In press in *Neuron*, 2024

Sensory choices as logistic classification

Matteo Carandini

UCL Institute of Ophthalmology, University College London, London WC1 6BT, UK

m.carandini@ucl.ac.uk

Logistic classification is a simple way to make choices based on a set of factors: give each factor a weight, sum the results, and use the sum to set the log odds of a random draw. This operation is known to describe human and animal choices based on value (economic decisions). There is increasing evidence that it also describes choices based on sensory inputs (perceptual decisions), presented across sensory modalities (multisensory integration) and combined with non-sensory factors such as prior probability, expected value, overall motivation, and recent actions. Logistic classification can also capture the effects of brain manipulations such as local inactivations. The brain may implement by thresholding stochastic inputs (as in signal detection theory) acquired over time (as in the drift diffusion model). It is the optimal strategy under certain conditions, and the brain appears to use it as a heuristic in a wider set of conditions.

Introduction

Many situations require choosing between options based on multiple factors. In a choice between walking and driving, the factors may include distance, traffic, weather, cost, and parking[1]. In a choice between picking a berry vs. moving on, they may include ripeness, hunger, and ease of reach. How does the brain combine such diverse factors to make a choice?

A large body of research indicates that the brain makes choices by logistic classification: perform a weighted sum of the factors, use the result to set the (log) odds of a coin, and toss that coin (Figure 1**a**). If there are more than two options, the coin is replaced by a die. Logistic classification was invented by statisticians in the secrecy surrounding WWII[2]. In the subsequent decades, it was linked to the theory of choice from mathematical psychology and to value-based decisions in economics[1,3,4]. Since then, it has become a workhorse of statistics and machine learning[5]. It characterizes value-based decisions by humans and animals[6-9] and the corresponding brain activity[7], suggesting that it was discovered first by evolution.

There is increasing evidence that logistic classification describes not only choices based on considerations of value (economic decisions) but also choices based on sensory input[2,10,11] (perceptual decisions) and on additional factors, alone or in combination. These factors include trial history and prior probability[12-22], disparities in value[7-9], signals from additional sensory modalities[23] (multisensory integration), and effects of brain manipulations such as local inactivations[24-26].

Here I briefly review this evidence, and I argue that logistic classification is an effective description of results that had been previously analyzed with a different framework (signal detection theory[27-29]). I show that it is an optimal strategy in some conditions and a useful heuristic in others, perhaps explaining why it is widely implemented by the brains of multiple species.

Results

When a choice involves only two alternatives, logistic classification is particularly simple (Figure 1**a**). First, give each of the factors $x_1, x_2, \ldots$ a weight $W_1, W_2, \ldots$ Second, sum them to obtain a *decision variable*

$$z = W_1 x_1 + W_2 x_2 + \cdots + B \qquad \text{Eq. 1}$$

where B represents a constant *bias*. Third, use that decision variable to bias a coin, i.e., to set the probability p that it will return Heads. Specifically, use it as the *log odds* of drawing Heads:

$$\log\left(\frac{p}{1-p}\right) = z \qquad \text{Eq. 2}$$

Fourth, flip the coin.

This process can be written as a single equation (Figure 1**a**):

$$p = \sigma(W_1 x_1 + W_2 x_2 + \cdots + B) \qquad \text{Eq. 3}$$

where σ is the *logistic sigmoid* (Figure 1**b**)

$$\sigma(z) = \frac{1}{1+\exp(-z)} \qquad \text{Eq. 4}$$

In statistics, these equations are often used for "logistic regression": to find the weights $W_1, W_2, \ldots$ and bias B that best relate the factors $x_1, x_2, \ldots$ to some observed outcomes[3,5]. We are here interested in the reverse: using the weights as given, to classify a new observation of the factors and thus make a choice. This is best called "logistic classification"[5].

In this description of logistic classification, a choice is stochastic (flipping a coin) and the factors x_i are deterministic. We will later see an equivalent description where the factors are stochastic and the choice is deterministic, as in signal detection theory[27-29], which is particularly apt when the factors x_i are sensory (Section S1). We will also see a description where the stochastic factors are acquired over time, as in the drift diffusion model[2,30-32] (Section S2). For most of the paper, however, this simple description will suffice.

Sensory choices between two alternatives

Though it is simple, logistic classification can do an excellent job at summarizing perceptual decisions. For example, in a classic motion discrimination task, a monkey views dots and estimates whether most of them move Left or Right (Figure 1**d**). The task difficulty depends on the coherence x of the dots, defined as positive if most dots move to the Right and negative otherwise. The *psychometric curve*, i.e., the probability of Right choices as a function of stimulus strength can then be fit with a logistic classification model (Eq. 3) with a single factor[10,11]:

$$p = \sigma(Wx + B)$$

With two free parameters (W and B) the model describes how nine values of coherence x determine the probability of choosing Right (Figure 1**e**).

The model and the data are easier to interpret if instead of probability we plot the logarithm of the odds (Figure 1**f**). In this representation (as shown in Eq. 1), the logistic classification model becomes a straight line:

$$z = Wx + B \qquad \text{Eq. 5}$$

and the measured odds fall neatly on the line (Figure 1**f**). In the rest of the paper, I will thus express the model in terms of z.

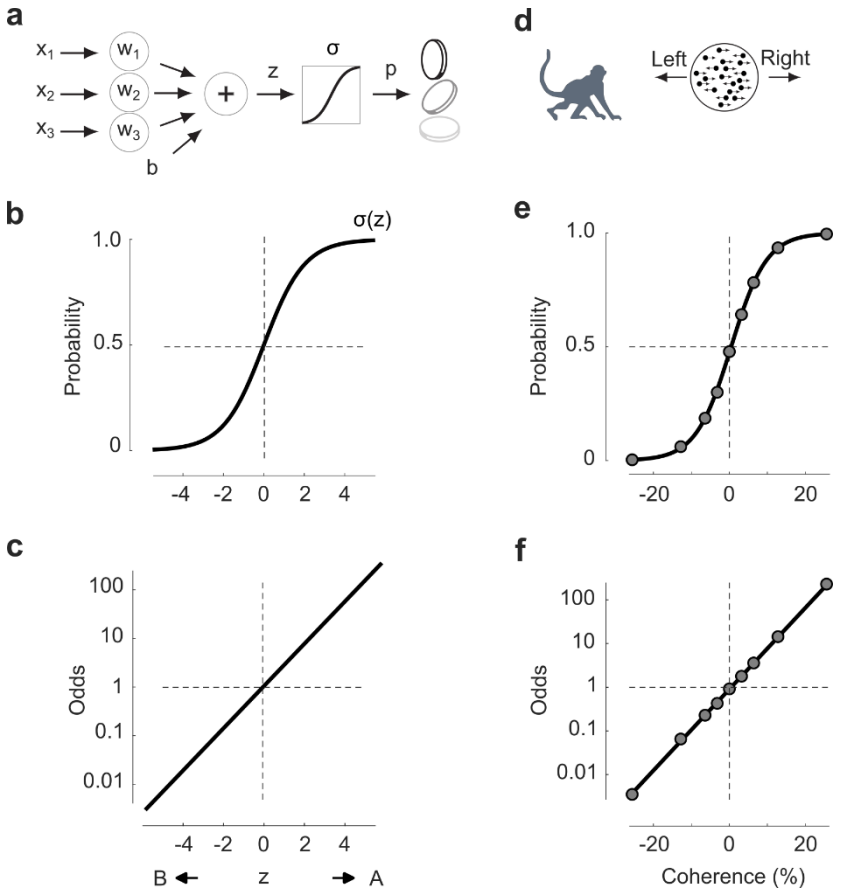

Figure 1. ***Logistic classification describes sensory choices between two alternatives****.* ***a*** *Diagram of logistic classification, in the case of 3 factors x_1, x_2, and x_3. The weighted factors are added to a constant bias b and passed through a logistic function. The result is used to bias a coin, which is then flipped.* ***b****. The logistic sigmoid, plotting probability of choice A vs. decision variable (negative if in favor of choice B, positive if in favor of choice A).* ***c.*** *The logistic sigmoid becomes a line when plotting the odds (in logarithmic scale) vs. the decision variable.* ***d*** *A monkey performing the random dots task, indicating whether most dots move Left or Right.* ***e****. Performance of an example monkey, showing probability of Right choices as a function of coherence (negative for leftward motion, positive for rightward motion). Data were taken from Ref.[11] using a data grabber (grabit.m). Curve shows fit of the logistic classification model.* ***f*** *Same data and fit, in log odds.*

In this example, the measure of stimulus strength x in Eq. 5 is the one used in the laboratory: the coherence of the dots[10,11]. It is often necessary, however, to allow an expansive or compressive deformation that converts the measure used in the laboratory to a scale used inside the brain. For instance, one could raise the stimulus strength to an exponent q, which is fit to the data[25,33,34]. For simplicity, I will write this as x^q, but the proper expression is $sign(x)|x|^q$, which maintains the sign of x. In other cases, one may want functions $f(x)$ that saturate more strongly[12,24].

As we will see later, the weights W for sensory factors are bounded by sensory reliability (the inverse of its standard deviation). By contrast, no bound exists for the weights for other factors, such as those based on value. So, one could give an infinite weight to chocolate – ensuring a deterministic choice of dessert – but not to the rightward motion of random dots.

Sensory choices influenced by priors

In the previous example task (Figure 1**e**,**f**), the stimuli were equally likely to move left or right. In other words, their prior probabilities were equal: $P_R = P_L$. Reflecting this symmetry, the monkey was unbiased ($B \approx 0$): when the stimulus was uninformative (zero coherence), it responded L or R with approximately equal probability.

In general, however, one of the choices may have a higher chance of being correct than the other. In this case, the optimal strategy is to set the bias B to the logarithm of the prior odds:

$$B = \log \frac{P_R}{P_L} \qquad \text{Eq. 6}$$

As proven in Section S3, this value of B is optimal in that it maximizes the overall reward (if the rewards for the two sides are equal, and if the evidence Wx is distributed uniformly).

Therefore, an observer that performs logistic classification and seeks to maximize reward should adjust its bias B according to the prior probability of the stimuli (Eq. 6), without changing its sensitivity W. The bias may not hit precisely the ideal level (Eq. 6), because the observer may not be able to estimate the true prior probabilities, but it should at least move in the right direction.

This prediction is verified in humans performing a biased motion discrimination task: if the odds of a stimulus going rightward P_R/P_L are 80:20 instead of 50:50 (Figure 2**a**), the observers make more Right choices[35] especially at zero coherence, when the stimulus is uninformative (Figure 2**b**). This shift is captured by the logistic classification model (Eq. 5) with the same weight W but a larger bias B (Figure 2**b**). The log odds of the choices fall on a line that is parallel but displaced from the 50:50 case (Figure 2**c**).

Analogous results are seen in mice that perform a biased position discrimination task[22] (Figure 2**d**). In this experiment, stimuli are divided in blocks having higher probability on the Left (80:20) or on the Right (20:80). As expected, the psychometric curves measured in the two blocks are displaced (Figure 2**e,f**), such that at zero contrast (no visual stimulus) Right choices are more likely in 20:80 blocks and Left choices more likely in 80:20 blocks. The model (Eq. 5) captures this difference with a bias B that is positive in the 20:80 block and negative in the 80:20 block, with no changes in sensory weight W.

Sensory choices without evidence

In the absence of sensory evidence, these data exhibit a phenomenon that is widely observed in non-sensory tasks[8,9,36-43], called "probability matching": the probabilities of the observer's choices tend to match the probabilities of those choices being correct. For instance, at zero coherence the human in the 20:80 motion task chose Right with probability ~0.8 (Figure 2**b**). At zero contrast, the mouse in the position task showed a similar (though less pronounced[9,44]) tendency (Figure 2**e**).

This phenomenon is puzzling because it is suboptimal[37,40]. Indeed, in the absence of evidence, one would obtain more rewards by always guessing Right if $P_R > P_L$ and always guessing Left otherwise (Section S3). In other words, one should set B to $+\infty$ if $P_R > P_L$ and to $-\infty$ otherwise (in practice, setting B to +10 or -10 would suffice).

The observer, however, does not know that some trials have no evidence. It may thus apply the optimal value of the bias (Eq. 6) to all trials. This strategy produces probability matching. Indeed, if the evidence is uninformative ($x = 0$), then the probability that the observer chooses R is

$$p = \sigma\left(\log \frac{P_R}{P_L}\right) = \sigma(\sigma^{-1}(P_R)) = P_R$$

where I used the fact that $\log \frac{p}{1-p} = \sigma^{-1}(p)$.

Perhaps brains use Eq. 6 as a heuristic (even when they should be setting $B = \pm\infty$), because in most situations in nature there is a correct answer that can be inferred from factors (present stimuli, past stimuli, actions, outcomes, etc.). This strategy is suboptimal when no factors are relevant (e.g. when gambling), and leads them to perform probability matching.

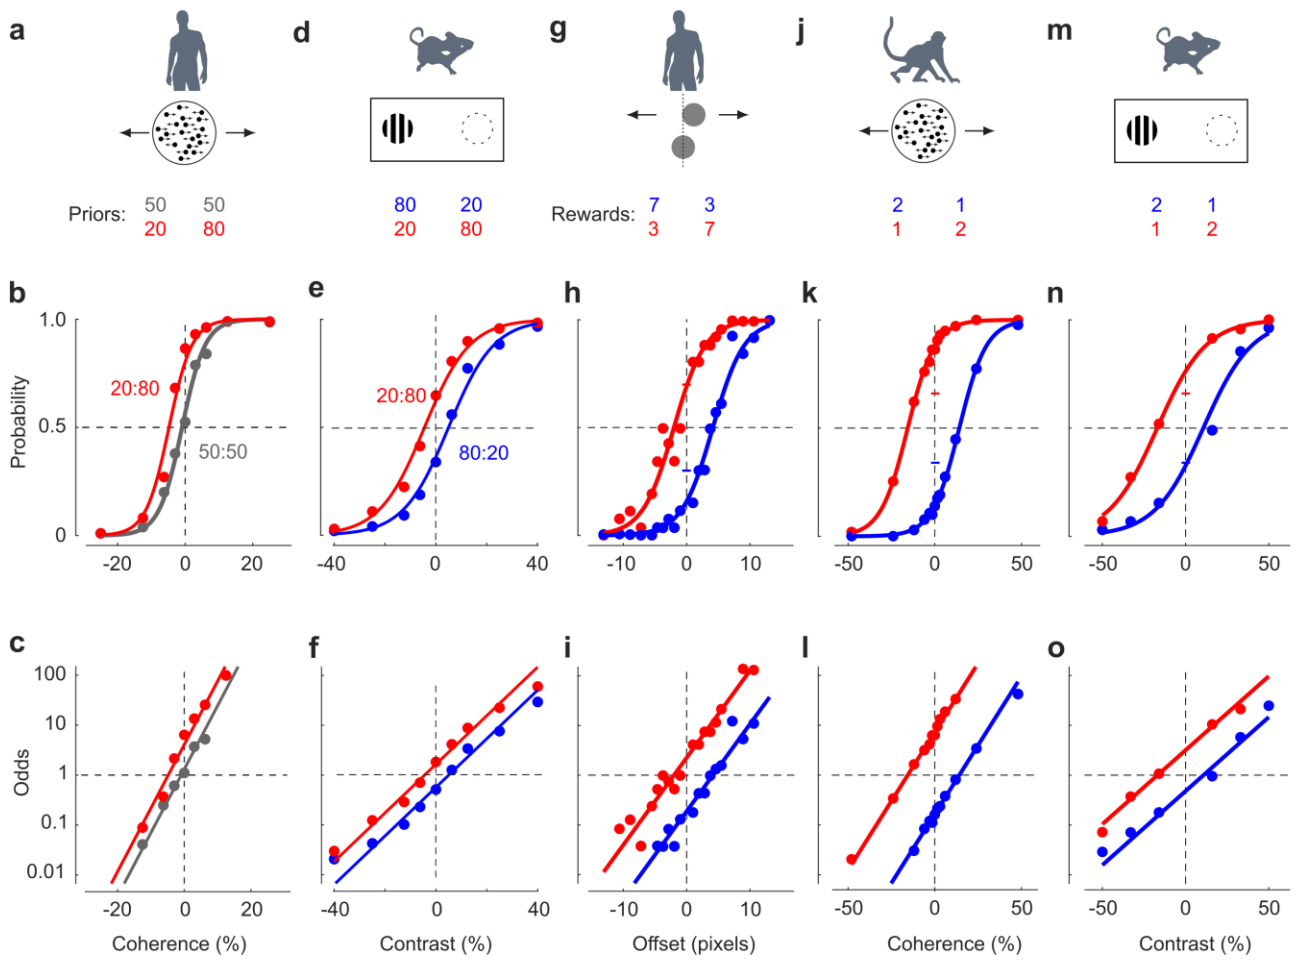

Figure 2. ***Logistic classification describes choices based on stimuli, priors, and values****.* ***a.*** *A random dots task where a human is told that rightward motion is more likely than leftward motion (blue) or that they are equally likely (gray).* ***b****. Data from an observer showing that probability of choosing Right depends not only on stimulus coherence (abscissa) but also on prior probability (blue vs. gray). Data grabbed from Ref.[35]. Curves show fit of the logistic classification model with same weight w but two different values for the bias b.* ***c.*** *The same data and curves plotted as odds of choosing Right vs. Left, in logarithmic scale.* ***d-f.*** *Same, for a mouse indicating whether a visual stimulus appears on the Left or the Right, presented in alternating blocks where Left was more likely (red) or less likely (blue) than Right. Data grabbed from Ref.[22].* ***g****. A position discrimination task where the observer indicates whether the top stimulus is to the Right of the bottom one, in conditions where correct Left responses were rewarded less (red) or more (blue) than correct Right responses.* ***h****. Data from an observer showing that probability of choosing Right depends not only on stimulus offset (abscissa) but also on reward condition (blue vs. red). Data grabbed from Ref.[45]. Curves show fit of the logistic classification model with same weight W but two different values for the bias B. The red and blue ticks in the central axis indicate the intercepts predicted by the optimal strategy.* ***i.*** *The same data and curves plotted as odds of choosing Right vs. Left, in logarithmic scale.* ***j-l****. Same, for a monkey performing a random dots task where stimuli are presented in alternating blocks where correct Left choices were rewarded less (red) or more (blue) than correct Right choices. Data grabbed from Ref.[46].* ***m-o.*** *Same, for a mouse indicating whether a visual stimulus appears on the Left or the Right, presented in alternating blocks where correct Left choices were rewarded less (red) or more (blue) than correct Right choices. Data grabbed from Ref.[47].*

Reassuringly, humans stop performing probability matching if they understand that they are gambling (guessing an independent random variable). In that case, they become more optimal, privileging the option that is more likely to be rewarded[40].

We have seen that if an observer performs logistic classification (Eq. 5) they should set their bias to the log prior odds (Eq. 6), which leads to probability matching (when the evidence is zero). This argument can also be inverted: as shown in Section S4, the only way to achieve probability matching (for any strength of evidence) is to perform logistic classification (Eq. 5) and to set the bias to the log prior odds (Eq. 6).

Sensory choices influenced by value

In the experiments examined up to here, the rewards for correct Left and Right choices were equal. In general, however, the rewards, V_L and V_R, could differ. In this general case, one can prove (Section S3) that the ideal bias B (the bias that maximizes reward) is

$$B = \log\frac{P_R}{P_L} + \log\frac{V_R}{V_L} \qquad \text{Eq. 7}$$

Therefore, if the probabilities of both options are equal ($P_R = P_L$) an optimal logistic classifier should engage in "value matching": its fraction of Right choices should match the ratio of Right rewards to total rewards. Indeed, if $x = 0$ the probabilities of Right choices should be $p = \sigma(\log(V_R/V_L)) = V_R/(V_R + V_L)$.

There is ample evidence that the brain follows this strategy when making purely economic decisions[7], even though it would obtain more reward by sticking to the option with higher value. Once again, logistic classification suggests an explanation for this behavior: the brain may assume that there is relevant evidence x to guide the choice (even when there cannot be any), so it may obey Eq. 7 even when not optimal.

More generally, the optimal logistic classifier (Eq. 7) captures how observers combine sensory signals with considerations of value. For example, when humans perform a position discrimination task (Figure 2**g**), they make more Left or more Right choices depending on which choice yields the larger reward[45]. The data are well fit by logistic classification (Eq. 5) with a single sensory gain W and different biases B in the two reward conditions (Figure 2**h**). Once again, the effects are easiest to judge in terms of log odds (Figure 2**i**): changing the stimulus moves the log odds along a line, and changing the reward ratio moves the log odds to a parallel line. Similar results are seen in monkeys doing the random dot motion task[46] (Figure 2**j-l**), and in mice performing the position discrimination task[47] (Figure 2**m-o**). In all these cases the changes in bias were at least as large as predicted by optimality (Eq. 7). In fact, they were generally larger than optimal[48], perhaps because of a non-linear utility scale that disproportionally favors large rewards[7,49].

Taken together, logistic classification (Eq. 5) and the ideal bias (Eq. 7) indicate how to combine sensory signals with relative probability and relative reward. They show that the logarithm of the odds provides a "natural currency for trading off sensory information, prior probability and expected value"[10].

Sensory choices influenced by motivation

In the experiments examined up to here, observers assigned a constant weight W to sensory evidence (resulting in a constant slope of the log odds line) in the face of changes in stimulus strength, probability, and value. However, paying attention to the evidence may require motivation. If observers are more motivated, they may give a higher weight W to the evidence.

For instance, in an experiment measuring rate discrimination in rats, changes in the ratio of rewards were yoked to changes in the total available reward[26]. Logistic classification provides good fits to those data (Supp Figure 1) and reveals changes both in bias B and in sensory weight W. The changes in bias are expected given the changes in ratio of rewards (Eq. 7). The changes in sensory weight, instead, might result from the changes in total available reward, which might affect task engagement[12] or sensory attention[50]. Variations in these factors may explain why the parameters of logistic classification often fluctuate over time[51,52].

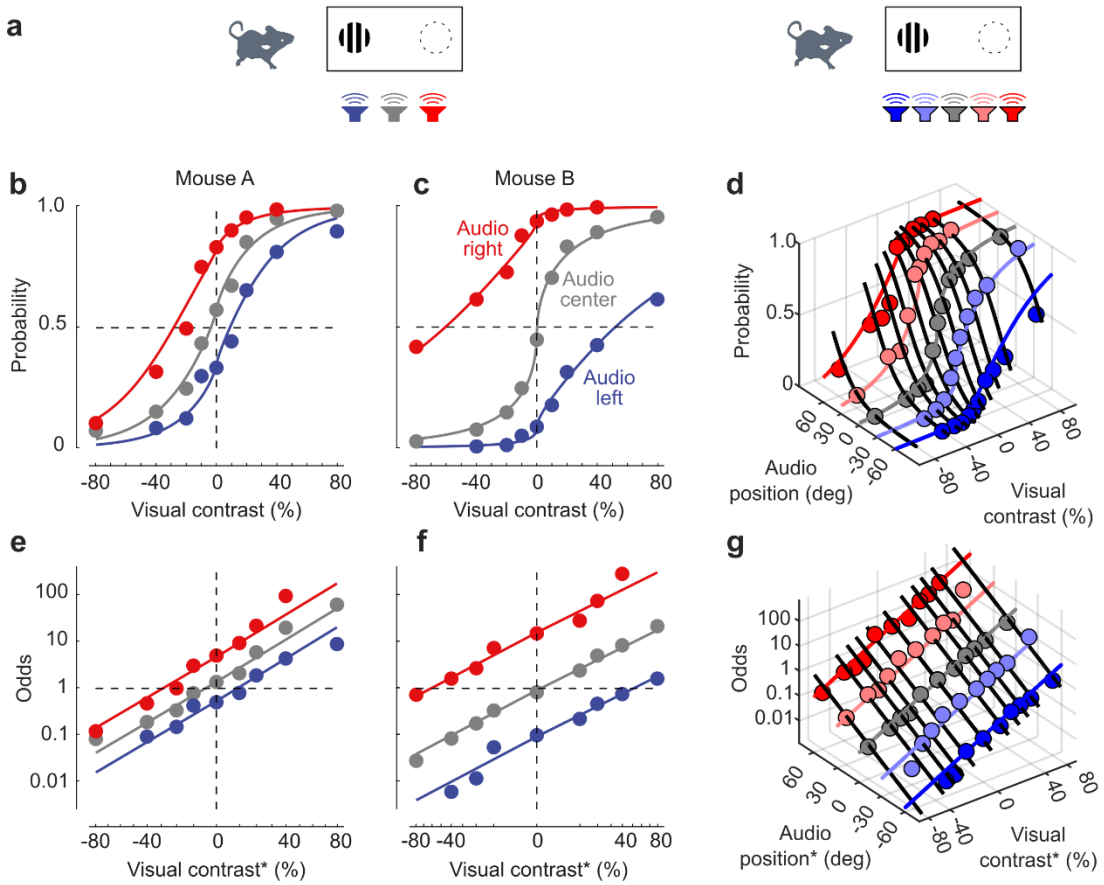

Figure 3. ***Logistic classification describes audiovisual choices****.* ***a****. The audiovisual task: images appear on the Left or right, and sounds are emitted at one of 3-5 positions.* ***b****. Probability that a mouse chooses Right based on the contrast of the visual stimulus (abscissa, negative for stimuli on the Left) and on the position of the auditory stimulus (colors as in a).* ***c****. Same, for a mouse with higher auditory sensitivity.* ***d****. Results in a task with 5 auditory positions.* ***e-g****. The data in b-d, replotted as odds of Right vs. Left choices, in logarithmic scale. The quantities on the abscissa (horizontal plane in g) are compressed (raised to an exponent q<1). Data from Ref.[23]. Therefore, the probability of choosing R is a logistic function* $p = \sigma(z)$ *with three terms: one that depends on vision, one that depends on audition, and a constant.*

Sensory choices across modalities

Logistic classification can describe not only sensory choices based on a single modality but also those based on multiple modalities (multisensory integration). As shown in Eq. 1, logistic classification naturally accommodates multiple factors as input (Figure 1**a**). If there are two sensory stimuli s_1, s_2 of different modalities (e.g., vision and audition), and these stimuli are statistically independent, Bayes' rule (as shown in Section S4) predicts that an observer should perform logistic classification based on their sum[23]:

$$z = \log\left(\frac{p}{1-p}\right) = f_1(s_1) + f_2(s_2) + B \qquad \text{Eq. 8}$$

This expression is just the sum of three factors, and yet it describes multisensory choices in a variety of species, tasks, and combinations of modalities.

For instance, Eq. 8 describes the choices of mice performing an audiovisual detection task[23]. In this task, a visual pattern appears on the Left or on the Right and a sound is played in one of 3 or 5 positions (Figure 3**a**). If the sound is in the middle, it is uninformative, and the task becomes purely visual (Figure 3**b,** *gray*). If it is on the Left or Right, the mouse is more likely to make those choices (Figure 3**b,** *blue* and *red*).

Logistic classification (Eq. 8) describes this behavior accurately, as the sum of three factors:

$$z = F_1 s_1^q + F_2 s_2^r + B \qquad \text{Eq. 9}$$

Here, the variables s_1 and s_2 are visual and auditory attributes, raised to exponents q, r (if necessary). The remaining parameters are the visual gain F_1, the auditory gain F_2, and the bias B. The first term describes the responses when the task is purely visual (Figure 3**b,** *gray*), and the second describes the responses when the task is purely auditory (Figure 3**b,** where abscissa is zero). The model provides excellent fits (Figure 3**b***, curves*), correctly describing the choices not only when the two modalities agree, but also when they contradict each other[23]. Similar results are seen in a mouse with better auditory localization (higher performance at zero contrast, Figure 3**c**), and in a mouse tested with more auditory positions (Figure 3**d**).

As usual, these results are easiest to understand in terms of log odds (Figure 3**e**-**g**). In this representation, and with a slightly different abscissa (s_1^q and s_2^r instead of s_1 and s_2) the results fall on parallel lines: changing visual contrast moves the points along the lines, and changing the auditory position shifts them to a different line.

For such high-dimensional data, one way to judge the performance of the model[31,53] is to plot the probability of all points as a function of the predicted decision variable z (Eq. 1) and see if they fall near the logistic curve $\sigma(z)$ (Eq. 4). Here they do (Supp Figure 3).

Logistic classification describes multisensory choices in multiple other species across multiple modalities. Studies of multisensory integration have typically analyzed the responses with signal detection theory (Section S1), and found that observers are roughly[54,55] Bayesian[23,26,56-61] (at least when all stimuli are presented for the same duration[62]): they weigh each modality according to its reliability[56,57]. Below I examine three such studies[55,59,61], which tested multisensory integration with different tasks, across different modalities (visual, auditory, tactile, and vestibular), in different species (rats, monkeys, and humans), and I show that their data are well described by logistic classification.

In a study that measured visuo-tactile integration[55], rats indicated whether the orientation of a stimulus was closer to horizontal or vertical by choosing between Left or Right response ports. The stimulus could be purely visual or purely tactile (Supp Figure 2**a,b**) or a combination of the two (Supp Figure 2**c**). I applied the logistic classification model to those data with 4 free parameters: visual gain F_1, tactile gain F_2, their common exponent q (there was no need for different exponents), and bias B. The model provided good fits (Supp Figure 2**a**-**c**), and would have done fairly well even with just two free parameters (setting $q = 1$ and $B = 0$, not shown). The model's performance is easy to evaluate if the data are plotted as log odds: aside from the (small) bias, the line for multisensory stimuli (Supp Figure 2**f**) is just the sum of the lines for unisensory stimuli (Supp Figure 2**d**,**e**). An additional check verifies that the predicted decision variables for all these measured data describe a single logistic curve (Supp Figure 3**a-c**).

Sensory choices for multimodal features

The other two studies[59,61] tested multisensory integration in more detail, by varying both the amplitudes and the features of two modalities, and by introducing offsets between those features.

To discuss these experiments, it helps to define stimulus "feature" vs. "amplitude". The feature ϕ is the attribute on which the choice is based; it can be positive or negative, and is zero at the boundary between the two choices. The amplitude α, instead, is positive and is zero if the stimulus is absent. For instance, in the motion task (Figure 1**d**-**f** and Figure 2**a**-**c**), ϕ is direction (±1) and α is coherence (between 0 and 1). Similarly, in the position task (Figure 2**d**-**f**), ϕ is position and α is contrast.

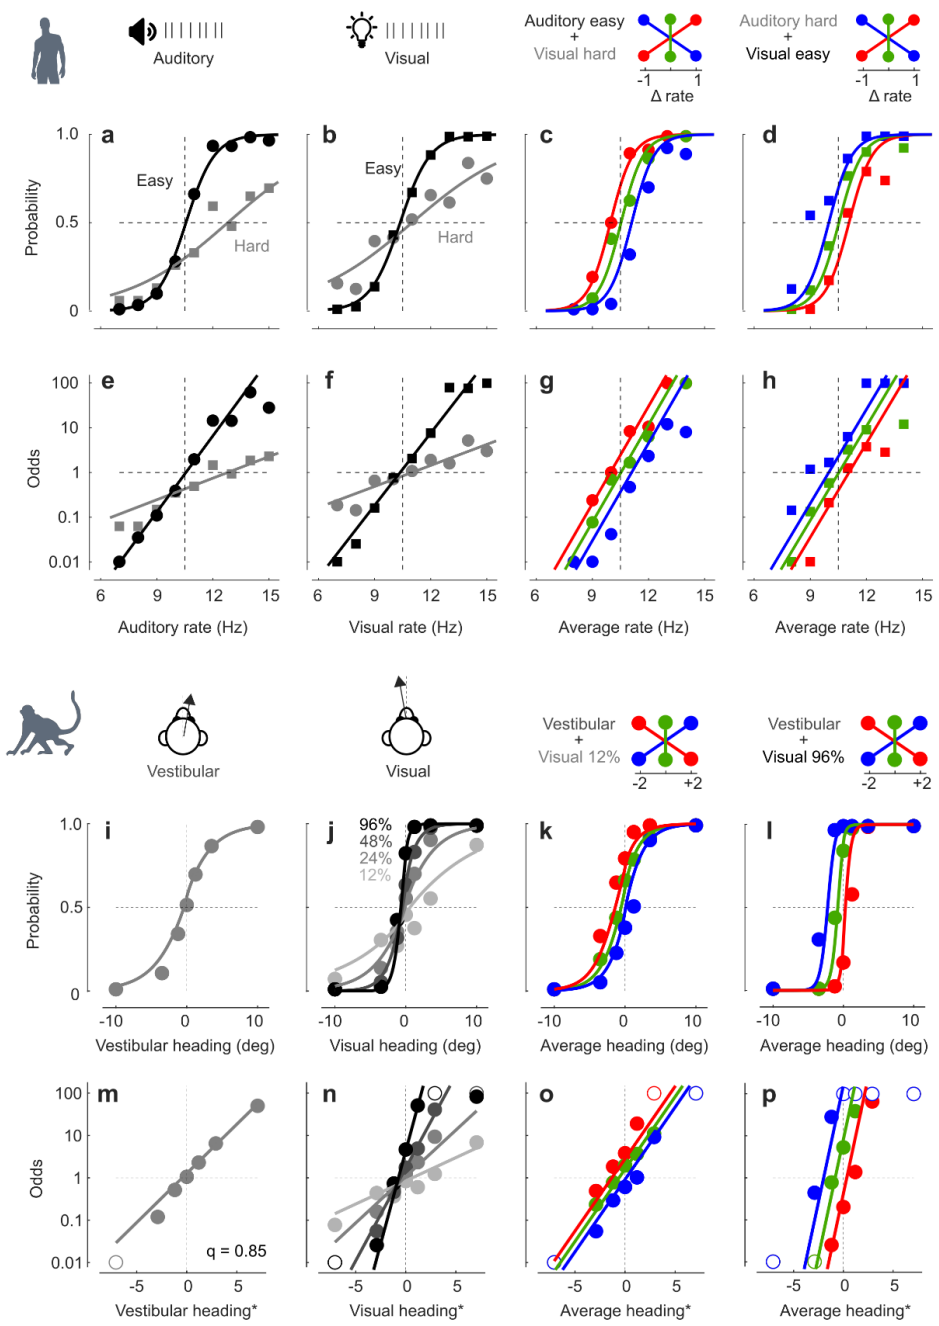

Figure 4. ***Logistic classification describes audiovisual and vestibulo-visual choices with stimuli differing in feature and amplitude****.* ***a*** *Average choices of a human performing a rate discrimination task based on auditory sequences. The ordinate plots the probability that the observer judges the rate of the sequence to be higher than a reference (10.5 Hz) when stimuli have low amplitude (Hard) or high amplitude (Easy). Curves in this panel and subsequent ones are fits of the logistic classification model.* ***b****. Same, for visual stimuli.* ***c****. Same, for stimuli obtained by summing an easy auditory stimulus to a hard visual stimulus. As illustrated in the inset, the rate of the auditory stimulus was lower (blue), equal (green), or higher (red) than the rate of the visual stimulus. The abscissa plots the average rate of the visual and auditory stimuli.* ***d****. Same, for stimuli obtained by summing a hard auditory stimulus to an easy visual stimulus.* ***e-h:*** *The same data as a-d, plotted in terms of odds, in logarithmic scale. Data are from Human 1 in Ref.[61].* ***i*** *Average choices of a monkey performing a heading discrimination task based on unisensory vestibular stimulation. The ordinate plots the probability that the observer judges the heading to be Right. Curves are predictions of the logistic classification model.* ***j****. Same, for unisensory visual stimuli varying in coherence from 12% (hardest) to 96% (easiest).* ***k****. Same, for multisensory stimuli obtained by summing the vestibular stimulus to a hard visual stimulus, with the former having a heading left of (blue), same as (green), or right of (red) the latter. The abscissa plots the average heading of the vestibular and visual stimuli.* ***l****. Same, for stimuli obtained by summing the vestibular stimulus to an easy visual stimulus.* ***m-p:*** *The same data as a-d, plotted in terms of odds, in logarithmic scale. Open symbols indicate values where the probability is close to 0 or 1, where the odds would require a high precision, and are thus likely to lie outside the range of the ordinate.* ***Data were grabbed from Ref.[59].***

If a stimulus s can have multiple combinations of amplitude and feature, we must define how those combinations affect the stimulus representation $f(s)$. A reasonable working hypothesis is that the stimulus representation is separable, i.e., the product of a function of amplitude and one of feature: $f(\phi, a) = g(\phi)h(\alpha)$. In the visual system, for example, this hypothesis would be appropriate for orientation and contrast[63] (but perhaps not for temporal frequency and contrast[64]).

Under this hypothesis, when there are two sensory modalities and thus two amplitudes α_1, α_2 and two features ϕ_1, ϕ_2, the model (Eq. 8) becomes

$$z = g_1(\phi_1)h_1(\alpha_1) + g_2(\phi_2)h_2(\alpha_2) + B \qquad \text{Eq. 10}$$

To test this model, I considered a task where humans received sequences of visual and auditory stimuli and judged whether their rate was above a criterion[60,61]. As expected, when stimuli were unisensory, the observers did well if the stimulus had high amplitude ("Easy"), and poorly if it had low amplitude ("Hard", Figure 4**a**,**b**). When the easy auditory stimulus was combined with the hard visual stimulus, it largely determined the observer's choices: if the two stimuli disagreed, the observer responded in a manner consistent with the auditory stimulus (Figure 4**c**). Similar results were obtained by combining the hard auditory stimulus with the easy visual stimulus (Figure 4**d**). In both cases the choices were largely determined by the easier stimulus (as expected from a Bayesian observer[56,57,61]).

The choices measured in this task agree with logistic classification, with amplitudes and features combined under the separable approximation (Eq. 10). To fit the data, I chose linear expressions for the functions: $h_i(\alpha_i) = \alpha_i$, and $g_i(\phi_i) = F_i\phi_i + B_i$ (with $i = 1,2$ denoting visual and auditory modalities). The result is the following expression:

$$z = \alpha_1(F_1\phi_1 + B_1) + \alpha_2(F_2\phi_2 + B_2) + B \qquad \text{Eq. 11}$$

With just 5 free parameters (visual gain F_1, visual bias B_1, auditory gain F_2, auditory bias B_2, and overall bias B) this function produced 10 psychometric curves that fitted 114 measurements (Figure 4**a**-**d**, *curves*).

As usual, model and data are easiest to inspect as log odds. In this scale, the unisensory responses are lines that are steeper for Easy stimuli than for Hard stimuli (Figure 4**e**,**f**). When the steep line from one modality is summed to the shallow line from the other, its largely determines the slope and intercept of the resulting line (Figure 4**g**,**h**).

In another multisensory task, monkeys judged whether their heading was Left or Right based on visual and vestibular cues[59]. The animals performed well when presented with purely vestibular stimuli (Figure 4**i**) and with purely visual stimuli with high coherence (Figure 4**j**). When the vestibular stimulus was combined with the hardest visual stimulus (lowest coherence), the former largely determined the observer's choices (Figure 4**k**). Conversely, when the vestibular stimulus was combined with the easy visual stimulus, the latter largely determined the choices (Figure 4**l**). As expected from a Bayesian observer[59], therefore, the choices were largely determined by the modality that was easier.

These data were well fit by the logistic classification model under the separable assumption (Eq. 11). Here, $\phi_1$indicates the visual heading and ϕ_2 indicates the vestibular heading, both raised to an exponent q to capture increased sensitivity around the heading of 0. This function provided good fits to the psychometric curves (Figure 4**i**-**l**), and as usual, the reasons for its success are easiest to see in terms of log odds (Figure 4**m**-**o**).

It is striking that logistic classification describes multisensory choices in so many species, tasks, and modalities, especially considering that the tasks do not conform to the conditions where logistic classification is optimal.

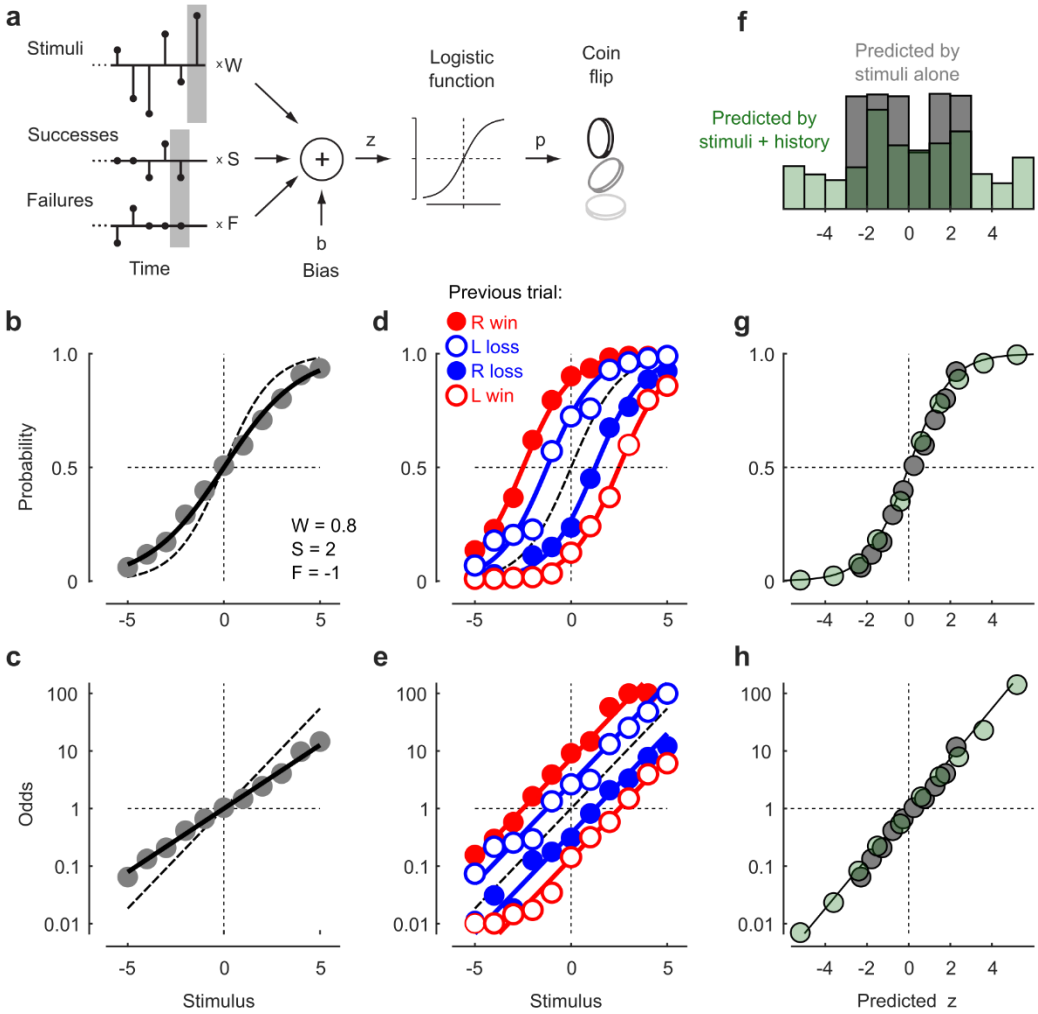

Figure 5. ***Logistic classification describes effects of history. a****. Model of logistic classification where choices are determined not only by a stimulus but also by prior wins or losses (Eq. 12). Adapted from Ref.[12]. The subsequent panels show results of simulations with W = 0.8, S = 2, F = -1, B = 0. See Fig. 2c of Ref. [74] for similar data.* ***b****. Simulated data showing the probability of R choices as a function of contrast, compared to what would be observed based on the sensory factor alone (dashed).* ***c****. Same, plotted as the logarithm of the odds.* ***d-e****. Same as b-c, contingent on previous trial being a success or a failure (red vs. blue) on the L or R (open vs. closed). The dashed curve shows the result that would be observed based on the sensory factor alone.* ***f****. Distribution of values of the decision variable z, as predicted by the stimuli alone (gray) or by the stimuli plus the history terms (green).* ***g.*** *Binned values of the decision variables for the two models, on top of a logistic curve.* ***h.*** *Same, plotted as the logarithm of the odds.*

Indeed, logistic classification (Eq. 8) is optimal if the modalities are independent[23] (Section S4), but here the two modalities had positive correlations. For instance, in the audiovisual position task (Figure 3) the modalities agreed more often than they disagreed[23]. More strikingly, in the visuotactile task (Supp Figure 2) the two modalities had a correlation of 1 when both present, because they always had the same orientation. Strong correlations between modalities were present also in the other two tasks (Figure 4). In these conditions, logistic classification may no longer be optimal, but the brain seems to adopt it nonetheless as a useful heuristic[65].

Sensory choices influenced by history

A key challenge in making a choice is to weigh only the factors that are relevant: giving weight to irrelevant factors decreases the chance of correct answers. And yet, superstitious behavior is widespread. For instance, in sensory tasks where the sequence of stimuli is random, humans and other animals[12-22,49,66-75] typically act as if it followed some order ("serial dependence").

The logistic classification model can capture this serial dependence by including history factors as regressors in addition to the sensory factors[12-22,73,76]. This approach was obtained to explain apparently erratic choices exhibited by mice in a visual task[12]. It merged previous logistic models that had separately used sensory factors[10] and history factors[7-9].

In this framework (Figure 5**a**), choices depend not only on the present sensory stimulus $x(t)$ as in Eq. 5, but also on whether the preceding choice was a success s or a failure f:

$$z(t) = Wx(t) + Ss(t-1) + Ff(t-1) + B \qquad \text{Eq. 12}$$

where $s(t) = sign\big(x(t)\big), f(t) = 0$ if trial t was a success, and $s(t) = 0, f(t) = -sign\big(x(t)\big)$ if it was a failure.

This equation describes just one of various possibilities, chosen for its simplicity. In the studies that used this approach[12-22,73,76], the history regressors typically extend further in the past, or are expressed in different combinations such as past choices (regardless of success) or past rewards (regardless of choice).

If serial dependence is present, ignoring it can make one underestimate stimulus sensitivity. Indeed, plotting the performance as a function of the present stimulus yields shallower psychometric curves (fewer correct answers) than would be observed based only on sensory sensitivity W (Figure 5**b**,**c**). A better estimate of sensitivity (and a measure of the effect of history) is obtained by plotting performance conditional not only on the present stimulus but also on the choice and outcome of the previous trial (Figure 5**d**). If the effect is additive (as in Eq. 12), it will simply shift the lines representing the log odds of the choices, because the history factors sum with the fixed bias B to yield a variable overall bias (Figure 5**e**). One can confirm that these history factors are in play by comparing a model that ignores them (Eq. 5) to one that includes them[53] (Eq. 12). In our example, the model that includes history factors predicts larger (positive or negative) values for the decision variable z (Figure 5**f**), while providing good fits (Figure 5**g**,**h**).

The strength of the history terms varies across individuals and in time[12]. It can change with learning[51] and can be increased or decreased by appropriate design, but it cannot be removed[16]. In humans, for instance, it correlates with education[16]: PhD candidates or graduates have strongly negative failure weights (lose-switch) but weak success weights, whereas the remaining observers have positive success weights (win-stay). This correlation is not necessarily causal and may reflect other unmeasured factors.

However, some effects of history are more complex than can be captured by a simple logistic classification model such as Eq. 12. For instance, past choices have a stronger effect if they were made with low confidence[71], a behavior that is more readily explained by models based on reinforcement learning[47]. Also, engagement in the trial typically fluctuates, and these fluctuations can be captured by allowing the weights to change across trials[51,52]. Variations in engagement will vary the fraction of successes, so observing data after a success vs. a failure may be a way to gauge the effect of engagement. This phenomenon may explain results observed in rats during rate discrimination[26], where past successes and failures affect not only the overall bias but also the sensory weight W, which was slightly higher after a success than after a failure (Supp Figure 1**g**-**l**). Given that successes tend to occur in an engaged state and failures in a disengaged state, the latter effect may be explained by a dependence of sensitivity on engagement.

Sensory choices during brain manipulations

Logistic classification is a good model not only for how the brain makes decisions, but also for how these decisions are modified by brain manipulations such as local inactivations[23-26]. If the effects of inactivating a brain region can be summarized by changes in a few weights, one can surmise that the region is necessary for processing the corresponding factors.

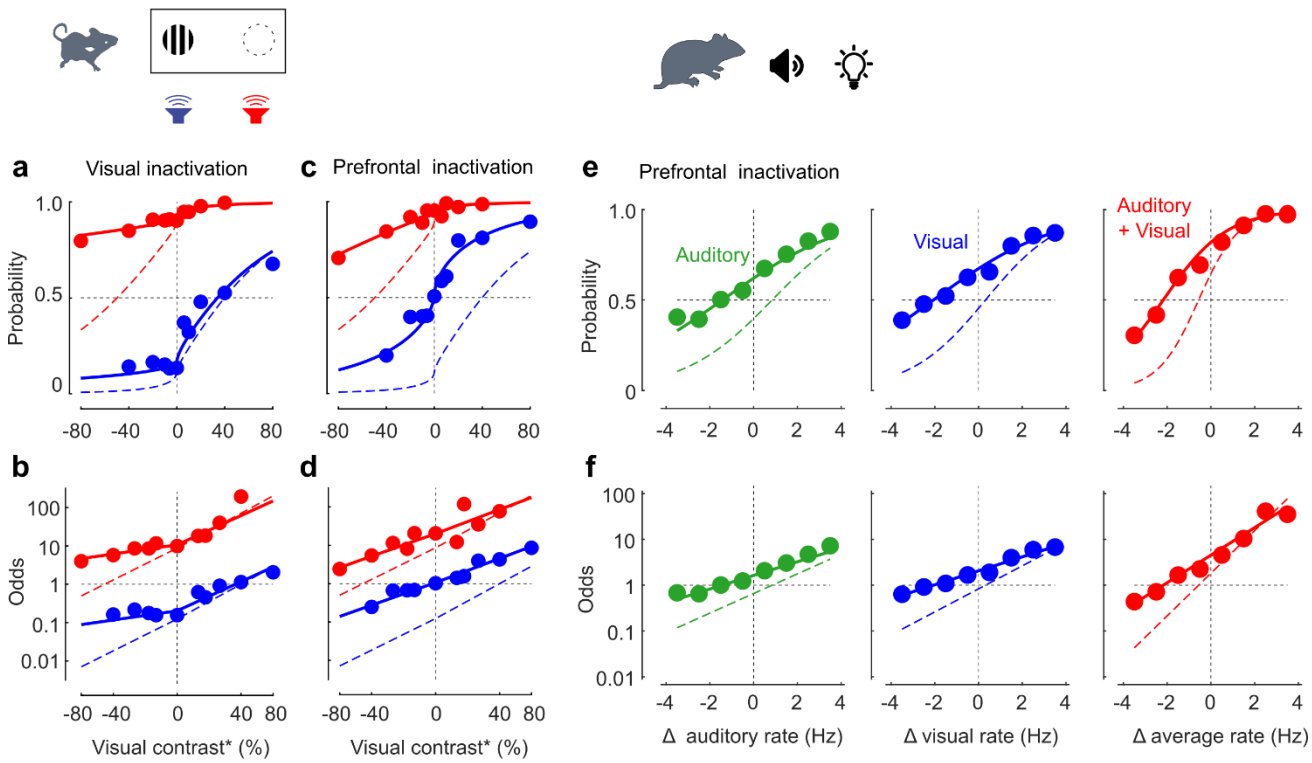

Figure 6. ***Logistic classification describes effects of brain manipulations. a****.Effects of inactivating right visual cortex in mice that perform an audiovisual localization task[23]. Showing the probability of Right choices as a function of visual contrast (abscissa) in the presence of sounds from the Left (blue) and Right (red), in control conditions (dashed) and during inactivation (data and continuous curves). The curves are predictions of the model (Eq. 14) allowing only one parameter to change with inactivation: the weight applied to visual stimuli on the left,* F_{1L}*.* ***b****. Same data, plotted as log odds.* ***c,d****. Effects of inactivating the right prefrontal cortex. The model captured them via changes not only in the contralateral visual weight but also in the contralateral auditory weight and in the overall bias. Replotted from Ref. [23].* ***e****. Effects of prefrontal inactivation in a rat performing rate discrimination based on auditory (green), visual (blue) and audiovisual (red) cues, showing the probability of choosing "Right" (for high rates) as a function of rate (expressed relative to the reference).* ***f****. Same data, in log odds. Data from Ref. [26].*

For example, in mice performing the audiovisual localization task (same as Figure 3), inactivation of the right visual cortex decreases the contralateral visual weight, but no other weights. As a result, it impairs the detection of visual stimuli on the left side, but does not affect visual detection on the right side or auditory detection on either side (Figure 6**a**). This effect is particularly easy to see in terms of log odds (Figure 6**b**).

To describe these asymmetric effects, it helps to extend Eq. 9 to allow for different weights for stimuli on the left vs. right[23]. To this end, we split the decision variable z into two variables z_L, z_R summarizing the evidence in favor of the two choices:

$$z = z_R - z_L + B \qquad \text{Eq. 13}$$

and we give each of them a visual weight ($F_{L,1}$ and $F_{R,1}$) and an auditory weight ($F_{L,2}$ and $F_{R,2}$):

$$\begin{aligned} z_L &= F_{L,1} s_{L,1}^q + F_{L,2} s_{L,2} \\ z_R &= F_{R,1} s_{R,1}^q + F_{R,2} s_{R,2} \end{aligned} \qquad \text{Eq. 14}$$

This model has 6 parameters (four sensory weights, the exponent, and the bias), and in principle a brain manipulation could change all those parameters relative to the control condition. Typically, however, the control condition is described by fewer parameters (e.g., the weights on the left and right are the same), and brain manipulations change only a few parameters[23-26]. For instance, inactivating right visual cortex affects only the

weight applied to left visual stimuli, $F_{L,1}$ (Figure 6**a**,**b**). This result indicates that right visual cortex is only necessary for processing visual stimuli, and only those on the left[23].

By contrast, inactivating right prefrontal cortex reduces both visual weights and auditory weights (regardless of side), and increases bias B thus increasing the proportion of Right choices (Figure 6**c**,**d**). This result indicates that prefrontal cortex processes bilateral visual and auditory signals and promotes contralateral choices[23].

Similar results are seen in rats performing the audiovisual rate discrimination task[26] (same as Figure 4): the control and inactivation conditions are well fit by logistic classification (Eq. 14), and inactivation reduces both auditory and visual weights (regardless of whether stimuli have low or high rate), and increases the bias (Figure 6**e**-**j**).

Sensory choices without flipping coins

These results indicate that it is useful to think of the brain as computing separate variables z_L, z_R summarizing the evidence in favor, and more generally the utility, of each choice (Eq. 13). Each of the variables is obtained by a weighted sum of factors (similar to Eq. 1). The overall decision variable z is obtained by subtracting the utilities of the two choices (Figure 7**a**) .

However, this implementation still involves a set of deterministic factors followed by a stochastic decision (a coin flip). The brain may be faced with the opposite situation: a set of stochastic factors, especially if these factors are sensory, followed by a deterministic choice. Logistic classification can be achieved also this way, in alternative implementations that result in the same equations used throughout this paper.

The first alternative implementation has strong similarities to signal detection theory[27-29] (Figure 7**b**). In this implementation, the two utilities z_L, z_R determine the means of two random variables, which vary from trial to trial (e.g. because of sensory noise) and are distributed according to a specific distribution (called a Gumbel distribution). The two utilities are then subtracted, and if the result is above zero the choice is R, otherwise it is L. This approach is used in economics[1,3] to model differences among consumers, and in signal detection theory[27-29] to model sensory noise (Section S1). In this implementation, the weight W that an ideal observer would assign to a sensory factor (Eq. 1) is proportional to the factor's reliability (the inverse of standard deviation of the noise). This represents an upper limit: a less ideal observer would assign a lower weight.

The second alternative implementation is analogous to the widely used drift diffusion model[2,30-32]. In this implementation, a single decision variable z controls the drift of a random walk, and the choice is made when the walk hits one of two boundaries (Figure 7**c**). This implementation is identical to the logistic classification model if the starting point is midway between the two bounds[31,77]. Indeed, as detailed in Section S2, if it takes N steps to reach a bound[78], and the probability to go up (vs. down) at each step is set to $P = \sigma(z/N)$, then the probability of hitting the upper bound is $p = \sigma(z)$. If one chooses z as the weighted sum of factors (Eq. 1), this is the logistic classification model.

The drift diffusion model predicts not only the probability but also the timing of the choices[2,30-32] (Section S2). For instance, it predicts that larger values of a decision variable result in faster decisions.

There are many ways to change the drift diffusion model, e.g. with an asymmetric starting point or with diverse relationships between P and sensory inputs and biases[79]. In some of those cases there may not be a closed form equation for the probability of hitting a boundary. In those cases, the psychometric functions must be computed numerically[80].

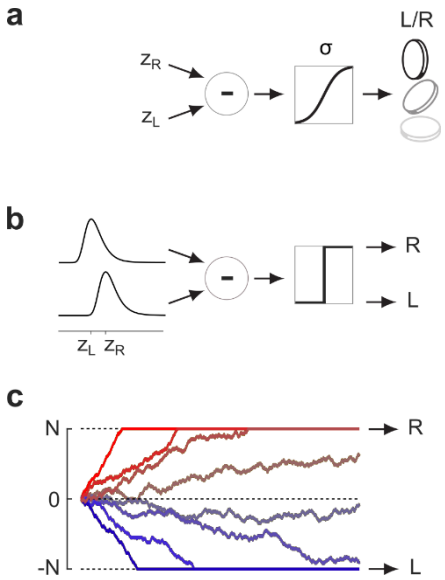

Figure 7. ***Three implementations of logistic classification. a.*** *The implementation described in most of this paper, based on deterministic utilities z_L, z_R informing a stochastic decision.* ***b****. An alternative implementation where the two utilities are stochastic across trials, and are then subtracted and thresholded. This is analogous to signal detection theory[27-29].* ***c****. An implementation where the decision variable z controls the drift of a random walk, and the choice is made when the walk hits one of two boundaries. This is the widely used drift diffusion model[2,30-32].*

Sensory choices between more alternatives

In all the analyses above, we have assumed that the observer faces a binary choice. However, this is not a strict requirement: The logistic classifier is readily extended to the case when the choice is one of n alternatives. First, one obtains a vector of decision variables $\boldsymbol{z} = z_1, z_2, \dots, z_n$, one for each alternative:

$$\boldsymbol{z} = \boldsymbol{W}^T \boldsymbol{x} + \boldsymbol{B}$$

where$\boldsymbol{x} = x_1, \dots, x_f$ is the usual vector of factors, $\boldsymbol{W}$ is a $f \times n$ matrix of weights, and $\boldsymbol{B} = B_1, B_2, \dots, B_n$ is a vector of biases, one for each possible choice. Then, one applies a smooth version of the argmax function, called *softmax*:

$$p_i = softmax(\boldsymbol{z})_i = \frac{\exp(z_i)}{\sum_j \exp(z_j)}$$

For example, if there are 3 possible choices, and we assume for simplicity that $z_3 = 0$, then the probability p_1 of the 1st choice is largest when $z_1 > 0$ and $z_1 > z_2$, dividing the space of probabilities in three regions where the three choices are more probable than the others (Figure 8**a**).

Reassuringly, when there are only two outcomes the softmax function reduces to the familiar logistic sigmoid $p = \sigma(z)$ with decision variable $z = z_2 - z_1$:

$$p_1 = \frac{\exp(z_1)}{\exp(z_1) + \exp(z_2)} = \frac{1}{1 + \exp(z_2 - z_1)} = \sigma(z)$$

In a network, the softmax can be thought of as a competitive process among units that each receive one z_i as external input[5]. Because the outputs of softmax sum to 1, an increase in the input to one unit decreases the activity of others. When the difference between the largest z_i and the others is large, this process becomes a winner-take-all.

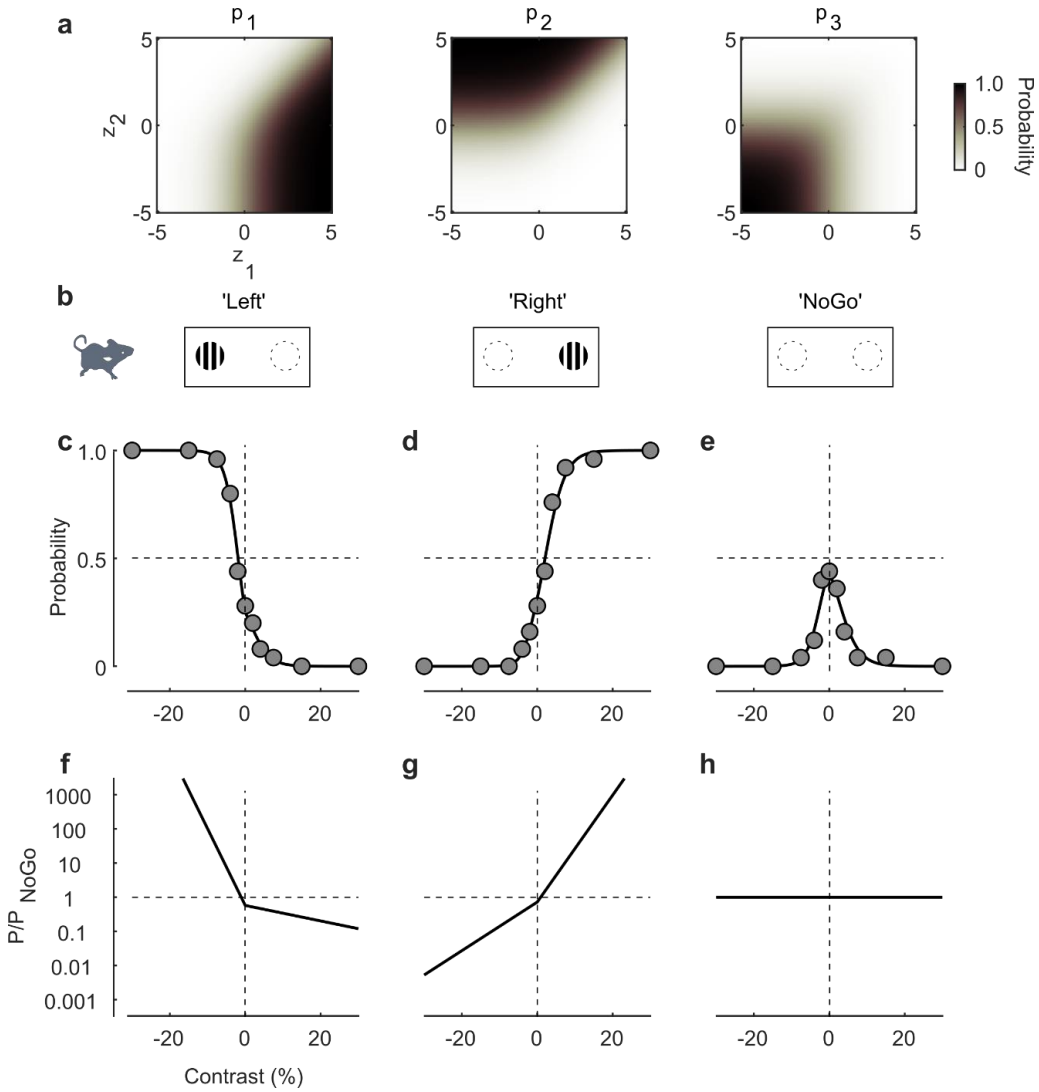

Figure 8. ***Logistic classification with more than two choices. a****. Multinomial logistic classification among 3 choices , showing probability of choices 1,2, and 3 (left to right) as a function of z1 (abscissa) and z2 (ordinate), when z3 = 0.* ***b****. The 3-choice version of the position detection task, where mice choose "Left" or "Right" to indicate stimulus position, and "NoGo" if the stimulus is absent.* ***c-e****. Probabilities of the three choices for a mouse in this task (Mouse I from Fig. 3 of Ref.[24]), as a function of stimulus contrast.* ***f-h****. The odds of the three choices relative to the third choice.*

A logistic classifier with multiple outcomes can provide excellent fits to psychophysical data[24,25]. For example, consider a 3-choice version of the position detection task[24], where the mouse chooses the stimulus position as "Left" or "Right" if the stimulus is present, and chooses "NoGo" if it is absent (Figure 8**b**). This three-choice design was originally devised in the framework of signal detection theory[81,82] but it is particularly easy to analyze in terms of logistic classification, where it boils down to two equations[24]:

$$z_L = \log\left(\frac{p_L}{p_{NoGo}}\right) = w_{LL}c_L + w_{RL}c_R + B_L$$

$$z_R = \log\left(\frac{p_R}{p_{NoGo}}\right) = w_{LR}c_L + w_{RR}c_R + B_R$$

where c_L and $c_R$are the stimulus contrasts on the left and right. These equations provide good fits to the mouse choices (Figure 8**c-e**). As usual, the model is easiest to understand in terms of log odds, i.e., by plotting the values of z_L and z_R (Figure 8**f-h**).

Discussion

A growing body of evidence indicates that the brain uses logistic classification to make choices based on a multitude of factors alone or in combination. These factors include relative value, relative probability, sensory inputs from one or more modalities, recent history, and brain manipulations.

To illustrate this body of evidence, I have followed a hybrid approach, in some cases reviewing existing results and in others applying the logistic classification model to data that were previously analyzed with other approaches. I acquired most of these data manually by clicking on the data points in the published plots. This introduced noise, and is particularly an issue when the probabilities are near 0 or 1, where small variations have large effects on the log odds. Given these limitations, the fits that I provided are evidence for plausibility, but not for superiority over other fits.

I have shown that an observer that performs logistic classification, and that believes that the factors are at least potentially informative, should set its bias to perform "probability matching" (or "value matching"), because this will maximize the reward. Indeed, as we have seen, observers typically adjust their bias to approximate this optimum value. If there are no informative factors, these strategies are suboptimal: it would be better to always choose the option that is more likely or has higher value (Section S3).

This suggests a reason for probability matching and value matching: the observer does not know that the evidence is irrelevant or absent (i.e. that they are gambling). Rather, they act as if they were performing a task based on evidence that is at least potentially informative. Indeed, we have seen that logistic classification is also widely used in cases where it is not optimal, for instance when stimuli are absent (and choices exhibit probability matching or value matching) or stimuli of two modalities are not independent. This suggests that logistic classification is a heuristic that the brain implements in a variety of conditions even when it is not strictly optimal[65].

Throughout the paper, I have described logistic classification as an addition of factors that each contribute to the log odds of a choice. However, we could have just as easily talked about multiplication of factors, each of which contributes to the odds, and we could have taken the logarithm only after the multiplication. Mathematically, the two are identical but numerically it is easier to deal with log odds that vary, say, between -5 and 5 (for a natural logarithm), than with the corresponding odds, which would vary between less than one hundredth and over a hundred. Using firing rates, it is easier to encode the former than the latter. Indeed, when decoding the activity of neurons in prefrontal cortex the activity seems more commensurate to log odds[13,23].

As to how the brain may perform logistic classification, I presented two implementations that correspond to well-established models. One echoes the framework of signal detection theory[27-29], as it involves taking the difference between two stochastic variables and thresholding the result. The other is the well-established drift diffusion model[2,30-32], which performs logistic classification if the starting point is midway between the bounds. Logistic classification, therefore, bridges two established models of decision making: signal detection theory and drift diffusion.

Finally, we have seen that logistic classification is a good model not only for how the brain makes decisions, but also for how these decisions are modified by manipulations such as local brain inactivations. Indeed, the effects of inactivations can often be summarized by changes in a few weights, suggesting that the inactivating regions are necessary for processing the corresponding factors.

Going forward, a promising direction is to develop the dynamic aspect of logistic classification. For instance, logistic classification can capture the effect of stimulus fluctuations, teasing apart the dynamics of cognitive computations[83,84]. At a longer time scale, it seems clear that the weights of the logistic classifier should not be considered constant, because they change with factors such as learning and brain state[51,52].

Another important goal is to understand what limits the weights assigned to sensory stimuli, and therefore what limits performance. In some sensory tasks, the weights are limited by sensory noise (as typically captured by signal detection theory, Section S1). However, in many other tasks, the limiting factor is the appropriateness of the decoder. For instance, in an orientation discrimination task, the limiting factor is not the ability of the primary visual cortex to represent orientation, which is spectacularly high even in mice[85], but the ability of the rest of the brain to decode that representation as required by the task. Another potential limiting factor is a general hesitance to assign large weights to sensory stimuli.

Finally, a promising direction is to use neural data as a predictor. Ultimately, a model of choice should be based on neural data, and a way to start on this path is to add neural activity to task factors[76] such as those considered in this paper. Ultimately, the neural activity should be superior and make any other factor irrelevant[7,86], because it is the sole determinant of the brain's choices.

Acknowledgments

I thank my lab colleagues Celian Bimbard, Pip Coen, Tim Sit, Valentin Schmutz, Max Shinn, and Flora Takacs for advice and help in obtaining data. I thank Kenneth D Harris for numerous insights and for suggesting the proof in Section S2, and Josh Gold for helpful suggestions. This work was funded by the Wellcome Trust (223144), by UKRI (EP/X022366/1), and by the NIH (U19NS123716). Earlier work was supported by the GlaxoSmithKline / Fight for Sight Chair in Visual Neuroscience.

License

This preprint bears a CC-BY-NC-ND license.

Methods

Most of the data were taken from the original papers using a manual data grabber (the MATLAB function *grabit.m*). The data thus obtained are somewhat noisy, and the interested reader should refer to the original papers for the definitive data. In a few cases (noted in the legends of the relevant figures), I accessed the original data.

I then fitted these data with the 2-choice logistic classification model using the MATLAB function *glmfit*. In the case of 3 choices, instead, I used the function *mnrfit*.

To describe sensory choices influenced by priors (Figure 2 **a-f**) I fit this equation:

$$p = \sigma(Vv + H(h - 0.5) + B) \qquad \text{Eq. 15}$$

where v is visual strength (coherence for the motion task, contrast for the position task), h is the probability of a Right stimulus (determined by the block), and V, H, B are parameters (visual gain, prior gain, and bias).

To describe sensory choices influenced by value in humans, monkeys, and mice (Figure 2**g-o**) I fit this equation:

$$p = \sigma(Vv + Rr + B) \qquad \text{Eq. 16}$$

where v is visual strength (position, coherence, or contrast, depending on the experiment), r is the relative reward offered for Left vs. Right choices (determined by the block) and V, R, B are parameters (visual gain, reward gain, and bias).

To describe audiovisual position choices in mice (Figure 3) and visuotactile orientation choices in rats (Supp Figure 2) I fit Eq. 9. To find the best values of the exponent, I explored a range of exponents between 0.1 and

2.0, and for each value in this range I fitted a logistic classification model applied to contrast raised to that exponent. I then selected the exponent that provided the best fits.

To describe audiovisual rate choices in humans (Figure 4**a**-**h**) I fit Eq. 11.

To describe vestibulo-visual heading choices in monkeys (Figure 4**i**-**p**) I fit the same equation but with an exponent q applied to both ϕ_1 and ϕ_2.

To illustrate the effects of stimulus history (Figure 5) I simulated the model in Eq. 12 using the parameters W = 0.8, S = 2, F = -1, B = 0.

To describe the effects of brain manipulations (Figure 6) I modified Eq. 14 to allow not only for different sensitivities on L vs R but also for possible effects of manipulations:

$$z = \sum_{i=1,2} \sum_{j\in\{L} (F_{ij} + \Delta_{ij}) s_{ij} + (B + \Delta_B) \qquad \text{Eq. 17}$$

where the variables s_{ij} are the stimuli, indexed by their modality ($i = 1$ for visual contrast raised to exponent q and $i = 2$ for auditory position), and by their position ($j = L$ or R). The parameters $F_{1L}, F_{2L}, F_{1R}, F_{2R}$ are the corresponding weights, and the five Δ parameters allow for the effects of the brain manipulation on the four weights F and on the bias B.

Supplementary Figures and Mathematical Appendices

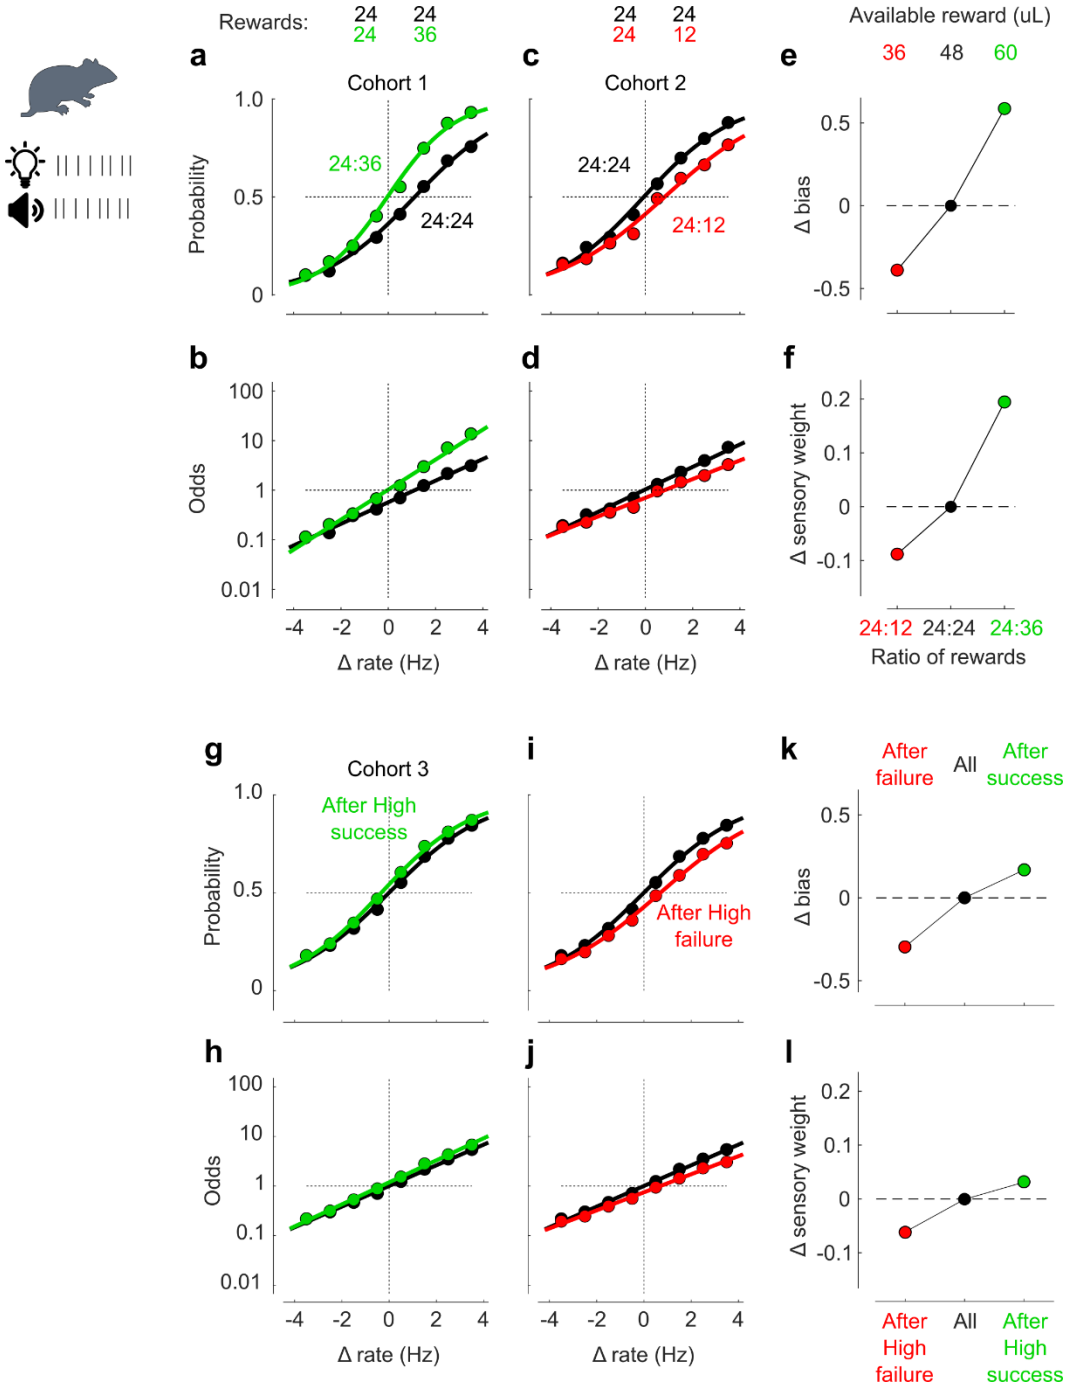

Supp Figure 1. ***Logistic classification describes the effects of unequal rewards and of recent history****. Related to Figs. 2 and 5.* ***a****. Probability of choosing "high" in a task where rats judge the rate of audiovisual stimuli, where the rewards for correct "high" choices are the same as correct "low" choices (24 µL, black) or higher (36 µL, green). Curves show fits of the logistic classification model where both the sensory weight W and the bias B are allowed to change when the rewards are unequal.* ***b****. Same data and curves, plotted in terms of the logarithmic odds of choosing "high".* ***c,d****. Same, for a cohort of rats that were run in conditions where the rewards for correct "high" choices are the same as correct "low" choices (24 µL, black) or lower (12 µL, green).* ***e****. As predicted by logistic classification under the hypothesis that observers seek to maximize reward, the bias parameter was affected in opposite directions by the two manipulations. The changes in bias predicted by the model (Eq. 7) would be log(12/24)=-0.69 and log(36/24)=0.41.* ***f****. In addition, the sensory weight increased when the ratio of rewards increased, perhaps because the total available reward also increased (abscissa). This may have caused an increase in attention to the task variable.* ***g-j****. Same as a-d for a cohort of rats that were run with equal rewards for "low" and "high" correct choices, distinguishing performance following a "high" choice that was correct (green) or incorrect (red).* ***k****. The bias term followed the predictions of logistic classification with history terms (Eq. 12) in a win-stay lose-switch strategy.* ***l****. The sensory weight was slightly higher after a success than after a failure, possibly because failures tended to occur during behavioral states where the rats were paying less attention to the stimulus. Data were obtained from Fig. 4 of Ref [26].*

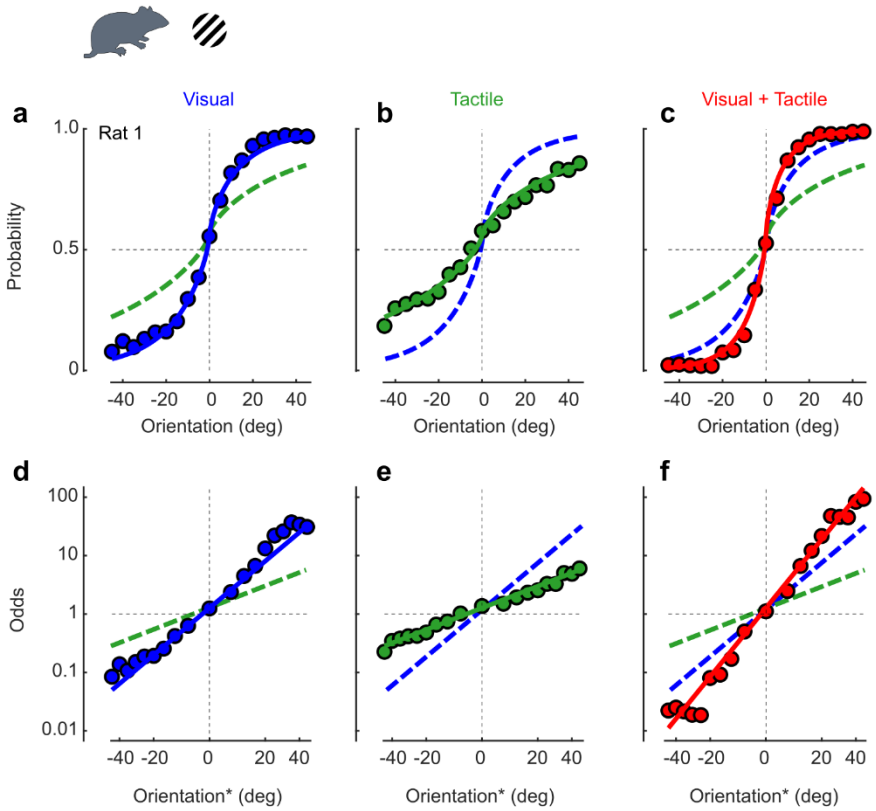

Supp Figure 2. ***Logistic classification describes visuotactile orientation choices in rats****. Related to Fig 3.* ***a****. Probability of Right choices for a rat performing the orientation discrimination task based on visual cues.* ***b****. Same, based on tactile cues.* ***c.*** *Same, based on stimuli that are both visual and tactile.* ***d-e****. Same data, plotted as odds of Right vs. Left in logarithmic scale. The quantity on the abscissa is compressed orientation (raised to an exponent q <1). Curves are from the logistic classification model. Data are from Rat 1 in Ref.[55].*

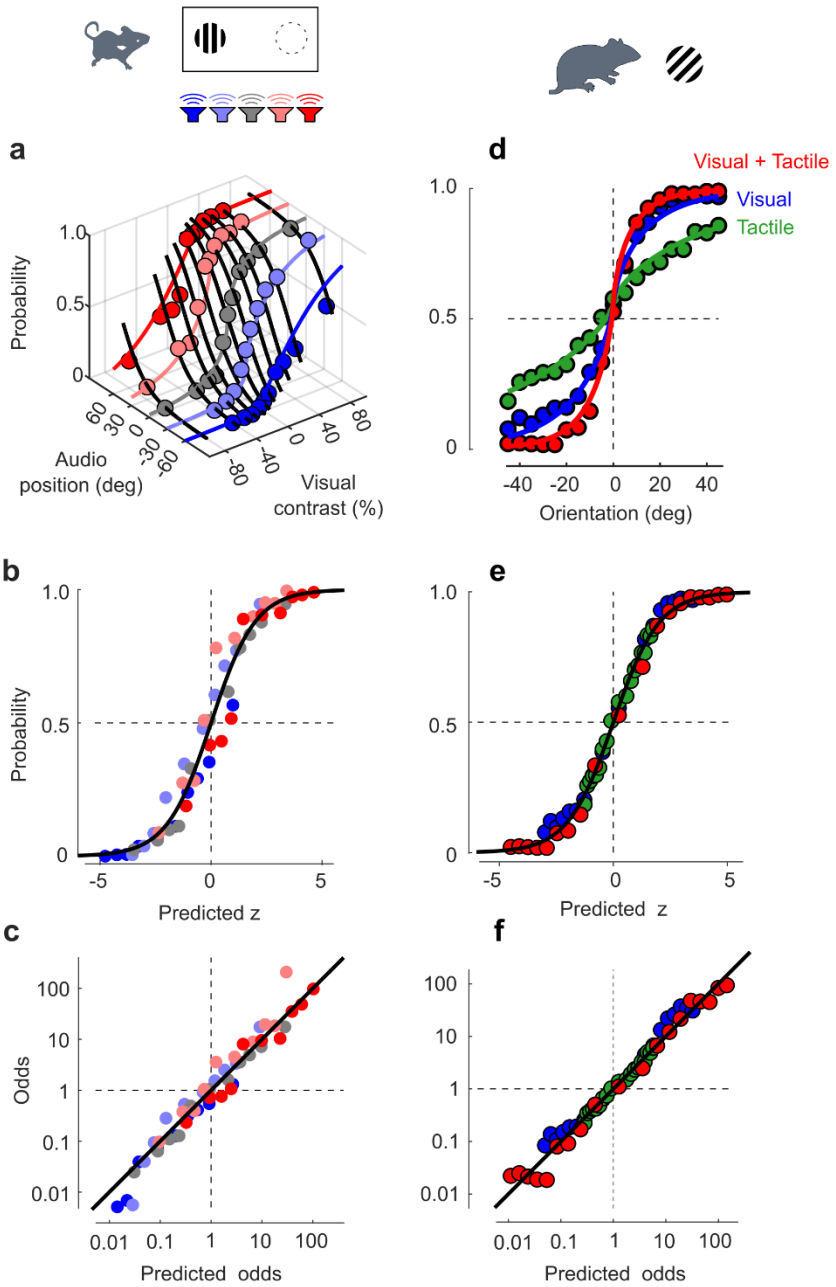

Supp Figure 3. ***Summarizing multiple conditions on a single logistic curve****. Related to Fig. 4.* ***a****. Probabilities of Right choices in an audiovisual task (replotted from Figure 3).* ***b****. The same data, plotted as a function of the decision variable z predicted by the model, plotted on top of a standard logistic curve.* ***c****. The same data, in log odds.* ***d****. Probabilities of Right choices in a visuotactile task (replotted from Supp Fig 2). Curves in a,b show fits of the logistic classification model.* ***e-f.*** *The same data plotted as in b,c.*

Section S1. Logistic classification and signal detection theory

We have seen that logistic classification can be obtained by taking the difference of two stochastic factors, and then making a deterministic choice by thresholding (Figure 7**b**). In this description, the utilities L, R in favor of Left and Right choices are each the sum of a deterministic value z_L, z_R and a random variable ϵ_L, ϵ_R (which varies across trials). The model subtracts the two and thresholds the result (Figure 9**a**). This description is mathematically equivalent to the one in the main text if the random variables are drawn from a Gumbel distribution (Figure 9**b**); in that case their difference is drawn from a logistic distribution (Figure 9**c**), and the model performs logistic classification[1,3] (Figure 9**d**).

In this framework, the weight W that appears in the logistic classification model (Eq. 1) is the inverse of the standard deviation of ϵ. For an ideal observer, this standard deviation is equal to the sensory noise, but for any other observer, it is larger. The weight W is thus bound by the sensory reliability (the inverse of the standard deviation). It can be lower, but not higher.

This description reveals that logistic classification (for two alternatives) is similar to signal detection theory[27,28] (SDT), which is the dominant formalism used to describe perceptual decisions. In SDT, the deterministic part of L and R is the "signal", and the random variable is "noise", which is thought to be sensory and related to limitations of sensory systems. In SDT, the random variable is commonly taken to be Gaussian, resulting in a psychometric curve that is a cumulative Gaussian function (Figure 9**e**).

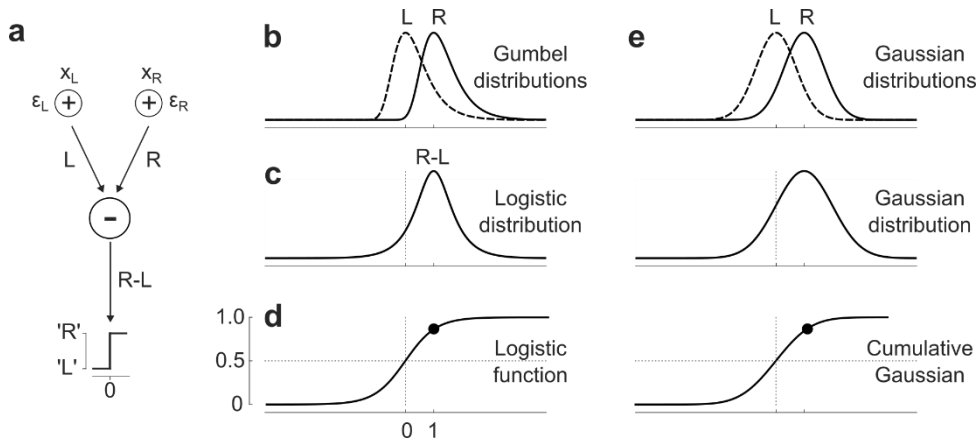

Figure 9. ***Logistic classification and signal detection theory. a****. An alternative description of the choice model, where the utilities L, R in favor of Left and Right choices are each the sum of a deterministic value x and a random variable ε. The model subtracts the two and thresholds the result to obtain a choice.* ***b****. Probability distributions for the utilities L and R in the case where $x_L = 0$ and $x_R = 1$, and the random variables ε are distributed according to the Gumbel distribution.* ***c****. The distribution of the difference R-L is logistic.* ***d****. After thresholding, the probability of choosing 'R' as a function of $x_R - x_L$ is a point on the logistic function.* ***e****. The same quantities, if the random variables ε are drawn from a Gaussian distribution. In that case the difference is Gaussian, and the probability of 'R' choice is a point on a cumulative Gaussian function. This is the formalism typically used in signal detection theory.*

When describing a psychometric curve determined by a single factor, the two frameworks (logistic and Gaussian) are hardly distinguishable[26,28,87-89]. For example, consider the probabilities of choosing 'R' in a two-alternative task (Figure 10**a**). These were generated by a logistic classification model $p = \sigma(Wx + B)$, where x is the stimulus and W, B are parameters (sensory weight and bias). The same probabilities can be fitted extremely well by the cumulative Gaussian $\Phi((x - \mu)/\sigma)$, where the two parameters σ, μ determine a "threshold" and a "bias" (Figure 10**b**).

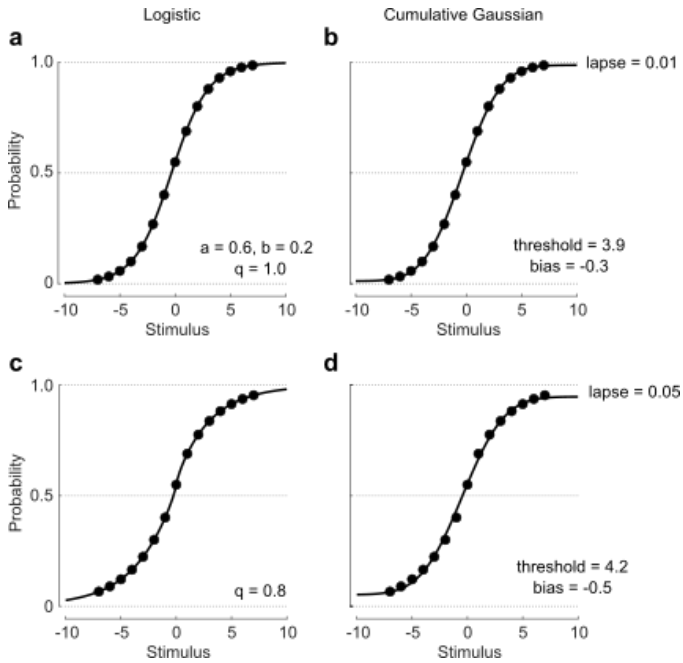

Figure 10. ***Logistic classification vs. fitting a cumulative gaussian. a****. Probabilities simulated by a two-alternative logistic classification model (data points and curve).* ***b****. The same data points, fitted by a cumulative Gaussian with a lapse rate (curve).* ***c****. Probabilities simulated by a logistic classification model with a slight compression of the stimulus axis (exponent q=0.8).* ***d.*** *The same data points, fitted with the cumulative Gaussian with a larger lapse rate (curve).*

When fitting actual data, however, the probabilities don't always reach zero or one, even for the easiest stimuli. To deal with this eventuality, when fitting a cumulative Gaussian one typically adds an extra parameter λ (or two) to account for a "lapse rate" λ (Figure 10**d**), to obtain

$$p = \lambda + (1 - \lambda)\Phi\left(\frac{x - \mu}{\sigma}\right)$$

With logistic classification, it is generally not necessary to add such a lapse rate. In many cases, one can use the logistic function as is, because it tends to zero or one more gradually than the cumulative Gaussian (the logistic distribution has long tails). The probabilities are then below 1 or above 0 simply because the observer does not give enough weight W to the sensory input x. In other cases, as shown in multiple occasions in the main text, the observer may process stimulus strength in a nonlinear fashion. For instance, the sensory input x may be given by the measure s of stimulus strength used in the laboratory, raised to the power $q < 1$ to obtain a compressive nonlinearity (Figure 10**c**). In other cases, one may need a nonlinearity that is even more compressive[12,24].

More generally, logistic classification outstrips signal detection theory (SDT) when it comes to representing more general choice environments. For instance, as I have shown in the main text, logistic classification makes it trivial to add predictors, both sensory and non-sensory, and places all these predictors on an equal footing; it has a simple linear interpretation (in terms of log odds); it makes it easy to describe the effects of brain manipulations; and it is readily extended to an arbitrary number of alternatives. Moreover, logistic classification incorporates the Bayesian assumption, without having to postulate that an observer keeps track separately of stimulus estimate and confidence (the standard deviation of the Gumbel distributions is always 1). SDT lacks some of these advantages. It can incorporate non-sensory predictors but gives them a special role[27,29] in setting the bias μ. Moreover, it makes it unwieldy to have more than one or two predictors or more than two choice alternatives.

Section S2: Logistic classification and drift diffusion

The drift diffusion model is an established model for decision making[2,30-32] (Figure 7**c**). In the model the decision variable controls the drift rate of a random walk. When the walk hits one of the two boundaries, the corresponding decision is made.

In the general drift diffusion model, the random walk can jump at each step by an arbitrary amount, determined by a probability distribution[2,30-32,77]. In the simpler drift diffusion model described here[78], at each moment in time the decision variable can only go up or down by 1 (Figure 7**c**).

Here I derive mathematical expressions for the probability of hitting the upper bound, and for the average time required to hit a bound. I focus on the simpler drift diffusion model (for the more general model, the derivations are more involved[31,77]).

Probability of hitting the upper bound

If the random walk starts from the middle between the two bounds, then the probability that it will hit one of the bounds (say, the upper one) is given by the logistic function. Specifically, here I show that given a decision variable z (e.g. given by Eq. 1), if the probability of going up at each step is set to $P = \sigma(z/N)$, where N is the distance from the starting point to the bounds, the random walk reaches the upper bound with probability $p = \sigma(z)$.

The first step to prove this is to recognize that the simplified drift diffusion model is identical to the well-known "gambler's ruin" problem, where players 1 and 2 have N_1 and N_2 coins, and repeatedly bet one coin, in bets where player 1 wins with probability P. As shown in a textbook (Ref. [90], chapter XIV), the probability that player 1 will win (i.e. that player 2 will go bankrupt) is

$$p = \frac{1 - r^{-N_1}}{1 - r^{-(N_1+N_2)}}$$

where $r = P/(1-P)$ and we have assumed that $r \neq 1$ (the drift is preferentially in one direction).

If the two players start with equal amount of money ($N_1 = N_2 = N$), then player 1 will win with probability

$$p = \frac{1}{1 + r^{-N}} = \frac{1}{1 + e^{-N\log(r)}} = \sigma(N\log(r))$$

where σ is the logistic function. Now, if $P = \sigma(z/N)$, then

$$z = N\log\frac{P}{1-P} = N\log(r)$$

and the above expression can be rewritten as

$$p = \sigma(z).$$

Expected number of steps before hitting a bound

Here I show that if the probability of going up at each step is set to $P = \sigma(z/N)$, where $N \gg z$ is the distance from the starting point to the bounds, the drift diffusion model reaches a bound in average in d steps, where

$$d \approx 2N^2\frac{\sigma(z) - \sigma(-z)}{z}$$

The proof starts from the same textbook (Ref. [90], chapter XIV), which derives the expected number of steps required for hitting a bound:

$$d = \frac{N}{1-2P}\left(1 - 2\frac{1}{1+r^{-N}}\right) = N\frac{1-2p}{1-2P}$$

Substituting $p = \sigma(z)$ and $P = \sigma(z/N)$ we obtain:

$$d = N\frac{1-2\sigma(z)}{1-2\sigma(z/N)} = N\frac{\sigma(z)-\sigma(-z)}{\sigma(z/N)-\sigma(-z/N)} \approx 2N^2\frac{\sigma(z)-\sigma(-z)}{z}$$

Where in the second equality I used the fact that $\sigma(x) = 1 - \sigma(-x)$, and then I used the Taylor series approximation for $\sigma(x)$ when $x \approx 0$ (because $z \ll N$) which is $\sigma(x) \approx 1/2 + x/4$.

Note that the case $z = 0$ is special: the expected delay becomes $d = N^2$.

Fitting psychometric and chronometric data

These results constitute a fundamental strength of the drift diffusion model[2,10,11,32,77]: it can account not only for how stimulus strength affects the probability of choices (the psychometric data) but also for the time of those choices (the "chronometric" data).

To see this in practice, let's return to data measured in the random dots task[11], where choices depend on stimulus coherence x, and let's consider both the psychometric data (Figure 11**a**) and the chronometric data (Figure 11**b**). The drift diffusion model can fit the psychometric data because it provides logistic classification:

$$z = Wx + B$$

$$p = \sigma(z)$$

Moreover, it can fit chronometric data with an additional equation

$$d = R + 2S\frac{\sigma(z)-\sigma(-z)}{z}$$

where R, S are parameters that describe a fixed "nondecision time" and the maximal additional time that can be added to the decision when there is no evidence. The parameters W, B affect both the psychometric and chronometric data. They can be fitted simultaneously to both sets of data[11].

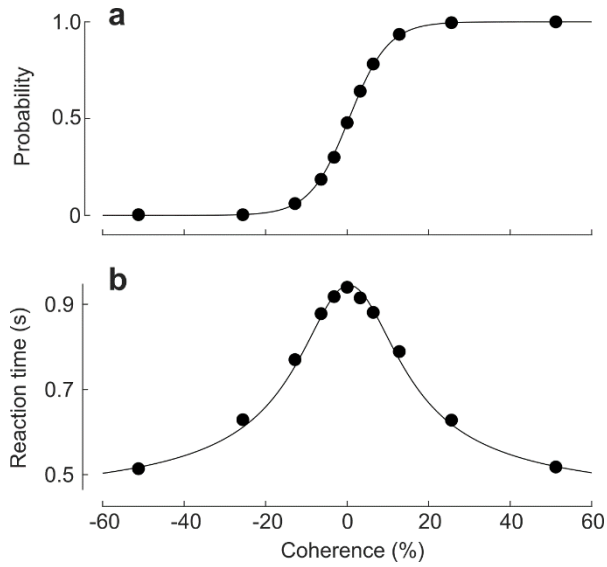

Figure 11. ***Psychometric and chronometric data fitted with the drift diffusion model. a.*** *Probability of rightward choices in the random dots task, as a function of stimulus coherence (same as Figure 1d-f). I used these data to obtain the parameters* $W = 0.209, B = -0.097$. ***b****. The corresponding reaction times. I used these data to obtain the parameters* $R = 0.420, S = 0.525$. *Data were grabbed from Ref. [11]. Curves show predictions of the drift diffusion model with the above parameters.*

Section S3: The optimal bias

Here I derive the ideal bias for an observer that performs logistic classification given two potentially different rewards for correct left vs. right choices V_L, V_R, and two potentially different prior probabilities of stimuli being on the left or right, P_L, P_R.

If an observer performs logistic classification based on evidence x with weight a, the probability that they will choose Right is $p = P('R'|x) = \sigma(ax + b)$.

What is the bias b that maximizes expected value? The answer depends on whether the evidence x is informative or not.

Uninformative evidence

If the evidence is uninformative ($ax = 0$), the best strategy is to always choose the Right option ($b = \infty$) if $V_R P_R > V_L P_L$ and always choose the Left option ($b = -\infty$) otherwise.

This is easy to prove. The expected value of a choice is the product of the probabilities of each answer being correct (i.e. rewarded) times the corresponding reward:

$$V = V_L P(L \,\&\, 'L') + V_R P(R \,\&\, 'R') \qquad \text{Eq. 18}$$

Because the choices are not guided by evidence, they are independent of the process that determines the rewarded side, the above equation can be written as

$$V = V_L P(L) P('L') + V_R P(R) P('R') = V_L P_L (1 - p) + V_R P_R p$$

So V lies on a line that goes from $V_L P_L$ to $V_R P_R$ as p varies between 0 and 1. If the former value is higher than the latter, the maximum is achieved at $p = 0$, which means $b = -\infty$. Otherwise, the maximum is achieved at $p = 1$, which means $b = \infty$.

Informative evidence

If instead the evidence is informative, the optimal bias b is a finite number. Specifically, if x is distributed uniformly (and a is not zero) the bias b that maximizes expected value is

$$b = \log\left(\frac{V_R}{V_L}\right) + \log\left(\frac{P_R}{P_L}\right) \qquad \text{Eq. 19}$$

This result is proven below. My derivation is inspired by an analogous result obtained with signal detection theory using cumulative gaussians[45]. A similar expression is also obtained in the context of signal detection theory[27,29].

$$V = V_L P('L' \,\&\, x < 0) + V_R P('R' \,\&\, x > 0)$$

The evidence x is a random variable with prior probability $p(x)$. We can thus write

$$V = V_L \int_{-\infty}^{0} P('L', x)\, dx + V_R \int_{0}^{\infty} P('R', x)\, dx =$$

$$= V_L \int_{-\infty}^{0} P('L'|x) p(x) dx + V_R \int_{0}^{\infty} P('R'|x) p(x) dx$$

We assume that the prior probability is different between Left and Right, but otherwise flat: $p(x) = P_L/T, x \in [-T, 0]$ and $p(x) = P_R/T, x \in [0, T]$. Later we will set $T \to \infty$. Using the fact that $P('R'|x) = \sigma(ax + b)$ and $P('L'|x) = 1 - \sigma(ax + b) = \sigma(-ax - b)$, we obtain:

$$V = \frac{P_L V_L}{T} \int_{-T}^{0} \sigma(-ax - b)dx + \frac{P_R V_R}{T} \int_{0}^{T} \sigma(ax + b)dx$$

To maximize V we need to find b that such that $dV/db = 0$. We can bring the derivative inside the integrals:

$$\frac{dV}{db} = \frac{P_L V_L}{T} \int_{-T}^{0} \frac{d}{db}\sigma(-ax - b)dx + \frac{P_R V_R}{T} \int_{0}^{T} \frac{d}{db}\sigma(ax + b)dx = 0$$

A nice property of the logistic function is that its derivative has a simple form: $d\sigma(x)/dx = \sigma(x)\sigma(-x)$. Using this property and rearranging the terms, we obtain:

$$\frac{P_R V_R}{P_L V_L} = \frac{\int_{-T}^{0} \sigma(ax + b)\sigma(-ax - b)dx}{\int_{0}^{T} \sigma(ax + b)\sigma(-ax - b)dx}$$

With a change of variable $z = ax + b$ the Eq. becomes

$$\frac{P_R V_R}{P_L V_L} = \frac{1/a \int_{-aT+b}^{b} \sigma(z)\sigma(-z)dz}{1/a \int_{b}^{aT+b} \sigma(z)\sigma(-z)dz} = \frac{\sigma(b) - \sigma(-aT + b)}{\sigma(aT + b) - \sigma(b)} \xrightarrow{T\to\infty} \frac{\sigma(b)}{1 - \sigma(b)}$$

Where in the second equality we have solved the integrals by using the expression for the derivative of σ in reverse, and in taking the limit we have used the fact that $\sigma(\infty) = 1$ and $\sigma(-\infty) = 0$. If we now take the logarithm of both sides, and recognize that $\sigma^{-1}(p) = \log\ p/(1 - p)$, we obtain:

$$log \frac{P_R V_R}{P_L V_L} = log \frac{\sigma(b)}{1 - \sigma(b)} = \sigma^{-1}\big(\sigma(b)\big) = b$$

which proves the result in Eq. 19.

Section S4: Probability matching implies logistic classification

We have seen that if an observer performs logistic classification (Eq. 5) they should set their bias to the log prior odds (Eq. 6), which leads to probability matching (when the evidence is zero).

Here I show that this argument can be inverted: the only way to achieve probability matching (for any strength of evidence) is to perform logistic classification (Eq. 5), with bias set to the log prior odds (Eq. 6).

Consider a task where a stimulus is on the Left or the Right with probabilities P_L and P_R, and the observer infers the stimulus position from a sensory input s.

Assume that the observer aims to perform "probability matching", i.e., to say that the stimulus is on the Right with probability $p = P(R|s)$, the probability that the stimulus s is indeed on the Right.

By Bayes' rule, the odds that the observer chooses Right are the product of two terms: one that depends on the sensory stimulus (the Bayes factor) and one that does not (the prior odds):

$$\frac{p}{1-p} = \frac{P(R|s)}{P(L|s)} = \frac{P(s|R)}{P(s|L)}\,\frac{P_R}{P_L}$$

Taking the logarithm yields a sum of two terms:

$$z = \log\frac{p}{1-p} = \log\frac{P(s|R)}{P(s|L)} + \log\frac{P_R}{P_L} \qquad \text{Eq. 20}$$

In other words, the decision variable is the sum of a function of the stimulus and a constant bias,

$$z = f(s) + B \qquad \text{Eq. 21}$$

where $f(s)$ is a (possibly nonlinear, and possibly complicated) function of the stimulus, and $B = \log(P_R/P_L)$.

We have thus obtained the equation for logistic classification, with a suitable choice of evidence $x = f(s)$ (Eq. 5), and the bias set to the log prior odds (Eq. 6).

A more general expression can be obtained if there two sensory stimuli s_1, s_2 of different modalities (e.g., vision and audition), and these stimuli are statistically independent[23]. In that case, using both Bayes rule and the definition of independence, $P(s_1, s_2|R) = P(s_1|R)P(s_2|R)$, one obtains

$$\frac{p}{1-p} = \frac{P(R|s_1, s_2)}{P(L|s_1, s_2)} = \frac{P(s_1|R)}{P(s_1|L)}\,\frac{P(s_2|R)}{P(s_2|L)}\,\frac{P_R}{P_L}$$

In other words,

$$z = \log\left(\frac{p}{1-p}\right) = f_1(s_1) + f_2(s_2) + B$$

thus proving Eq. 8.

References

1. McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In Frontiers in Econometrics, P. Zarembka, ed. (Academic Press), pp. 105-142.

2. Gold, J.I., and Shadlen, M.N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. Neuron *36*, 299-308. 10.1016/s0896-6273(02)00971-6.

3. Train, K.E. (2009). Discrete Choice Methods with Simulation, 2nd Edition (Cambridge University Press). https://doi.org/10.1017/CBO9780511805271.

4. McFadden, D. (2001). Economic choices. American Economic Review *91*, 351-378. DOI 10.1257/aer.91.3.351.

5. Goodfellow, I., Bengio, Y., and Courville, A. (2016). Deep Learning (MIT Press).

6. Cramer, J.S. (2002). The origins of logistic regression.

7. Padoa-Schioppa, C. (2022). Logistic analysis of choice data: A primer. Neuron. 10.1016/j.neuron.2022.03.002.

8. Lau, B., and Glimcher, P.W. (2005). Dynamic response-by-response models of matching behavior in rhesus monkeys. J Exp Anal Behav *84*, 555-579.

9. Corrado, G.S., Sugrue, L.P., Seung, H.S., and Newsome, W.T. (2005). Linear-Nonlinear-Poisson models of primate choice dynamics. J Exp Anal Behav *84*, 581-617.

10. Gold, J.I., and Shadlen, M.N. (2001). Neural computations that underlie decisions about sensory stimuli. Trends Cogn Sci *5*, 10-16.

11. Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci *25*, 10420-10436. 10.1523/JNEUROSCI.4684-04.2005.

12. Busse, L., Ayaz, A., Dhruv, N.T., Katzner, S., Saleem, A.B., Schölvinck, M.L., Zaharia, A.D., and Carandini, M. (2011). The detection of visual contrast in the behaving mouse. J Neurosci *31*, 11351-11361. 10.1523/JNEUROSCI.6689-10.2011.

13. Padoa-Schioppa, C. (2013). Neuronal origins of choice variability in economic decisions. Neuron *80*, 1322-1336. 10.1016/j.neuron.2013.09.013.

14. Fründ, I., Wichmann, F.A., and Macke, J.H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. J Vis *14*. 10.1167/14.7.9.

15. Akaishi, R., Umeda, K., Nagase, A., and Sakai, K. (2014). Autonomous mechanism of internal choice estimate underlies decision inertia. Neuron *81*, 195-206. 10.1016/j.neuron.2013.10.018.

16. Abrahamyan, A., Silva, L.L., Dakin, S.C., Carandini, M., and Gardner, J.L. (2016). Adaptable history biases in human perceptual decisions. Proc Natl Acad Sci U S A *113*, E3548-3557. 10.1073/pnas.1518786113.

17. St. John-Saaltink, E., Kok, P., Lau, H., and de Lange, F.P. (2016). Serial Dependence in Perceptual Decisions Is Reflected in Activity Patterns in Primary Visual Cortex. J Neurosci.

18. Urai, A.E., Braun, A., and Donner, T.H. (2017). Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. Nat Commun *8*, 14637. 10.1038/ncomms14637.

19. Hwang, E.J., Dahlen, J.E., Mukundan, M., and Komiyama, T. (2017). History-based action selection bias in posterior parietal cortex. Nat Commun *8*, 1242. 10.1038/s41467-017-01356-z.

20. Akrami, A., Kopec, C.D., Diamond, M.E., and Brody, C.D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. Nature *554*, 368-372. 10.1038/nature25510.

21. Braun, A., Urai, A.E., and Donner, T.H. (2018). Adaptive History Biases Result from Confidence-Weighted Accumulation of past Choices. J Neurosci *38*, 2418-2429. 10.1523/JNEUROSCI.2189-17.2017.

22. International Brain Laboratory, Aguillon-Rodriguez, V., Angelaki, D., Bayer, H., Bonacchi, N., Carandini, M., Cazettes, F., Chapuis, G., Churchland, A.K., Dan, Y., et al. (2021). Standardized and reproducible measurement of decision-making in mice. eLife *10*. 10.7554/eLife.63711.

23. Coen, P., Sit, T.P.H., Wells, M.J., Carandini, M., and Harris, K.D. (2023). Mouse frontal cortex mediates additive multisensory decisions. Neuron *111*, 2432-2447 e2413. 10.1016/j.neuron.2023.05.008.

24. Burgess, C.P., Lak, A., Steinmetz, N.A., Zatka-Haas, P., Bai Reddy, C., Jacobs, E.A.K., Linden, J.F., Paton, J.J., Ranson, A., Schröder, S., et al. (2017). High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. Cell Rep *20*, 2513-2524. 10.1016/j.celrep.2017.08.047.

25. Zatka-Haas, P., Steinmetz, N.A., Carandini, M., and Harris, K.D. (2021). Sensory coding and the causal impact of mouse cortex in a visual decision. eLife *10*. 10.7554/eLife.63163.

26. Pisupati, S., Chartarifsky-Lynn, L., Khanal, A., and Churchland, A.K. (2021). Lapses in perceptual decisions reflect exploration. eLife *10*. 10.7554/eLife.55490.

27. Green, D.M., and Swets, J.A. (1966). Signal detection theory and psychophysics (Wiley).

28. McNicol, D. (2005). A primer of signal detection theory (Psychology Press).

29. Swets, J., Tanner, W.P., Jr., and Birdsall, T.G. (1961). Decision processes in perception. Psychol Rev *68*, 301-340.

30. Ratcliff, R., Smith, P.L., Brown, S.D., and McKoon, G. (2016). Diffusion Decision Model: Current Issues and History. Trends Cogn Sci *20*, 260-281. 10.1016/j.tics.2016.01.007.

31. Link, S.W. (1992). The wave theory of difference and similarity (Lawrence Erlbaum Associates).

32. Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput *20*, 873-922. 10.1162/neco.2008.12-06-420.

33. Stevens, S.S. (1960). The Psychophysics of Sensory Function. American Scientist *48*, 226-253.

34. Stevens, S.S. (1970). Neural events and the psychophysical law. Science *170*, 1043-1050. 10.1126/science.170.3962.1043.

35. Hanks, T.D., Mazurek, M.E., Kiani, R., Hopp, E., and Shadlen, M.N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. J Neurosci *31*, 6339-6352. 10.1523/JNEUROSCI.5613-10.2011.

36. Brunswik, E. (1939). Probability as a determiner of rat behavior. J Exp Psych *25*, 175.

37. Grant, D.A., Hake, H.W., and Hornseth, J.P. (1951). Acquisition and extinction of a verbal conditioned response with differing percentages of reinforcement. J Exp Psych *42*, 1.

38. Hake, H.W., and Hyman, R. (1953). Perception of the statistical structure of a random series of binary symbols. J Exp Psych *45*, 64.

39. Estes, W.K., and Straughan, J. (1954). Analysis of a verbal conditioning situation in terms of statistical learning theory. J Exp Psych *47*, 225.

40. Goodnow, J.J. (1955). Determinants of choice-distribution in two-choice situations. Am J Psychol *68*, 106-116.

41. Herrnstein, R.J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. J Exp Anal Behav *4*, 267-272. 10.1901/jeab.1961.4-267.

42. Vulkan, N. (2000). An economist's perspective on probability matching. J Econ Surv *14*, 101-118.

43. Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. Science *304*, 1782-1787. 10.1126/science.1094765.

44. Baum, W.M. (1974). On two types of deviation from the matching law: bias and undermatching. J Exp Anal Behav *22*, 231-242. 10.1901/jeab.1974.22-231.

45. Whiteley, L., and Sahani, M. (2008). Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. J Vis *8*, 2 1-15. 10.1167/8.3.2.

46. Feng, S., Holmes, P., Rorie, A., and Newsome, W.T. (2009). Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? PLoS Comput Biol *5*, e1000284. 10.1371/journal.pcbi.1000284.

47. Lak, A., Okun, M., Moss, M.M., Gurnani, H., Farrell, K., Wells, M.J., Reddy, C.B., Kepecs, A., Harris, K.D., and Carandini, M. (2019). Dopaminergic and Prefrontal Basis of Learning from Sensory Confidence and Reward Value. Neuron *105*, 700-711 e706. 10.1016/j.neuron.2019.11.018.

48. Fan, Y., Gold, J.I., and Ding, L. (2018). Ongoing, rational calibration of reward-driven perceptual biases. eLife *7*. 10.7554/eLife.36018.

49. Constantinople, C.M., Piet, A.T., and Brody, C.D. (2019). An Analysis of Decision under Risk in Rats. Curr Biol *29*, 2066-2074 e2065. 10.1016/j.cub.2019.05.013.

50. Cameron, E.L., Tai, J.C., and Carrasco, M. (2002). Covert attention affects the psychometric function of contrast sensitivity. Vis Res *42*, 949-967. 10.1016/s0042-6989(02)00039-1.

51. Roy, N.A., Bak, J.H., International Brain, L., Akrami, A., Brody, C.D., and Pillow, J.W. (2021). Extracting the dynamics of behavior in sensory decision-making experiments. Neuron *109*, 597-610 e596. 10.1016/j.neuron.2020.12.004.

52. Ashwood, Z.C., Roy, N.A., Stone, I.R., International Brain, L., Urai, A.E., Churchland, A.K., Pouget, A., and Pillow, J.W. (2022). Mice alternate between discrete strategies during perceptual decision-making. Nat Neurosci *25*, 201-212. 10.1038/s41593-021-01007-z.

53. Roy, N.A., Bak, J.H., Akrami, A., Brody, C.D., and Pillow, J.W. (2018). Efficient inference for time-varying behavior during learning. Adv Neural Information Processing Systems *31*.

54. Battaglia, P.W., Jacobs, R.A., and Aslin, R.N. (2003). Bayesian integration of visual and auditory signals for spatial localization. J Opt Soc Am A *20*, 1391-1397. 10.1364/josaa.20.001391.

55. Nikbakht, N., Tafreshiha, A., Zoccolan, D., and Diamond, M.E. (2018). Supralinear and Supramodal Integration of Visual and Tactile Signals in Rats: Psychophysics and Neuronal Mechanisms. Neuron *97*, 626-639 e628. 10.1016/j.neuron.2018.01.003.

56. Alais, D., and Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. Curr Biol *14*, 257-262. 10.1016/j.cub.2004.01.029.

57. Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. Nature *415*, 429-433. 10.1038/415429a.

58. Gu, Y., Angelaki, D.E., and Deangelis, G.C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. Nat Neurosci *11*, 1201-1210. 10.1038/nn.2191.

59. Fetsch, C.R., Turner, A.H., DeAngelis, G.C., and Angelaki, D.E. (2009). Dynamic reweighting of visual and vestibular cues during self-motion perception. J Neurosci *29*, 15601-15612. 10.1523/JNEUROSCI.2574-09.2009.

60. Raposo, D., Sheppard, J.P., Schrater, P.R., and Churchland, A.K. (2012). Multisensory decision-making in rats and humans. J Neurosci *32*, 3726-3735. 10.1523/JNEUROSCI.4998-11.2012.

61. Sheppard, J.P., Raposo, D., and Churchland, A.K. (2013). Dynamic weighting of multisensory stimuli shapes decision-making in rats and humans. J Vis *13*. 10.1167/13.6.4.

62. Drugowitsch, J., DeAngelis, G.C., Klier, E.M., Angelaki, D.E., and Pouget, A. (2014). Optimal multisensory decision-making in a reaction-time task. eLife *3*. 10.7554/eLife.03005.

63. Busse, L., Wade, A.R., and Carandini, M. (2009). Representation of concurrent stimuli by population activity in visual cortex. Neuron *64*, 931-942. 10.1016/j.neuron.2009.11.004.

64. Carandini, M., Heeger, D.J., and Movshon, J.A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. J Neurosci *17*, 8621-8644. 10.1523/JNEUROSCI.17-21-08621.1997.

65. Gardner, J.L. (2019). Optimality and heuristics in perceptual neuroscience. Nat Neurosci *22*, 514-523. 10.1038/s41593-019-0340-4.

66. Fernberger, S.W. (1920). Interdependence of judgments within the series for the method of constant stimuli. J Exp Psych *3*, 126.

67. Leopold, D.A., Wilke, M., Maier, A., and Logothetis, N.K. (2002). Stable perception of visually ambiguous patterns. Nat Neurosci *5*, 605-609. 10.1038/nn0602-851.

68. Fischer, J., and Whitney, D. (2014). Serial dependence in visual perception. Nat Neurosci *17*, 738-743. 10.1038/nn.3689.

69. Gold, J.I., Law, C.T., Connolly, P., and Bennur, S. (2008). The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. J Neurophysiol *100*, 2653-2668. 10.1152/jn.90629.2008.

70. Lueckmann, J.M., Macke, J.H., and Nienborg, H. (2018). Can Serial Dependencies in Choices and Neural Activity Explain Choice Probabilities? J Neurosci *38*, 3495-3506. 10.1523/JNEUROSCI.2225-17.2018.

71. Lak, A., Hueske, E., Hirokawa, J., Masset, P., Ott, T., Urai, A.E., Donner, T.H., Carandini, M., Tonegawa, S., Uchida, N., and Kepecs, A. (2020). Reinforcement biases subsequent perceptual decisions when confidence is low, a widespread behavioral phenomenon. eLife *9*. 10.7554/eLife.49834.

72. Hermoso-Mendizabal, A., Hyafil, A., Rueda-Orozco, P.E., Jaramillo, S., Robbe, D., and de la Rocha, J. (2020). Response outcomes gate the impact of expectations on perceptual decisions. Nat Commun *11*, 1057. 10.1038/s41467-020-14824-w.

73. Duque, D., and de la Rocha, J. (2022). Isolating perceptual biases caused by trial history during auditory categorization. bioRxiv, 2022.2001.2017.476581. 10.1101/2022.01.17.476581.

74. Tsunada, J., Cohen, Y., and Gold, J.I. (2019). Post-decision processing in primate prefrontal cortex influences subsequent choices on an auditory decision-making task. eLife *8*. 10.7554/eLife.46770.

75. Green, D.M. (1964). Consistency of Auditory Detection Judgments. Psychol Rev *71*, 392-407. 10.1037/h0044520.

76. Seidemann, E. (1998). Neuronal mechanisms mediating conversion of visual signals into perceptual decisions in a direction discrimination task. PhD (Stanford University).

77. Shadlen, M.N., Hanks, T.D., Churchland, A.K., Kiani, R., and Yang, T. (2006). The Speed and Accuracy of a Simple Perceptual Decision: A Mathematical Primer. In Bayesian Brain: Probabilistic Approaches to Neural Coding, K. Doya, S. Ishii, A. Pouget, and R.P.N. Rao, eds.

78. Ratcliff, R., and McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. Psychol Rev *104*, 319-343. 10.1037/0033-295x.104.2.319.

79. Urai, A.E., de Gee, J.W., Tsetsos, K., and Donner, T.H. (2019). Choice history biases subsequent evidence accumulation. eLife *8*. 10.7554/eLife.46331.

80. Shinn, M., Lam, N.H., and Murray, J.D. (2020). A flexible framework for simulating and fitting generalized drift-diffusion models. eLife *9*. 10.7554/eLife.56938.

81. Sridharan, D., Steinmetz, N.A., Moore, T., and Knudsen, E.I. (2014). Distinguishing bias from sensitivity effects in multialternative detection tasks. J Vis *14*. 10.1167/14.9.16.

82. DeCarlo, L.T. (2012). On a signal detection approach to -alternative forced choice with bias, with maximum likelihood and Bayesian approaches to estimation. J Math Psychol *56*, 196-207. 10.1016/j.jmp.2012.02.004.

83. Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220-223. 10.1038/nature14066.

84. Lorteije, J.A.M., Zylberberg, A., Ouellette, B.G., De Zeeuw, C.I., Sigman, M., and Roelfsema, P.R. (2015). The Formation of Hierarchical Decisions in the Visual Cortex. Neuron *87*, 1344-1356. 10.1016/j.neuron.2015.08.015.

85. Stringer, C., Michaelos, M., Tsyboulski, D., Lindo, S.E., and Pachitariu, M. (2021). High-precision coding in visual cortex. Cell. 10.1016/j.cell.2021.03.042.

86. Peixoto, D., Verhein, J.R., Kiani, R., Kao, J.C., Nuyujukian, P., Chandrasekaran, C., Brown, J., Fong, S., Ryu, S.I., Shenoy, K.V., and Newsome, W.T. (2021). Decoding and perturbing decision states in real time. Nature. 10.1038/s41586-020-03181-9.

87. Ingleby, J. (1973). The separation of bias and sensitivity in multiple-choice tasks. Perception *2*, 295-305.

88. Treisman, M., and Faulkner, A. (1985). On the Choice between Choice Theory and Signal-Detection Theory. Quarterly Journal of Experimental Psychology Section a-Human Experimental Psychology *37*, 387-405. Doi 10.1080/14640748508400941.

89. DeCarlo, L.T. (1998). Signal detection theory and generalized linear models. Psychol Methods *3*, 186-205. Doi 10.1037/1082-989x.3.2.186.

90. Feller, W. (1968). An Introduction to Probability Theory and its Applications, 3rd Edition (John Wiley & Sons).