributions of parameter *a* (Fig 3A) were not different between human and rat data sets, P=5.18e-01 (N=104,115). Distributions of parameter *t* (Fig 3B) were different, P=5.85e-32 (N=104,115). The drift rate parameters cannot be compared, as the coherences were different and the visual stimuli were non-comparable in other respects. For the DDM with variability parameters (Fig 5), distributions of parameter *a* (Fig 5A) were different, P=2.86e-12 (N=104,115). Distributions of parameter *t* (Fig 5B) were different, P=1.89e-33 (N=104,115). Distributions of parameter σ_t (Fig 5E) were different, P=5.17e-19 (N=104,115). Distributions of parameter σ_v (Fig 5F) were different, P=2.85e-26 (N=104,115). Distributions of parameter ratio σ_z/a (Fig 5D) were different, P=7.09e-14 (N=104,115).

Details on the statistics of model fitting, including convergence, high density confidence interval, Monte Carlo error, and the reproducibility of parameter fits across fitting runs, may be found in S1 Appendix. The complete set of parameters for every epoch from a representative fitting run, including the STD and quantiles over training iterations, may be found in S2 Data.

Model Performance Comparison

For model comparison we used the deviance information criterion (DIC) as provided by HDDM, normalized by the number of trials in the respective data set. The DIC is a hierarchical modeling generalization of the Akaike Information Criterion (AIC) used for model comparison in this context [76]. The variable DDM (Fig 5) consistently had a lower DIC than the standard DDM (Fig 1F) for both humans (P=7.5e-09) and rats (P=2.3e-07). For humans, the standard DDM DIC=0.566+/-0.000 N=2 runs, the variable DDM DIC=0.423+/-0.003 N=5 runs. For rats, the standard DDM DIC=1.450+/-0.000 N=2 runs, the variable DDM DIC=1.197+/-0.005 N=4 runs. On these grounds, the variable DDM would be selected as the better model for both humans and rats. The standard and variable DDM fits were also compared by cross validation error, see next section.

In principle, it should not be possible for the variable DDM to fit the data worse than the standard DDM, because the standard DDM is an available solution within the variable DDM model (by setting all variability parameters to zero). Therefore we compared the un-normalized sum squared error of two models with a one-tailed, paired sample test (H_0: the model fits are equal; vs. H_1: variable DDM is better). For the human data set, the variable DDM (mean error over epochs = 0.035) fit better than the standard DDM (mean error 0.075) in 54 of 63 epochs, P=9.42e-10 by one-tailed Sign Rank test. For the rat data set, the variable DDM (mean error 0.042) fit better than the standard DDM (mean error 0.056) in 55 of 98 epochs, P=8.06e-04.

We obtained similar results when we ran this entire anlaysis on a previously specified data subset that contained only one unbiased, low-lapse psychometric epoch from each subject (the bestUnbiasedPsychometricEpochs data set at [51], not shown.

Acknowledgements

Carly Shevinsky shared the human subject data set in advance of publication.

References

1. Britten KH, Shadlen MN, Newsome WT, Movshon JA. The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci. 1992;12(12):4745-65. Epub 1992/12/01. PubMed PMID: 1464765; PubMed Central PMCID: PMCPMC6575768.
2. Shadlen MN, Britten KH, Newsome WT, Movshon JA. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J Neurosci. 1996;16(4):1486-510. Epub 1996/02/15. PubMed PMID: 8778300; PubMed Central PMCID: PMCPMC6578557.
3. Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci. 1996;13(1):87-100. Epub 1996/01/01. doi: 10.1017/s095252380000715x. PubMed PMID: 8730992.
4. Britten KH, Shadlen MN, Newsome WT, Movshon JA. Responses of neurons in macaque MT to stochastic motion signals. Vis Neurosci. 1993;10(6):1157-69. Epub 1993/11/01. doi: 10.1017/s0952523800010269. PubMed PMID: 8257671.
5. Gold JI, Shadlen MN. Neural computations that underlie decisions about sensory stimuli. Trends Cogn Sci. 2001;5(1):10-6. Epub 2001/02/13. doi: 10.1016/s1364-6613(00)01567-9. PubMed PMID: 11164731.
6. Gold JI, Shadlen MN. The neural basis of decision making. Annu Rev Neurosci. 2007;30:535-74. Epub 2007/06/30. doi: 10.1146/annurev.neuro.29.051605.113038. PubMed PMID: 17600525.
7. Huk AC, Shadlen MN. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci. 2005;25(45):10420-36. Epub 2005/11/11. doi: 10.1523/JNEUROSCI.4684-04.2005. PubMed PMID: 16280581; PubMed Central PMCID: PMCPMC6725829.
8. Mazurek ME, Roitman JD, Ditterich J, Shadlen MN. A role for neural integrators in perceptual decision making. Cereb Cortex. 2003;13(11):1257-69. Epub 2003/10/25. doi: 10.1093/cercor/bhg097. PubMed PMID: 14576217.
9. Palmer J, Huk AC, Shadlen MN. The effect of stimulus strength on the speed and accuracy of a perceptual decision. J Vis. 2005;5(5):376-404. Epub 2005/08/16. doi: 10.1167/5.5.1. PubMed PMID: 16097871.
10. Roitman JD, Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci. 2002;22(21):9475-89. Epub 2002/11/06. PubMed PMID: 12417672; PubMed Central PMCID: PMCPMC6758024.
11. Shadlen MN, Newsome WT. Motion perception: seeing and deciding. Proc Natl Acad Sci U S A. 1996;93(2):628-33. Epub 1996/01/23. doi: 10.1073/pnas.93.2.628. PubMed PMID: 8570606; PubMed Central PMCID: PMCPMC40102.
12. Shadlen MN, Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J Neurophysiol. 2001;86(4):1916-36. Epub 2001/10/16. doi: 10.1152/jn.2001.86.4.1916. PubMed PMID: 11600651.
13. Stone M. Models for choice-reaction time. Psychometrika. 1960;25:251-26.

14. Ratcliff R, McKoon G. The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. 2008;20(4):873-922. Epub 2007/12/19. doi: 10.1162/neco.2008.12-06-420. PubMed PMID: 18085991; PubMed Central PMCID: PMCPMC2474742.
15. Ratcliff R, Smith PL, Brown SD, McKoon G. Diffusion Decision Model: Current Issues and History. Trends Cogn Sci. 2016;20(4):260-81. Epub 2016/03/10. doi: 10.1016/j.tics.2016.01.007. PubMed PMID: 26952739; PubMed Central PMCID: PMCPMC4928591.
16. Ratcliff R, Tuerlinckx F. Estimating parameters of the diffusion model: approaches to dealing with contaminant reaction times and parameter variability. Psychon Bull Rev. 2002;9(3):438-81. Epub 2002/11/05. doi: 10.3758/bf03196302. PubMed PMID: 12412886; PubMed Central PMCID: PMCPMC2474747.
17. Ashby FG. A Biased Random-Walk Model for 2 Choice Reaction-Times. J Math Psychol. 1983;27(3):277-97. PubMed PMID: WOS:A1983SN03200004.
18. Busemeyer JR, Townsend JT. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. Psychological review. 1993;100(3):432-59. PubMed PMID: Medline:8356185.
19. Usher M, McClelland JL. The time course of perceptual choice: the leaky, competing accumulator model. Psychological review. 2001;108(3):550-92. PubMed PMID: Medline:11488378.
20. Brown SD, Heathcote A. The simplest complete model of choice response time: linear ballistic accumulation. Cogn Psychol. 2008;57(3):153-78. PubMed PMID: Medline:18243170.
21. Sun P, Landy MS. A Two-Stage Process Model of Sensory Discrimination: An Alternative to Drift-Diffusion. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2016;36(44):11259-74. PubMed PMID: Medline:27807167.
22. Huang Y, Rao RP. Reward optimization in the primate brain: a probabilistic model of decision making under uncertainty. PLoS One. 2013;8(1):e53344. Epub 2013/01/26. doi: 10.1371/journal.pone.0053344. PubMed PMID: 23349707; PubMed Central PMCID: PMCPMC3551910.
23. Purcell BA, Heitz RP, Cohen JY, Schall JD, Logan GD, Palmeri TJ. Neurally constrained modeling of perceptual decision making. Psychological review. 2010;117(4):1113-43. PubMed PMID: Medline:20822291.
24. Scott BB, Constantinople CM, Erlich JC, Tank DW, Brody CD. Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. Elife. 2015;4:e11308. PubMed PMID: Medline:26673896.
25. White CN, Servant M, Logan GD. Testing the validity of conflict drift-diffusion models for use in estimating cognitive processes: A parameter-recovery study. Psychon B Rev. 2018;25(1):286-301. PubMed PMID: Medline:28357629.
26. Usher M, Tsetsos K, Yu EC, Lagnado DA. Dynamics of decision-making: from evidence accumulation to preference and belief. Front Psychol. 2013;4:758. PubMed PMID: Medline:24151480.
27. Heathcote A, Love J. Linear deterministic accumulator models of simple choice. Front Psychol. 2012;3:292. PubMed PMID: Medline:22936920.
28. Tsetsos K, Gao J, McClelland JL, Usher M. Using Time-Varying Evidence to Test Models of Decision Dynamics: Bounded Diffusion vs. the Leaky Competing Accumulator Model. Front Neurosci. 2012;6:79. PubMed PMID: Medline:22701399.

29. Wang X-J. Probabilistic decision making by slow reverberation in cortical circuits. Neuron. 2002;36(5):955-68. PubMed PMID: Medline:12467598.
30. Rao RPN. Decision making under uncertainty: a neural model based on partially observable markov decision processes. Front Comput Neurosci. 2010;4:146. PubMed PMID: Medline:21152255.
31. Pleskac TJ, Busemeyer JR. Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. Psychological review. 2010;117(3):864-901. PubMed PMID: Medline:20658856.
32. Bowman NE, Kording KP, Gottfried JA. Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. Neuron. 2012;75(5):916-27. PubMed PMID: Medline:22958830.
33. Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A. The cost of accumulating evidence in perceptual decision making. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2012;32(11):3612-28. PubMed PMID: Medline:22423085.
34. Tajima S, Drugowitsch J, Pouget A. Optimal policy for value-based decision-making. Nat Commun. 2016;7:12400. PubMed PMID: Medline:27535638.
35. Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2011;31(17):6339-52. PubMed PMID: Medline:21525274.
36. Boehm U, Hawkins GE, Brown S, van Rijn H, Wagenmakers E-J. Of monkeys and men: Impatience in perceptual decision-making. Psychon B Rev. 2016;23(3):738-49. PubMed PMID: Medline:26518307.
37. Hawkins GE, Forstmann BU, Wagenmakers EJ, Ratcliff R, Brown SD. Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. J Neurosci. 2015;35(6):2476-84. Epub 2015/02/13. doi: 10.1523/JNEUROSCI.2410-14.2015. PubMed PMID: 25673842; PubMed Central PMCID: PMCPMC6605613.
38. Ratcliff R, Starns JJ. Modeling confidence judgments, response times, and multiple choices in decision making: recognition memory and motion discrimination. Psychol Rev. 2013;120(3):697-719. Epub 2013/08/07. doi: 10.1037/a0033152. PubMed PMID: 23915088; PubMed Central PMCID: PMCPMC4106127.
39. Ditterich J. Evidence for time-variant decision making. Eur J Neurosci. 2006;24(12):3628-41. PubMed PMID: Medline:17229111.
40. Ditterich J. Stochastic models of decisions about motion direction: behavior and physiology. Neural Netw. 2006;19(8):981-1012. PubMed PMID: Medline:16952441.
41. Reinagel P. Speed and accuracy of visual motion discrimination by rats. PloS one. 2013;8(6):e68505. PubMed PMID: Medline:23840856.
42. Shevinsky CA, Reinagel P. The Interaction Between Elapsed Time and Decision Accuracy Differs Between Humans and Rats. Front Neurosci. 2019;13:1211. PubMed PMID: Medline:31803002.
43. Reinagel P. Speed and accuracy of visual image discrimination by rats. Front Neural Circuits. 2013;7:200. PubMed PMID: Medline:24385954.
44. Sriram B, Li L, Cruz-Martin A, Ghosh A. A Sparse Probabilistic Code Underlies the Limits of Behavioral Discrimination. Cerebral cortex (New York, N Y : 1991). 2020;30(3):1040-55. PubMed PMID: Medline:31403676.
45. McCormack PD, Swenson AL. Recognition memory for common and rare words. J Exp Psychol. 1972;95(1):72-7. PubMed PMID: Medline:5070296.

46. Ratcliff R, Rouder JN. Modeling response times for two-choice decisions. Psychol Sci. 1998;9(5):347-56. doi: Doi 10.1111/1467-9280.00067. PubMed PMID: WOS:000076649300003.
47. Stirman JN, Townsend LB, Smith SL. A touchscreen based global motion perception task for mice. Vision Res. 2016;127:74-83. doi: 10.1016/j.visres.2016.07.006. PubMed PMID: WOS:000384633100009.
48. Long MH, Jiang WQ, Liu DC, Yao HS. Contrast-dependent orientation discrimination in the mouse. Sci Rep-Uk. 2015;5. doi: ARTN 15830
10.1038/srep15830. PubMed PMID: WOS:000363626000001.
49. Heitz RP. The speed-accuracy tradeoff: methodology, and behavior. Front Neurosci. 2014;8. doi: ARTN 150
10.3389/fnins.2014.00150. PubMed PMID: WOS:000346489100001.
50. Hanks TD, Summerfield C. Perceptual Decision Making in Rodents, Monkeys, and Humans. Neuron. 2017;93(1):15-31. doi: 10.1016/j.neuron.2016.12.003. PubMed PMID: WOS:000396427800006.
51. Reinagel P, Shevinsky CA. Human and Rat Motion Discrimination Reaction Time Task Data. Harvard Dataverse2020.
52. Karsilar H, Simen P, Papadakis S, Balci F. Speed accuracy trade-off under response deadlines. Front Neurosci. 2014;8. doi: ARTN 248
10.3389/fnins.2014.00248. PubMed PMID: WOS:000346509800001.
53. Keramati M, Dezfouli A, Piray P. Speed/accuracy trade-off between the habitual and the goal-directed processes. PLoS Comput Biol. 2011;7(5):e1002055. Epub 2011/06/04. doi:
10.1371/journal.pcbi.1002055. PubMed PMID: 21637741; PubMed Central PMCID: PMCPMC3102758.
54. Kepecs A, Uchida N, Zariwala HA, Mainen ZF. Neural correlates, computation and behavioural impact of decision confidence. Nature. 2008;455(7210):227-U55. doi: 10.1038/nature07200. PubMed PMID: WOS:000259090800048.
55. Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Scholvinck ML, et al. The Detection of Visual Contrast in the Behaving Mouse. Journal of Neuroscience. 2011;31(31):11351-61. doi:
10.1523/Jneurosci.6689-10.2011. PubMed PMID: WOS:000293459400025.
56. Hwang EJ, Dahlen JE, Mukundan M, Komiyama T. History-based action selection bias in posterior parietal cortex. Nat Commun. 2017;8(1):1242. Epub 2017/11/02. doi: 10.1038/s41467-017-01356-z. PubMed PMID: 29089500; PubMed Central PMCID: PMCPMC5663966.
57. Lavan D, McDonald JS, Westbrook RF, Arabzadeh E. Behavioural Correlate of Choice Confidence in a Discrete Trial Paradigm. Plos One. 2011;6(10). doi: ARTN e26863
10.1371/journal.pone.0026863. PubMed PMID: WOS:000296521400041.
58. Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. Science. 2004;304(5678):1782-7. doi: DOI 10.1126/science.1094765. PubMed PMID: WOS:000222089500040.
59. Goldfarb S, Wong-Lin K, Schwemmer M, Leonard NE, Holmes P. Can post-error dynamics explain sequential reaction time patterns? Front Psychol. 2012;3. doi: ARTN 213
10.3389/fpsyg.2012.00213. PubMed PMID: WOS:000208863900224.
60. Frund I, Wichmann FA, Macke JH. Quantifying the effect of intertrial dependence on perceptual decisions. J Vision. 2014;14(7). doi: Artn 9
10.1167/14.7.9. PubMed PMID: WOS:000338844900009.

61. Hattori R, Danskin B, Babic Z, Mlynaryk N, Komiyama T. Area-Specificity and Plasticity of History-Dependent Value Coding During Learning. Cell. 2019;177(7):1858-72.e15. PubMed PMID: Medline:31080067.
62. Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. Spontaneous behaviors drive multidimensional, brainwide activity. Science (New York, N Y). 2019;364(6437):255. PubMed PMID: Medline:31000656.
63. Stringer C, Pachitariu M, Steinmetz N, Carandini M, Harris KD. High-dimensional geometry of population responses in visual cortex. Nature. 2019;571(7765):361-5. PubMed PMID: Medline:31243367.
64. Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, et al. The importance of mixed selectivity in complex cognitive tasks. Nature. 2013;497(7451):585-90. PubMed PMID: Medline:23685452.
65. Parthasarathy A, Herikstad R, Bong JH, Medina FS, Libedinsky C, Yen S-C. Mixed selectivity morphs population codes in prefrontal cortex. Nat Neurosci. 2017;20(12):1770-9. PubMed PMID: Medline:29184197.
66. Meier P, Flister E, Reinagel P. Collinear features impair visual detection by rats. J Vis. 2011;11(3). Epub 2011/04/02. doi: 10.1167/11.3.22. PubMed PMID: 21454857.
67. Brainard DH. The Psychophysics Toolbox. Spatial vision. 1997;10(4):433-6. PubMed PMID: Medline:9176952.
68. Kleiner M, Brainard D, Pelli D. What's new in Psychtoolbox-3? Perception. 2007;36:14-. PubMed PMID: WOS:000250594600049.
69. Kleiner M. Visual stimulus timing precision in Psychtoolbox-3: Tests, pitfalls and solutions. Perception. 2010;39:189-. PubMed PMID: WOS:000289740100580.
70. Wiecki TV, Sofer I, Frank MJ. HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. Front Neuroinform. 2013;7. doi: ARTN 14 10.3389/fninf.2013.00014. PubMed PMID: WOS:000209207300013.
71. Brooks S, Gelman A. General methods for monitoring convergence of iterative simulations. J Computational and Graphical Statistics. 1997;7:434-55.
72. Gelman A, Rubin D. Inference from iterative simulation using multiple sequences. Statistical Science. 1992;7:457-511.
73. Boehm U, Annis J, Frank MJ, Hawkins GE, Heathcote A, Kellen D, et al. Estimating across-trial variability parameters of the Diffusion Decision Model: Expert advice and recommendations. J Math Psychol. 2018;87:46-75. PubMed PMID: CCC:000453341900004.
74. Milosavljevic M, Malmaud J, Huth A, Koch C, Rangel A. The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure. Judgm Decis Mak. 2010;5(6):437-49. PubMed PMID: WOS:000283646400003.
75. Matzke D, Wagenmakers EJ. Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. Psychon B Rev. 2009;16(5):798-817. doi: 10.3758/Pbr.16.5.798. PubMed PMID: WOS:000271686100003.
76. Spiegelhalter DJ, Best NG, Carlin BR, van der Linde A. Bayesian measures of model complexity and fit. J Roy Stat Soc B. 2002;64:583-616. doi: Doi 10.1111/1467-9868.00353. PubMed PMID: WOS:000179221100001.

Supporting information

S1 Appendix. Parameter estimation confidence. Details of model fitting convergence assessment and reproducibility of the estimated parameters.

S2 Data. Complete parameter sets. Example final parameter output files including the epoch-level parameters and fitting statistics for both models and both species.

S3 Code. Python code for parameter fitting. Scripts for use with HDDM package to fit either model to either species along with ancillary utility files and setup instructions.

S4 Appendix. Preliminary simulations. Parameter sweeps of the DDM with variability terms, and simulation of the standard DDM with response bias.

S5 Appendix. Excluding epochs with lapse and bias. Alternative versions of Fig 2, 4, and 7 showing that our findings hold true within the subset of epochs that were free of lapse or bias.