Idiosyncratic choice bias in decision tasks naturally emerges from intrinsic stochasticity in neuronal network dynamics

Lior Lebovich^{1*}, Ran Darshan^{1,2}, Yoni Lavi^{3,4}, David Hansel^{5†} and Yonatan Loewenstein^{1,3,6†}

¹*The Edmond and Lily Safra Center for Brain Sciences, The Hebrew University, Jerusalem, Israel*

²*Janelia Farm Research Campus, Ashburn, VA, USA*

³*The Alexander Silberman Institute of Life Sciences, The Hebrew University, Jerusalem, Israel*

⁴*Code Institute, Dublin, Ireland*

⁵*Center for Integrative and Cognitive Neuroscience, Learning and Memory Lab; CNRS-UMR8002, 45 Rue des Saints Pères, 75270, Paris France*

⁶*Dept. of Cognitive Sciences and The Federmann Center for the Study of Rationality, The Hebrew University, Jerusalem, Israel*

*Correspondence should be addressed to Lior Lebovich (lebovich.lior@gmail.com).

†These authors contributed equally: David Hansel and Yonatan Loewenstein.

Idiosyncratic tendency to choose one alternative over others in the absence of an identified reason is a common observation in two-alternative forced-choice experiments. It is tempting to account for it as resulting from the (unknown) participant-specific history and thus treat it as a measurement noise. Here we quantify idiosyncratic choice biases in a perceptual discrimination task and a motor task. We report substantial and significant biases in both cases that cannot be accounted for by the experimental context. Then, we present theoretical evidence that even in idealized experiments, in which the settings are symmetric, idiosyncratic choice bias is expected to emerge from the dynamics of competing neuronal networks. We thus argue that idiosyncratic choice bias reflects the microscopic dynamics of choice and therefore is virtually inevitable in any comparison or decision task.

Decision making is the cognitive process of choosing an action among a set of alternatives. Decision making is often studied in experiments, composed of trials, each associated with a single decision. While a decision in a trial is primarily determined by the relevant features of the alternatives in that trial, biases are commonly observed¹. Of specific relevance to this work are *participant-specific* tendencies to prefer one alternative over the other(s). Such biases, which we term *idiosyncratic choice biases* (ICBs) have been described as early as half a century ago in perceptual discrimination²⁻⁴ and operant learning tasks⁵⁻⁷.

In discrimination tasks, the ICBs interfere with the estimate of perceptual noise. In operant learning experiments these biases mask the learning behavior. That is why such biases are typically considered as nuisance. When analyzing choice behavior, these biases are often accounted for by adding an *ad-hoc* participant-specific bias parameter³ or by counterbalancing the choices to average them out.

Many factors can contribute to ICBs. For example, in perceptual discrimination tasks, a stimulus in a given trial is often perceived as being more similar to the stimuli presented in previous trials⁸⁻¹⁰. Similarly, participants tend to choose those actions that were previously more often rewarded^{11,12}. Finally, participants may exhibit a preference towards an alternative because the corresponding motor action requires the least effort. Heterogeneity between the participants along any of these factors is sufficient to generate ICBs. One may thus expect that these biases would be diminished if these factors are controlled for in the experimental design or are factored out in the analysis. In contrast to this expectation, here we argue that even in an idealized gedanken experiment, in which symmetry between subjects in all the above factors is kept, substantial ICBs are expected. These ICBs that cannot be accounted for by the experimental context are the subject of this study.

We quantify ICBs in a perceptual discrimination task and in a novel sensory-motor task, in which sequential and operant factors are controlled for. We then analyze the ICBs in the framework of a Drift Diffusion Model (DDM) and show that they are primarily the result of biased drift rates. Finally, we show analytically and numerically that ICBs naturally emerge from the intrinsic stochasticity of the dynamics of competing populations of spiking neurons. Our work thus suggests that ICBs are inevitable unless they are actively suppressed, e.g. by the reward schedule.

Results

ICBs in the bisection discrimination task

We quantified ICBs in the bisection discrimination task depicted in Fig. 1a (inset). In each trial, a vertical transected line was presented on the screen and participants were instructed to indicate the offset direction of the transecting line (see Materials and Methods). Fig. 1a depicts the fraction of an ‘Up’ response, p_{up} as a function of the offset for three participants. As expected, the probability of a correct response increased with the magnitude of the offset $\Delta L/L \equiv (L^U - L^D)/(L^U + L^D)$, where L^U and L^D denote the lengths of the Up and Down segments of the vertical line. However, the responses differed between the three participants: the blue psychometric curve is shifted to the right of the black curve, whereas the red curve is shifted to its left.

We considered the choices of the participants in 20 “impossible” trials (1/6 of the trials), in which the line was transected at its midpoint ($\Delta L = 0$). The participant whose psychometric curve is plotted in black in Fig. 1a responded ‘Up’ in 11/20 impossible trials, which is statistically indistinguishable from chance ($p=0.82$, two-sided Binomial test). By contrast, the two other participants (red and blue in Fig. 1a) exhibited significant choice biases, responding ‘Up’ in 18/20 and 1/20 of the trials, respectively ($p<0.001$, two-sided Binomial test). Overall, 48% of the participants ($n=100$) exhibited a significant choice bias (24% significant ‘Up’, $p < 0.05$, two-sided Binomial test; 24% significant ‘Down’, $p < 0.05$, two-sided Binomial test). These ICBs were not restricted to the impossible trials. Rather, they were also observed in the possible trials albeit to a lesser degree. Biases in the possible and impossible trials were highly correlated (Fig. S1; two-sided Pearson’s $\rho=0.63$, $p<10^{-12}$). At the population level, we could not detect a global bias. The fraction of ‘Up’ choices in the impossible trials across all participants was 0.505, which is not significantly different from chance (95% CI 0.45-0.56, bootstrap).

To quantify the heterogeneity of these ICBs across the population, we computed for each participant the difference between the fraction of ‘Up’ and ‘Down’ responses in the impossible trials. This measure quantifies the bias because it vanishes for unbiased choices (ICB = 0 for $p_{\text{up}} = 0.5$) and its magnitude is maximal if choices are deterministic (ICB = -1 for $p_{\text{up}} = 0$; ICB = 1 for $p_{\text{up}} = 1$). The distribution of ICBs

across the participants is depicted in Fig. 1b. Its width is a measure of *idiosyncrasy* of these biases across the participants. We found that the variance of the distribution is significantly larger than expected by chance ($p < 10^{-6}$, one-sided bootstrap test, Bernoulli process). These results further establish the existence of ICBs in the vertical bisection task.

As mentioned in the Introduction, operant effects can contribute to ICBs. To minimize the contribution of feedback to the ICBs, participants received only sparse feedback every 30 trials on their accumulated performance until that point. Another potential contributor to ICBs is a propensity to repeat in a trial the actions taken in the previous trials. To minimize sequential effects, the impossible trials were always preceded by three irrelevant trials (see Materials and Methods). Indeed, the probability that a participant would repeat in an impossible trial the action she took in the previous (possible) trial was 0.50 ± 0.01 (average over participants $\pm$ SEM) (see also Fig. S2).

We then analyzed the ICBs in the framework of the DDM. According to the DDM, noisy evidence in favor of each alternative is integrated over the course of the trial. The difference of these evidence, a quantity known as the *decision variable*, is computed and a decision is reached once this variable reaches one of two decision thresholds. The DDM has been extensively used to explain both behavioral and neurophysiological data^{13–18}. In this framework, ICBs in the impossible trials can emerge via two mechanisms. In the first mechanism, the bias results from the initial condition of the decision variable being not equidistant from the two thresholds. In the second mechanism, the bias results from a drift bias of the decision variable, which is unrelated to the veridical evidence^{19–22}.

We investigated which of these two mechanisms best accounts for the ICBs which we observed experimentally. To that end, we fit choices and reaction-times of participants in the impossible trials to four versions of the DDM. The goodness of each fit was assessed using the Deviance Information Criterion (DIC; Materials and Methods). The first model was a baseline DDM with symmetric, *i.e.*, equidistant initial condition and no drift bias. By construction, there are no ICBs in this model and it was only used as a baseline for comparison with the other three models. To dissect the relative contributions of the drift and initial condition to the ICBs, we added to this baseline DDM (1) idiosyncratic drift rates ('drift bias' DDM), (2) idiosyncratic initial conditions ('IC

bias' DDM) or (3) both idiosyncratic drift rates and idiosyncratic initial conditions ('IC+drift bias' DDM). The DICs of all three models were compared to the DIC of the baseline model. As depicted in Fig. 1c, all three models did better than the baseline model. The 'drift bias' DDM (green) did substantially better than the 'IC bias' DDM (purple). These results suggest that in the framework of the DDM, bias in initial condition contributes less to the observed ICBs than the bias in the drift rate. We further dissected the relative contributions of the drift and IC to the ICBs in the 'IC+drift' DDM (black), which did better than the other models (Fig. 1c). To that goal, we computed for each participant the ICB expected from the DDM with parameters extracted from the 'IC+drift' DDM. As shown in Fig. 1d (black Xs), the expected and observed ICBs are in good agreement. They are also in good agreement when instead of the extracted initial conditions, symmetric ones are used (green circles). This indicates that asymmetry in the initial conditions does not play an important role in the generation of the ICBs. Indeed, when using the extracted initial conditions but unbiased drift we failed to account for the ICBs (purple squares).

Figure 1: ICBs in the vertical bisection task. **a**, Psychometric curves of three participants: the observed fraction of responding 'Up', p_{up} , vs. the sensory offset $\Delta L/L$. Error bars denote the standard error of the mean (SEM). Curves are best-fit logistic functions. Inset, a schematic illustration of the stimulus in a single trial. **b**, Distribution of ICBs (ICB = $p_{up} - p_{down}$) of all participants ($n = 100$). The ICBs of the three participants in (a) are denoted in the histogram by arrows of corresponding colors. **c**, Model comparison using DIC (Materials and Methods). The DIC of the 'IC bias' DDM (purple), 'drift bias' DDM (green) and 'IC+drift bias' DDM (black) were measured relative to the baseline DDM, $\Delta DIC = DIC_{model} - DIC_{baseline}$. Error bars are SEM, based on three repetitions of the fitting procedures. Results indicate that the 'IC+drift' DDM accounts for the data slightly better than the 'drift bias' DDM and much better than the 'IC bias' DDM. **d**, Relative contributions of the drift-bias and IC-bias to the ICBs in the 'IC+drift bias' DDM. Each symbol depicts a single participant. Abscissa: the observed p_{up} in the impossible trials. Ordinate: expected p_{up} , based on average posteriors of each participant in the biased 'IC+drift' DDM (Eq. 6 in Materials and Methods). Black Xs: both initial conditions and drifts were taken from the 'IC+drift' DDM. Green circles: Drifts were taken from the 'IC+drift' DDM with symmetric initial conditions ($z = 0.5$ in Eq. 6). Purple squares: Initial conditions were taken from the 'IC+drift' DDM assuming no drift ($A = 0$ in Eq. 6). Gray line is the diagonal. Slopes of best-fit orthogonal regressions are: black Xs, 0.96; green circles, 0.93; purple squares, 0.04.

ICB in the motor task

Next, we constructed a novel motor task, in which ICBs are unlikely to emerge from idiosyncratic sensory asymmetries. In each trial, two adjacent colored dots were

displayed on a white circular background (Inset in Fig. 2a, see also Fig. S3a). Participants were instructed to drag, as fast as possible, these two dots into a central region indicated by a larger black disk. To ensure that the participants would make two temporally-separated reaching movements, we introduced a 1.1 sec delay after the completion of the dragging of the first colored dot (Materials and Methods). The task was presented to the participants as a motor-speed task, in which faster movements are more rewarded (see Materials and Methods). However, the behavioral parameter that we were interested in was the order in which participants chose to execute the two dragging movements. In that sense, this paradigm is a symmetric, binary, implicit, decision-making task. In this task, choice bias manifests as a particular preference in the order in which the two dots are dragged. Each participant was presented with 10 pairs of dots, each pair differing in colors and locations. Each of these pairs was presented 20 times in a pseudorandom order. The ICB of a participant for a given pair of dots was defined as the difference between the fraction of trials in which the clockwise (CW) and counterclockwise (CCW) dot was dragged first ($ICB = p_{CW} - p_{CCW}$). This allowed us to measure 10 different ICBs (one for each pair) for each participant.

Figure 2a depicts the distribution of choice biases across the participants for a particular pair of dots (inset). At the population level, we could not detect a global bias. The fraction of clock-wise choices across all participants was 0.55, which is not significantly different from chance (95% CI 0.40-0.70, bootstrap). Nevertheless, 65% of the participants exhibited significant ICB for this pair (35% significant preference towards choosing first the clock-wise dot and 30% significant preference in favor of choosing first the counter-clockwise dot; $p < 0.05$ two-sided Binomial test). Consistent with that, the variance of the distribution of preferences in that pair was significantly larger than expected by the population-average preference ($p < 10^{-6}$, one-sided bootstrap test, Bernoulli process). Variance of the distributions of preferences that is significantly larger than expected by the population-average preference was observed in all ten pairs (Fig. S3b).

We then analyzed the ICBs using the DDM framework. As in the bisection task, the 'drift bias' model (green) did substantially better than the 'IC bias' model (purple) for all ten pairs, indicating a smaller contribution to behavior of the biased initial conditions relative to the contribution of biased drift rates. The DIC of the 'IC+drift bias' DDM

(black) was comparable to the DIC of the 'drift bias' DDM (Fig. 2c). In half of the pairs, the DIC of the 'IC+drift bias' DDM model was the lowest, whereas in the other half, the DIC of the 'drift bias' DDM was the lowest. We further dissected the contributions of drift biases and initial conditions to the observed ICBs using the 'IC+drift bias' DDM (Fig. 2d). As in the bisection task, we found that drift bias, rather than asymmetry in the initial conditions, is the dominant contributor to ICBs in the motor task.

Figure 2: ICBs in the motor task. $ICB = p_{cw} - p_{ccw}$ where p_{cw} and p_{ccw} are the probabilities of choosing first the clockwise and the counter clockwise dot. **a**, The distribution of ICBs of all participants ($n = 20$) for the pair of dots in the inset. **b**, The distribution of ICBs for all 10 pairs of dots in the experiment (color-coded as in Fig. S3b). **c**, Model comparison using DIC, as Fig. 1c. Model fits were performed separately on each pair of dots. Bars and error bars denote the average and SEM, ΔDIC , over the 10 pairs of dots. **d**, Same as Fig. 1d, demonstrating that the ICBs in the motor task are dominated by the drift-biases. Slopes of best-fit orthogonal regressions are: black Xs, 0.98; green circles, 0.99; purple squares, -0.01.

ICBs in the Poisson network model

What underlies participant-to-participant differences in drift rates in the bisection and motor tasks? To address this question, we constructed a simple neuronal network model of decision making and used it to study behavior in the bisection task. It consists of two populations of neurons representing 'Up' and 'Down' choices, denoted by ' U ' and ' D ' (Fig. 3a, left). Each population is made of $N/2$ independent Poisson neurons, such that the spike train of each neuron in a trial is an independent homogeneous Poisson process. The firing rates of the neurons depend on the offset in the input (ΔL) such that the firing rates of the U neurons increase with ΔL , whereas

The psychometric curve of an example network is depicted in Fig. 4a (center; black). Because of the dependence of the firing rate distributions on ΔL , the larger ΔL the more likely it is that the network would choose 'Up'. However, the outcome of this decision process is not deterministic. Because spiking is stochastic, $\Delta n(t)$ occasionally reaches the threshold that is incongruent with the stimulus, resulting in an error. More generally, because of this stochasticity, the psychometric curve is a smooth sigmoidal function of ΔL rather than a step function. Note that in the black psychometric curve of Fig. 4a (center), the network's perceptual decision in the "impossible trials" is approximately at chance level. Thus, this particular network does not exhibit a substantial ICB.

The black psychometric curve in Fig. 4a (center) was obtained for a particular realization of the network. The red and blue lines in Fig. 4a (center) depict the psychometric curves of two other realizations of the network. Despite the fact that the three networks were constructed in the same way, *i.e.*, by randomly drawing the firing rates of the neurons from the same distributions, the red and blue curves are horizontally shifted relative to the black psychometric curve. Thus, in contrast to the "black" network, the "red" and "blue" networks exhibit ICBs in favor and against responding 'Up'. The distribution of the ICBs across networks is depicted in Fig. 4b (center). It demonstrates that a wide distribution of ICBs naturally emerges in this decision network model.

It is possible to mathematically prove that for a large network, the behavior of this Poisson network model is equivalent to that of a DDM with a biased drift (Materials and Methods). The accumulation of the difference in the spike counts can be mapped to the accumulation of noisy evidence in the DDM; trial-to-trial variability results from the stochasticity in the neuronal firing; the drift bias stems from the heterogeneity in the firing rates in the two populations. In what follows, we provide an intuitive explanation for the emergence of ICBs in the Poisson network model.

A wide distribution of ICBs in a network consisting of a *small* number of neurons is easy to understand. Let us consider the impossible trials in a network composed of only two Poisson neurons (Eq. 1 in Materials and Methods), each representing one choice ('Up' or 'Down'). The firing rates of the two neurons are independently drawn from the same lognormal distribution. However, the actual firing rates of these two neurons will, in general, differ in any given network. In some realizations of the network

the firing rate of the '*U*' neuron will be higher than that of the '*D*' neuron, whereas in others, it will be lower. Choice is determined by the first threshold-reaching of the accumulated difference in the number of spikes fired by the neurons. It will more often be congruent with the neuron whose firing rate is higher. However, because the firing of spikes in the model is stochastic, the decision in a minority of trials decision will be incongruent with that neuron. This argument implies that this two-neuron network exhibits an ICB, which results from the interplay between the Poisson noise and the heterogeneity in the firing rates of the two neurons. The spiking stochasticity decreases the bias, whereas the firing-rate heterogeneity increases it.

It is thus clear why ICB is natural in such small decision-making networks. However, it is not immediately clear why ICBs are observed in Fig. 4a, in which the number of neurons in the network is large ($N = 200,000$). In this case, the difference between the *population averaged* firing rates of the *U* and *D* neurons is vanishingly small. This is because in large networks, this difference is of the order of $1/\sqrt{N}$, where N is the number of neurons. Since the competition between the two populations, which underlies the decision making, is a macroscopic process, one may expect that only average properties of the two populations would affect its outcome. Thus, heterogeneity in the firing rates should not play a significant role in the decision process in large networks. One should note, however, that the sensitivity of the decision-making process to the firing rate heterogeneity increases in proportion to $\sqrt{N}$. This is because the spike times are independent between neurons and thus the fluctuations in spike count decrease with N . As a result, the larger the network, the more sensitive it is to the firing rate heterogeneity. Because the sensitivity to the heterogeneity increases in proportion to $\sqrt{N}$ while the heterogeneity itself is proportional to $1/\sqrt{N}$, the effect of the heterogeneity in firing rates on behavior is independent of N (in the limit of large N). Thus, even if N is very large, the distribution of ICBs is wide (Materials and Methods).

Unlike network size, the decision threshold has a large effect on the magnitude of the ICB. This is depicted in Figs. 4a, where the psychometric curves of three networks, only differing in the value of the decision threshold, are plotted. The larger the threshold, the steeper is the psychometric curves. This is because the time it takes the network to reach a decision increases with the threshold (Fig. 4c). Thus, a larger threshold results in the integration of spikes over longer durations before a decision is

made. Therefore, the decision outcome is less sensitive to the Poisson noise. On the other hand, the network heterogeneity is independent of decision time. Because the magnitude of the ICB is determined by the interplay of the Poisson noise and networks heterogeneity, the larger the threshold is, the broader will be the distribution of ICBs (Fig. 4b; see also eq. (5) and Fig. S4c-d).

Figure 4: ICBs in the Poisson network model. Left, center and right correspond to network behaviors with low, intermediate and high thresholds (See Materials and Methods). **a**, The psychometric curves of three networks. Each color corresponds to a different network and the same color in different panels corresponds to the psychometric curves of the *same* network with different thresholds (Eq. 2 in Materials and Methods). **b**, ICB distributions (Eq. 3 in Materials and Methods). The ICBs of the three networks in (a) are denoted in the histogram by arrows of corresponding colors. **c**, Distribution of reaction times (RTs) (Eq. 4 in Materials and Methods). Number of neurons in the network is $N = 200,000$, $\theta^* = 0.65 \cdot \sqrt{N}$.

ICB in the recurrent spiking network

Our analysis of the Poisson model suggests that decision making networks exhibit ICBs if (1) the constituting neurons fire irregularly (e.g., Poisson), (2) the neuronal firing

rates are heterogeneous (e.g., log-normally distributed) and (3) decision is based on competition (e.g., threshold crossing of the difference in spike counts). In the Poisson model, these three ingredients are introduced *ad hoc*. Here we investigate a spiking network model in which these features are emergent properties of the network dynamics.

This model builds on previous studies that have shown that recurrent networks of excitatory and inhibitory neurons connected by numerous and strong synapses readily operate in a regime in which excitation is dynamically balanced by inhibition²⁶. Two hallmarks of this regime are (1) Poisson-like temporal variability of spike timing (Fig. 5b and S5) and (2) approximately log-normally distribution of firing rates (Fig. 5c). These features emerge from the intrinsic deterministic dynamics of the network (even when the neurons are identical and receive the same external input)^{27,28}. Similar to previous models of decision making^{29,30}, competition between the alternative actions in our model is mediated by inhibition.

Our model consists of 32,000 excitatory and 8,000 inhibitory Leaky Integrate and Fire (LIF) neurons (Fig. 5a; see Materials and Methods). All neurons receive a feedforward input, which is selective to the stimulus. For half of the neurons ('U' neurons), this input linearly increases with ΔL , whereas for the other half, ('D' neurons), it is a decreasing function of ΔL (Fig. 5a, bottom). When the two segments are of equal length (impossible trials), the 'U' and 'D' neurons receive the same feedforward input. All neurons are recurrently connected by strong synapses in a random and non-specific manner, *i.e.*, independent of the selectivity properties of the pre- and post-synaptic neurons. The competition between the 'U' and the 'D' neurons is mediated by an additional set of inhibitory connections, which are functionally specific, less numerous but stronger than the unspecific ones (Fig. 5a center; black, specific; gray, non-specific). To investigate the dynamics of this model we performed numerical simulations (See Materials and Methods).

Figure 5d depicts the spike times of 3,200 excitatory neurons in one 'impossible' trial. Before the stimulus is presented ($t < 0$), the activities of the 'U' and 'D' neurons are similar (Fig. 5d). In response to the sensory stimulus ($t = 0$), the neurons in both populations increase their firing rates. Because of the competition induced by the specific inhibitory connectivity, population 'D' inhibits population 'U' and as a result

population 'U' disinhibits the excitatory neurons in population 'D'. In our model, the decision occurs once the relative difference in the average firing rates of the excitatory neurons of the 'U' and the 'D' populations exceeds a threshold. After the decision is made, the feedforward stimulus-dependent input ceases and the network activity reverts to its baseline levels (Fig. 5d; see also Materials and Methods).

Figure 5: The recurrent spiking network model. **a**, Schematic illustration of the network architecture. The network consists of recurrently connected excitatory (red) and inhibitory (blue) LIF neurons receiving stimulus-selective feed-forward input (bottom). Direction of triangle indicates the selectivity. **b**, Spontaneous activity of example excitatory (red) and inhibitory (blue) neurons in the network. **c**, Distribution of the spontaneous firing rates of the 32,000 excitatory neurons in the network. **d**, Raster plot (10% of the excitatory neurons, top) and the average firing rates (bottom) of the excitatory 'U' (purple) and 'D' (green) neurons in response to a stimulus (gray region). The decision is made when the relative difference between the firing rates of the excitatory neurons of the two populations crosses the decision threshold (after which the feed-forward input ceases).

Fig. 6a (center, black) depicts the psychometric curve of the specific network simulated in Fig. 5d. When the magnitude of ΔL is large, the perceptual decision of the network is almost always correct; as ΔL decreases, the error rate increases. Considering the "impossible trials" ($\Delta L = 0$), the network's perceptual decision is approximately at chance level. However, different realizations of the connectivity matrix yield psychometric curves, which are laterally shifted (red and blue curves in Fig. 6a). In contrast to the "black" network, the "red" and "blue" networks exhibit substantial ICBs.

To estimate the distribution of ICBs in our recurrent network model, we simulated 200 networks, which only differed in their realizations of the connectivity matrix. We computed the ICB of each network from its choices in 500 “impossible” trials. The center panel in Fig. 6b depicts the distribution of these ICBs across the 200 networks. It is significantly wider than expected by chance ($p < 10^{-6}$, one-sided bootstrap test, fair Bernoulli process).

The level of competition in our model is determined by the strength of the functionally specific inhibition, g (see Materials and Methods). Figure 6b depicts the distribution of ICBs for three values of g . As g increases, the width of the distribution decreases and its shape changes from concave to convex. The distribution of decision times also varies with g . The larger g the faster is the average decision time (Fig. 6c). When the recurrent network model is analyzed in the framework of the ‘drift bias’ DDM, decreasing the specific inhibition g manifests primarily as an increase in the decision threshold (Fig. S6).

Figure 6: ICBs in the recurrent spiking network model. The strength of the selective inhibition is: Left, $g = 3.4$; Center, $g = 3.0$; Right, $g = 2.7$. **a**, The psychometric curves. Each color corresponds to a different network and the same color in different panels corresponds to the psychometric curves of the same network but with a different g . Each point is an average over 500 trials. Error bars correspond to SEM. **b**, Distribution of ICBs for 200 networks. Arrows correspond to the specific psychometric curves in (a) (same color coded). **c**, Distribution of reaction times (RTs).

It is instructive to analyze the recurrent network model in the framework of the DDM, as we did for the experimental data. Fitting the DDM to the recurrent network simulations, we found that a ‘drift bias’ DDM better explains the network dynamics than ‘IC bias’ DDM (Fig. S7a). Moreover, when the relative contribution of the drift bias and IC bias are tested in the ‘IC+drift bias’ DDM, the contribution of drift bias dominates the emergent ICBs (Fig. S7b). These results are similar to those observed in the behavioral data (Figs. 1d and 2d).

The conditional bias function

Responses in decision-making tasks can also be analyzed using the *conditional bias function* (CBF). This function quantifies the relationship between bias magnitude and reaction time within the responses of the decision-maker¹⁹. In the 'bias drift' DDM, the distributions of decision times for congruent and incongruent choices are equal^{17,31–33}. Therefore, bias is independent of the decision time. By contrast, in the 'IC bias' DDM, the bias decreases with decision time^{19,20}.

We applied the CBF analysis to the impossible trials in our tasks. Figure 7a depicts the CBF of the responses in the spiking network model, averaged over the 200 networks of Fig. 6. We found that the magnitude of the choice bias decreases with reaction time. The larger g , the more negative is the slope of the CBF (Fig. 7a, inset). Because a dependence of the bias on reaction time is expected from asymmetric initial conditions in the DDM, we compared the CBF to that predicted from the fitted 'IC+drift bias' DDM. We found that the dependence of the bias on the reaction time in the recurrent model is not well-captured by the fitted 'IC+drift bias' DDM (Fig. 7b).

We hypothesize that the discrepancy between the CBF of the recurrent network model and that of the corresponding 'IC+drift bias' DDM as resulting from a qualitative difference in the decision process in the DDM and the recurrent network. Competition in the recurrent network but not in the DDM results in an effective positive feedback. As a result, evidence accumulated in the beginning of the decision process has a larger effect on the responses than the evidence accumulated later in that process.

This discrepancy prompted us to compute the CBF in our behavioral data. As shown in Fig. 7c, in both the bisection and motor tasks the dependence of the bias on the reaction time is stronger than expected from the fitted 'IC+drift bias' DDM. These results indicate that the recurrent network model captures additional features of the behavior, beyond the 'IC+drift bias' DDM.

Figure 7 Conditional bias functions. **a**, The recurrent network model. We numerically simulated the recurrent network model in the impossible trials. Responses of each network were divided into 5 quantiles (quintiles) according to the reaction time (RTs). Within each quintile, the fraction of choices congruent with the overall bias of the network, p_{bias} , was computed and averaged over the different networks ($n = 200$ networks, 500 trials per network). This analysis was performed independently for different values of g (different colors; data for $g = 2.7, 3.0, 3.4$ are the same as in Fig. 6). Inset, the slopes of the different curves, defined as the change in p_{bias} per percentile of reaction time, $dp_{bias}/d\%$. **b**, 'IC+drift bias' DDM fits to the recurrent networks. Same analysis as in (a) was performed on the fitted 'IC+drift bias' DDM to the networks' responses using 10,000 simulated responses for each network. **c**, Responses of human participants. Solid lines, the conditional bias function for the impossible trials in the bisection task (black) and the motor task (gray). Dashed lines, the corresponding conditional bias functions for the fitted 'IC+drift bias' DDMs, based on 2,000 simulated responses for each decision-maker. Error bars are SEM.

Discussion

We experimentally investigated human ICBs in a discrimination task and in a motor task. We analyzed the behavior of the participants in the framework of the DDM. We found that in this framework, idiosyncratic biases in the drift rate account for these ICBs. We proved mathematically in a particular model that ICBs due to idiosyncratic drift biases naturally emerge in a network characterized by (1) irregular firing of the neurons, (2) heterogeneity of their firing rates and (3) competition. Finally, we constructed a recurrent network model of spiking neurons, in which these three features are the result of the deterministic dynamics. We numerically simulated this network and demonstrated that it exhibits ICBs, whose features are similar to those observed experimentally. Taken together, our results show that ICBs naturally emerge from the intrinsic dynamics of decision-making neural circuits.

Bias in the DDM

In the framework of the DDM, choice bias can emerge either from a drift-bias in the decision variable or from asymmetry in its initial condition^{19–22}. There is an active area of research that maps different factors affecting choice bias to these two mechanisms. In perceptual tasks, stimuli affect responses via the drift rate¹³. Other factors such as decision cutoff^{19,21} and the bias induced by the action taken in the previous trial³⁴ have both been primarily associated with biases in the drift rate. Arousal levels also affect the magnitude of choice biases through the drift rate³⁵. By contrast, asymmetries in the prior distribution of stimuli or in the reward schedule predominantly manifests as an asymmetric initial condition^{19,21,22,36} (but see³⁷). Heterogeneity among the participants along any of these factors is expected to result in ICBs. Our modeling work predicts the existence of additional, irreducible, ICBs. These ICBs cannot be explained by the experimental context and manifest as drift rate idiosyncrasies in the DDM. Our experimental work reports such ICBs in a discrimination task and a motor task.

The temporal-scale of stochasticity

In the two models that we have investigated, ICBs emerge from the interplay of two sources of stochasticity: (1) Stochasticity in the timing of action potentials; (2) Heterogeneity in the neuronal firing rates. Stochasticity in the timing of action potentials differs between trials and therefore we can refer to it as *fast* noise. By contrast, the second source of stochasticity is the same in all trials and therefore, we can refer to it as *frozen* noise. Cortical dynamics exhibits additional time-scales³⁸. Incorporating additional time-scales to the models will not qualitatively affect the results, as long as the contributions of these additional sources of stochasticity are of the order of $1/\sqrt{N}$ (where N is the number of neurons in the network).

To identify the potential contribution of stochasticity at minutes' time-scale, we tested whether ICBs differed between the first and second halves of our experiments. We did not observe changes in the ICBs in this time-scale that are statistically significant (vertical bisection + motor: two-sided permutation test identified significant, differences in only 17/300 of the pairs, $p < 0.05$, not corrected for multiple comparisons). It will be interesting to quantify the dynamics of ICBs over longer time-scales.

The effect of correlations

In the Poisson model, spikes are uncorrelated in time and between neurons. As a result, for sufficiently large networks, the magnitudes of both fast and frozen sources of stochasticity decrease as $1/\sqrt{N}$. Their ratio and hence the distribution of ICBs become independent of N for sufficiently large networks. The two sources of stochasticity also satisfy these scalings in the recurrent network model. This is because the network operates in the balanced regime^{26,39}. Noise correlations in the spike count of the neurons are therefore very weak^{40–42} and the firing rates are widely distributed and are uncorrelated between neurons²⁷. In network exhibiting correlations in the neuronal activity^{41,42}, averaging over neurons may not decrease the fast noise and the heterogeneities as $1/\sqrt{N}$. If the dependence of the two on N is very different, one source of stochasticity could dominate, resulting in deterministic or unbiased choices.

Alternative interpretation of the observed ICBs

Idiosyncratic postures can affect visual stimuli or the effort associated with motor responses. One cannot rule out the possibility that such idiosyncrasies contributed to the ICBs in the bisection task, which was performed online. In the motor task, by contrast, all participants were dextral and the positions of the chair, screen and mouse pad were kept the same for all participants. Nevertheless, we cannot exclude the possibility that differences in participants' anatomy, *e.g.*, their arm length, contributed to the ICBs.

Reinforcers can affect choice preferences^{43–45}. The specific history of stimuli also influences preferences in perceptual tasks^{9,46}. Along these lines, it is natural to attribute ICBs to the specific histories of the participants during the experiment. We therefore designed our tasks to minimize operant and sequential effects. Nevertheless, we cannot exclude the possibility that the observed ICBs are the result of operant or sequential effects which occurred *before* the experiment. For example, considering the impossible trials in our bisection task, participants may prefer to press the Down key because they are accustomed to pressing taskbar icons that are located at the bottom of their computer monitor. Other participants may prefer the Up key because they are used to a taskbar located at the top of the screen. In such a view, ICBs in the vertical

bisection task can be attributed to idiosyncratic histories of computer usage prior to the experiment.

All of the above effects can contribute to ICBs. However, we showed that these and similar explanations are not necessary. All that is required for ICBs are minute differences between the populations encoding the two alternatives. Such differences are almost inevitable in any 2-alternative task, in which the two alternatives are represented by different populations of neurons.

Substantial ICBs were observed in genetically-identical flies that were reared in the same environment⁴⁷. The results of that study suggest that biases can emerge from effects that are unpredictable from genetic, environmental or anatomical variables. This is in line with our study that showed that the random differences in the fine structure of connectivity between the neuronal populations involved in decision-making are sufficient to account for the ICBs. In conclusion, the occurrence of ICBs in a cortical-based decision task is thus almost inevitable. It would be therefore surprising to find a decision task that is devoid of the ICBs, unless they are actively suppressed, e.g., by penalizing them.

Materials and Methods

The perceptual discrimination task

The study was approved by the Hebrew University Committee for the Use of Human Subjects in Research. Recruitment was based on the online labor market Amazon Mechanical Turk⁴⁸. Data were collected from 100 participants (51 males, 49 females; 91 dextrals, 7 sinistral, 2 ambidextrous; mean age = 39 years, min = 22 years, max = 71 years). All participants were Mechanical Turk's Masters, located in the United States of America. All participants reported normal or corrected to normal vision and no history of neurological disorders. The experiment was described as an academic survey of visual acuity. A base monetary compensation was given to all applied participants for the participation. In order to encourage good performance throughout the experiment, an additional bonus fee was given for every correct response and another bonus was guaranteed to 10% of participants with highest scores.

Procedure

Participants were instructed to indicate the offset direction of the transecting line, out of two alternative responses. Possible responses were either 'Left' or 'Right', for the horizontal discrimination task, or 'Up' or 'Down', for the vertical discrimination task. Participants were asked to answer as quickly and accurately as possible.

In each trial, a 200 pixel-long white line, transected by a perpendicular 20 pixel-long white line was presented on a black screen (Fig. 1a, inset). The stimuli were limited to a 400-pixel X 400-pixel square at the center of the screen. Window resolution was verified for each participant individually, to make sure that it did not exceed the centric box in which all stimuli were presented. The horizontal location of all vertical bisection lines and the vertical location of all horizontal bisection lines were centered. After 1 sec, the stimulus was replaced by a decision screen composed of two arrows, appearing in opposite sides of the screen, and a middle 4-squares submit button. The participants indicated their decision by moving the initially centered cursor to one of the arrows, pressing it, and finalizing their decision by pressing the 'submit' button. No feedback was given regarding the correct response. The participants were, however, informed about the accumulated bonus fee every 30 trials.

The experiment consisted of 240 trials, 120 horizontal and 120 vertical. Trials were ordered in 80 alternating blocks of 3 horizontal and 3 vertical transected lines. Unbeknown to the participants, there were 20 impossible horizontal and 20 impossible vertical trials ($\frac{1}{6}$ of the trials). To minimize sequential effects in the impossible vertical bisection trials, each impossible vertical bisection trial was preceded by three horizontal bisection trials. The order of the trials was pseudorandom but identical for all participants. For the possible trials, the deviation from the veridical midpoint was uniformly distributed between 5 and 10 pixels ($|\Delta L|/L$ between 0.05 and 0.1, where $\Delta L/L \equiv (L^U - L^D)/(L^U + L^D)$ and L^U and L^D denote the lengths of the Up and Down segments of the vertical line). with an equal number of offsets in each direction. Because it is well established that in the horizontal bisection task participants exhibit a global bias (attributed to pseudoneglect⁴⁹), we focused on the vertical bisection trial in quantifying ICBs and performance. Mean performance in the possible vertical trials was $96.4\% \pm 4.6\%$ (standard deviation), range 71% – 100%. No participants were

excluded from the analysis. In the DDM analysis we excluded trials, in which the reaction time was longer than 3 sec. This excluded 1% of the vertical bisection trials.

To verify that the participants understood the instructions, they were required prior to the experiment to successfully complete a horizontal-bisection practice session and a vertical-bisection practice session. A session consisted of blocks of 4 easy trials ($|\Delta L|/L = 0.2$) with feedback and balanced polarity of ΔL . The main experiment started after the participant completed one horizontal and one vertical block successfully. Responses in this practice session were not included in the analysis.

The motor task

The study was approved by the Hebrew University Committee for the Use of Human Subjects in Research. The experiment was described as an academic survey testing speed of motion. Data were collected from 20 participants (13 males, 7 females; all dextrals; mean age = 25 years, min = 19 years, max = 41 years) who were recruited using on-campus advertising. All participants reported normal or corrected to normal vision and no history of neurological disorders.

Procedure

In each trial, a pair of dots, equally distant from a central black disk, were presented on a background of a larger white disk (Figs. 2a and S3a). Participants were instructed to drag as quickly as possible the two dots into the black disk using the mouse cursor. Each trial started with a forced delay period of 0.75 sec. Then, the mouse cursor appeared in the center of the disc. The participant used the mouse to move the cursor to one of the dots. She then dragged the chosen dot to the central black disk by pressing the mouse and moving it. If accurate, a release of the dot on the central black disk resulted in a 1.1 sec “swallowing” of the dot animation, indicating a successful drag. The dragging time (measured from the time of clicking on the dot to the time of its release) appeared on the screen. It disappeared after a forced delay of 1.1 sec and the cursor reappeared in the center of the disk. The participant processed the second dot in the same way as the first dot. We used 10 different pairs of dots, each presented 20 times. Each pair of dots was of equal distance from the center of the black disk, but of a different color and a different angular location (Fig. S3b). The order of presentation

was pseudorandom such that in every consecutive group of 10 trials all pairs appeared. Decision time in a trial is defined as the time elapsed from cursor appearance to the beginning of the dragging of the first dot. The positions of the chair, screen and mouse pad were fixed and identical for all participants in order to minimize heterogeneity between participants.

A base monetary compensation was given to all participants for their participation. An additional bonus fee was given based on dragging times in order to encourage good performance throughout the experiment.

In the DDM analysis we excluded trials, in which the reaction time was longer than 3 sec. This excluded 2% of the motor trials.

The Poisson model

We consider two populations of neurons, denoted by ‘U’ and ‘D’, representing choice ‘Up’ and ‘Down’ (Fig. 3a). Each population consists of $N/2$ independent Poisson neurons. The stimulus-dependent feedforward inputs to neuron i ($i \in \{1, \dots, N/2\}$) in population α ($\alpha \in \{U, D\}$) is given by: $\mu_i^\alpha = k^\alpha \cdot \Delta L/L + z_i^\alpha$, where $k^U = -k^D = k$ is a parameter and z_i^α is stimulus- and trial-independent, independently drawn (once) from a zero-mean Gaussian distribution with variance σ^2 , $\langle z_i^\alpha \rangle = 0$, $\langle (z_i^\alpha)^2 \rangle = \sigma^2$, where $\langle \dots \rangle$ denotes average. The firing rate v_i^α , different for each neuron, is

$$v_i^\alpha = \bar{v} \cdot e^{\gamma \cdot \mu_i^\alpha} \quad (1)$$

where $\bar{v}$ is a baseline firing rate and γ is the gain²⁷. Due to the exponential transfer function and the normal distribution of inputs, the firing rates are log-normally distributed. In each trial, the cumulative number of spikes, $n^U(t)$ and $n^D(t)$, emitted by populations ‘U’ and ‘D’ up to time t in a trial is counted (Fig. 3a). A decision is made when (t^*) the absolute value of the difference in the numbers of spikes, $|\Delta n(t^*)| = |n^U(t) - n^D(t)|$, reaches a given threshold $\theta = \sqrt{N} \cdot \tilde{\theta}$, where $\tilde{\theta}$ is a parameter.

For $N \gg 1$ and neglecting the threshold effect, the difference in spike count at time t is given by $\Delta n(t) \sim \mathcal{N}(\Delta v \cdot t, \Sigma v \cdot t)$, where $\Delta v = \sum_i v_i^U - \sum_i v_i^D$ and $\Sigma v = \sum_i v_i^U + \sum_i v_i^D$. Because $N \gg 1$, both Δv and Σv are normally distributed:

$$\Delta v \sim \mathcal{N}\left(N\bar{v}e^{\frac{\gamma^2\sigma^2}{2}}\sinh\left(\gamma k\frac{\Delta L}{L}\right), N\bar{v}^2(e^{\gamma^2\sigma^2}-1)e^{\gamma^2\sigma^2}\cosh\left(2\gamma k\frac{\Delta L}{L}\right)\right)$$

$$\Sigma v \sim \mathcal{N}\left(N\bar{v}e^{\frac{\gamma^2\sigma^2}{2}}\cosh\left(\gamma k\frac{\Delta L}{L}\right), N\bar{v}^2(e^{\gamma^2\sigma^2}-1)e^{\gamma^2\sigma^2}\cosh\left(2\gamma k\frac{\Delta L}{L}\right)\right).$$

Note that Δn and Δv are different stochastic processes: the stochasticity of Δn stems from trial-by-trial variability, conditioned on the firing rates of the neurons. By contrast, the stochasticity of Δv reflects heterogeneity in these firing rates across different realizations of the decision-making network.

The standard deviation of the distribution of Σv is of $O(\sqrt{N})$, whereas its mean is $O(N)$ even when $\Delta L \rightarrow 0$. Therefore, in the limit $N \gg 1$, $\Sigma v \approx N\bar{v}e^{\frac{\gamma^2\sigma^2}{2}}\cosh\left(\gamma k\frac{\Delta L}{L}\right)$. By contrast, in the regime in which $|\Delta L/L| = O(1/\sqrt{N})$, the mean and standard deviations of the distribution of Δv are comparable, both are $O(\sqrt{N})$.

The probability of an ‘Up’ decision is obtained by solving a first-passage problem, yielding

$$p \equiv \Pr(\text{‘Up’}) = \left(1 + e^{-2\Delta v\sqrt{N}\cdot\tilde{\theta}/\Sigma v}\right)^{-1} = \left(1 + e^{-2\frac{\Delta v}{\sqrt{N}}\cdot\frac{\tilde{\theta}}{\bar{v}\cdot\cosh\left(\gamma k\frac{\Delta L}{L}\right)\cdot e^{\frac{\gamma^2\sigma^2}{2}}}}\right)^{-1} \quad (2)$$

Substituting the dependence of Δv on ΔL in Eq. (2) yields the psychometric curve. In particular, when the two networks are symmetric, $\Delta v = N\bar{v}e^{\frac{\gamma^2\sigma^2}{2}}\sinh\left(\gamma k\frac{\Delta L}{L}\right)$, and

$$p = \left(1 + e^{-2\sqrt{N}\cdot\tilde{\theta}\tanh\left(\gamma k\frac{\Delta L}{L}\right)}\right)^{-1} \approx \left(1 + e^{-2\sqrt{N}\cdot\tilde{\theta}\cdot\gamma k\frac{\Delta L}{L}}\right)^{-1}$$

More generally, when the two networks are only drawn from the same distribution, the psychometric curve will be horizontally shifted relative to the identical networks case.

To compute the distribution of ICBs, we consider the case in which the external input is symmetric, $\Delta L = 0$ and thus $\Delta v \sim \mathcal{N}(0, N\bar{v}^2(e^{\gamma^2\sigma^2}-1) \cdot e^{\gamma^2\sigma^2})$. After a change of variables,

$$\Pr(p) = \frac{1}{p \cdot (1-p)} \cdot \frac{1}{\sqrt{8\pi(e^{\gamma^2\sigma^2}-1)\tilde{\theta}^2}} e^{-\frac{(\log(p)-\log(1-p))^2}{8(e^{\gamma^2\sigma^2}-1)\tilde{\theta}^2}}$$

Because $\text{ICB} = 2p - 1$,

$$\Pr(\text{ICB}) = \frac{2}{1-\text{ICB}^2} \cdot \frac{1}{\sqrt{8\pi(e^{\gamma^2\sigma^2}-1)\tilde{\theta}^2}} e^{-\frac{\left(\log\left(\frac{1+\text{ICB}}{1-\text{ICB}}\right)\right)^2}{8(e^{\gamma^2\sigma^2}-1)\tilde{\theta}^2}} \quad (3)$$

The corresponding distribution of decision times is computed by averaging the drift-conditioned distribution of first-passage times over the distribution of Δv , yielding^{17,31,32}:

$$f(t) = \frac{\pi}{2\theta^2} \bar{v} e^{\frac{\gamma^2\sigma^2}{2}} \frac{1}{\sqrt{1+t\bar{v}e^{\frac{\gamma^2\sigma^2}{2}}(e^{\gamma^2\sigma^2}-1)}} \exp\left(\frac{1}{2} \frac{\tilde{\theta}^2}{t\bar{v}e^{\frac{\gamma^2\sigma^2}{2}} + \frac{1}{(e^{\gamma^2\sigma^2}-1)}}\right) \\ \times \sum_{m=1}^{\infty} m \sin\left(\frac{\pi m}{2}\right) \exp\left(-t \frac{m^2\pi^2}{8} \frac{\bar{v}}{\tilde{\theta}^2} e^{\frac{\gamma^2\sigma^2}{2}}\right) \quad (4)$$

Two points are worthwhile noting:

- (1) The neuronal gain parameter γ affects $\Pr(p)$ through the term $(e^{\gamma^2\sigma^2} - 1)\tilde{\theta}^2$. This implies that increasing the gain is effectively equivalent to increasing the threshold parameter $\tilde{\theta}$, and thus is likely to broaden the distribution of ICBs.
- (2) The assumption of a lognormal distribution of firing rates is not essential to our analysis. For a general distribution of firing rates, Eq. (3) becomes

$$\Pr(p) = \frac{\lambda}{\sqrt{\pi p(1-p)}} \cdot e^{-(\lambda(\log(p)-\log(1-p)))^2}$$

and

$$\Pr(\text{ICB}) = \frac{2\lambda}{\sqrt{\pi(1-\text{ICB}^2)}} \cdot e^{-\left(\lambda \log\left(\frac{1+\text{ICB}}{1-\text{ICB}}\right)\right)^2} \quad (5)$$

where $\lambda = \frac{E[v_i^\alpha]}{\sqrt{8} \cdot \tilde{\theta} \cdot \sqrt{V[v_i^\alpha]}}$ and $E[v_i^\alpha]$ and $V[v_i^\alpha]$ are the mean and variance of the distribution

of the firing rates in the impossible trials.

The parameters used in all simulations are $\bar{\nu} = 1.26$ Hz, $\gamma = 1$, $k = 0.133$ and $\sigma^2 = 1$. For $\Delta L = 0$, the average and standard deviation firing rate are 2.1 Hz and 2.7 Hz. These numbers are compatible with experimental data in the cortex^{25,50}. The dependence of the width of the ICB distributions on the model parameters are depicted in Fig. S4.

The spiking network model

The model is a recurrent network of N leaky-integrate-and-fire (LIF) neurons, $N^E = 0.8N$ excitatory and $N^I = 0.2N$ inhibitory (the superscript denotes neuron type, excitatory or inhibitory).

Single neuron dynamics: The sub-threshold dynamics of the membrane potential, $V_i^\alpha(t)$, of neuron i in population α ($i = 1, \dots, N^\alpha$; $\alpha = E, I$) follow:

$$\tau_m \frac{dV_i^\alpha(t)}{dt} = -(V_i^\alpha(t) - V_L) + I_{rec,i}^\alpha(t) + I_{FF,i}^\alpha(t) + I_b^\alpha$$

where τ_m is the neuron membrane time constant, V_L is the reversal potential of the leak current. Inputs to the neuron are modeled as currents: $I_{rec,i}^\alpha(t)$ is the recurrent input into neuron (i, α) , due to its interactions with other neurons in the network, $I_{FF,i}^\alpha(t)$ is the feedforward input into that neuron elicited upon presentation of the stimulus, and I_b^α is a background feedforward input, independent of the stimulus, identical for all the neurons and constant in time. These subthreshold dynamics are supplemented by a reset condition: if at $t = t_i^\alpha$ the membrane potential of neuron (i, α) reaches the threshold, $V_i^\alpha(t_i^{\alpha-}) = V_T$, the neuron fires an action potential and its voltage resets to $V_i^\alpha(t_i^{\alpha+}) = V_R$.

The feedforward input: Each population, excitatory or inhibitory, consists of two types of neurons, namely U - and D -selective. In the absence of stimulus, the feedforward input $I_{FF,i}^\alpha(t) = 0$ for all the neurons. Upon presentation of a stimulus for which $\Delta L > 0$, $I_{FF,i}^\alpha(t)$ into U -selective neurons is stronger than $I_{FF,i}^\alpha(t)$ into D -selective neurons. The opposite is true when $\Delta L < 0$. Specifically, we take:

$$I_{FF,i}^\alpha(t) = I_0^\alpha + \varepsilon \frac{\Delta L}{L} I_1^\alpha$$

where I_0^α and I_1^α are constants and positive and ε characterizes the selectivity of the neuron: $\varepsilon = +1$ for '*U*' neurons and $\varepsilon = -1$ for '*D*' neurons. We denote the set of *U*-selective (resp. *D*-selective) neurons in population $\alpha = E, I$ by U^α (resp. D^α). Neuron $(i, \alpha) \in U^\alpha$ if $i = 1 \dots \frac{N^\alpha}{2}$ and $(i, \alpha) \in D^\alpha$ if $i = \frac{N^\alpha}{2} + 1 \dots N^\alpha$.

The recurrent input: The connectivity has two components. One is functionally specific and the other is not. The non-specific component is fully random (Erdős-Renyi graph) and does not depend of the selectivity of the pre- and post-synaptic neurons. The corresponding $N^\alpha \times N^\beta$ connectivity matrix, $C_{NS}^{\alpha\beta}$, is such that $C_{NS,ij}^{\alpha\beta} = 1$ with probability K/N^β and $C_{NS,ij}^{\alpha\beta} = 0$ otherwise, where K is the average number of non-specific inputs that a neuron receives from neurons in population β . The strength of the non-specific connections depends solely on α, β yielding: $J_{NS,ij}^{\alpha\beta} = J_{NS}^{\alpha\beta} C_{NS,ij}^{\alpha\beta}$ where $J_{NS}^{\alpha E} > 0$ (excitation) and $J_{NS}^{\alpha I} < 0$ (inhibition).

The competition between the '*U*' and the '*D*' selective neurons is mediated by an additional set of connections. These connections are specific and are much less numerous but stronger than the unspecific ones. The corresponding connectivity matrices, $C_{S,ij}^{\alpha\beta}$, are such that:

- 1) $C_{S,ij}^{\alpha E} = 0$ i.e. we assume no specific excitation.
- 2) $C_{S,ij}^{\alpha I} = 0$ if i and j have the same selectivity properties.
- 3) $C_{S,ij}^{\alpha I} = 1$ with probability $2\sqrt{K}/N^I$ if i and j have different selectivity properties.

Therefore, each neuron (excitatory as well as inhibitory) receives, on average, $\sqrt{K}$ connections from inhibitory neurons whose selectivities are different from its own (compared with, on average, K non-selective inhibitory connections).

The strength of the specific connections depends solely on the neurons' type $J_{S,ij}^{\alpha I} = J_S^{\alpha I} C_{S,ij}^{\alpha I}$; $g = J_S^{\alpha I} / J_{NS}^{\alpha I}$.

The total current into neuron (i, α) due to the recurrent interactions is

$$I_{rec,i}^{\alpha}(t) = \sum_{j,\beta} (J_{S,ij}^{\alpha\beta} + J_{NS,ij}^{\alpha\beta}) S_j^{\alpha\beta}(t)$$

where $S_i^{\alpha\beta}(t)$ are synaptic variables, which follow the dynamics

$$\tau_s \frac{dS_i^{\alpha\beta}(t)}{dt} = -S_i^{\alpha\beta}(t) + \sum_{\{t_j^{\beta}\}} \delta(t - t_j^{\beta})$$

Here, τ_s is the synaptic time constant (assumed to be the same for all synapses) and the sum is over all spikes emitted at times $t_j^{\beta} < t$.

Decision-making and decision criterion: In response to the sensory stimulus, the activities of the *U*-selective and *D*-selective neurons change differently (Fig. 5d). We compute at every time step the population-averaged activity of all the excitatory neurons in the set *a*, (*a*-selective), denoted by v_a , $a \in \{U, D\}$, by convolving the spike times with an exponential filter with a time constant of 50 msec. Decision is based on the ratio: $\frac{v_U - v_D}{v_U + v_D}$. If $\frac{v_U - v_D}{v_U + v_D} > \phi$, the decision provided by the network is that upper segment is longer than the lower one, whereas for $\frac{v_D - v_U}{v_D + v_U} > \phi$ it is the opposite, where $\phi > 0$ is the decision threshold.

The ability of the network to make a decision depends on the network parameters. In particular, it depends on the parameter g , which characterize the strength of the competition between '*U*' and '*D*' neurons, on the value chosen for the threshold ϕ as well as on the stimulus parameters, $\bar{I}_0^{\alpha}$ and $\bar{I}_1^{\alpha}$.

Numerical integration: The dynamics of the model circuit were numerically integrated using the Euler method supplemented with an interpolation estimate of the spike times⁵¹. In all simulations the integration time step was 0.1 msec. We verified the validity of the results by performing complementary simulations with smaller time steps.

Model parameters: A systematic study of the dependence of the network dynamics on the parameters is beyond the scope of the present paper. The parameters used in all the simulations are: $V_L = -60\text{mV}$; $\tau_m = 10\text{msec}$; $V_T = 10\text{mV}$; $V_R = -60\text{mV}$; $J_{EE} = 35\text{mV} \cdot \text{ms}$, $J_{IE} = 233.3\text{mV} \cdot \text{ms}$, $J_{EI} = -175\text{mV} \cdot \text{ms}$, $J_{II} = -233.3\text{mV} \cdot \text{ms}$, $I_b^E =$

840mV, $I_b^I = 560\text{mV}$, $I_0^E = 840\text{mV}$, $I_0^I = 560\text{mV}$, $I_1^E = 140\text{mV}$, $I_1^I = 140\text{mV}$, $\tau_s = 3\text{msec}$. The total number of neurons and average non-specific connectivity is $N = 40,000$, $K = 400$, $\phi = 0.4$. The sole parameter we vary is g .

The single-neuron parameters and the average number of inputs per neuron are as in⁴¹. The network size and fraction of inhibitory neurons are as in²⁸. The strengths of $E \rightarrow E$ connections, as well as the unstructured components of $I \rightarrow E$ and $I \rightarrow I$ interactions are as in⁴¹.

The background external inputs as well as the $E \rightarrow I$ connection strength were chosen to obtain spontaneous and evoked firing rates that are comparable with the experimental data. The decision threshold was chosen so that decision occurs only when the differences in the firing rates of the two populations is comparable to the experimentally observed¹⁵.

Relation to other spiking network models: In previous LIF network models of decision making^{30,52}, the recurrent connectivity of the competing populations encoding for the decision outcome is all-to-all and symmetric. In contrast, the connectivity in our network is sparse and random. In these previous studies, irregular firing is due to extrinsic stochastic input. In our model it is generated by the intrinsic recurrent dynamics. Interestingly, it is reported in the Supplementary Methods of⁵² that in their model, "The random connectivity within the sensory circuit and across circuits (i.e. bottom-up and top-down) can cause that the network's behavioral responses exhibit a bias towards one of the two choice options". We hypothesize that the mechanism underlying that bias can be understood using our framework.

DDM analysis

According to the DDM^{13–15,17,18,53–55}, noisy evidence in favor of choosing each of the two alternatives is integrated over the course of the trial. The difference of these evidence, a quantity known as the decision variable, is then computed. Mathematically, $dx/dt = A + \xi$, where x is the decision variable, A is the drift rate, t is time within the trial and ξ denotes white noise such that $E[\xi(t)] = 0$ and $E[\xi(t)\xi(t')] = \delta(t - t')$. In the free-response version of the DDM, which has proven useful for modeling choices even when the stimulus is presented for a fixed duration^{56–58}, a decision is made once the

decision variable reaches one of two decision thresholds, 0 or $a > 0$. The initial condition is set to $x(t = 0) = z \cdot a$, where $0 < z < 1$.

We focus on the impossible trials in which $\Delta L = 0$. Evidence for $A \neq 0$ in those trials is interpreted as drift bias; evidence for $z \neq 0.5$ is interpreted as initial condition bias (IC bias). The two bias mechanisms exhibit distinct patterns of dependence of bias on reaction-times. The effect of 'IC bias' is mostly prominent early in the trial and it therefore predicts that faster decisions are more biased than slower ones. By contrast, 'drift bias' affects evidence accumulation throughout the trial and the resulting bias affects both fast and slow decisions^{19,20}. Therefore, it is possible to dissect the two mechanisms by incorporating the decision times in the analysis.

We fit four different variants of the DDM to the behavioral data and simulations. (1) A baseline DDM with $A = 0$ (because we consider only the impossible trials) and $z = 0.5$. This model has a single parameter, the decision threshold, a . To fit the model to the data, a second parameter, which accounts for the component of the reaction-time that is independent of the decision process, T_{er} , is added^{56,59,60}. By construction, there are no ICBs in this model and it was used as a baseline for comparison with the other three models. (2) In the 'IC bias' DDM $A = 0$, as in the baseline DDM. However, by contrast, z is estimated from the data, this is in addition to a and T_{er} . (3) In the 'drift bias' DDM, we assumed that $z = 0.5$ and estimated A , a and T_{er} . (4) In the 'IC+drift bias' DDM, both z and A were estimated from the data, in addition to a and T_{er} .

Hierarchical Bayesian estimation of the DDM parameters

The dataset of the vertical bisection task includes 20 impossible trials performed by 100 participants. The motor task includes 10 datasets (each pair of dots was considered a task and was analyzed separately). Each task was tested on 20 participants, each performing 20 decisions. The recurrent network simulations included 8 datasets, each corresponding to a different level of specific inhibition. Each of these dataset consisted of the responses made by 200 networks, each tested on 500 impossible trials.

We fit each of the four DDM variants to each of the datasets using the HDDM Python toolbox, which allows for the construction of Bayesian hierarchical DDMs⁶¹. HDDM

uses Bayesian Markov-chain Monte Carlo sampling for generating posterior distributions over both subject-level and group-level model parameters, rather than point estimates of only subject-level parameters. To accomplish this, HDDM uses informative prior distributions on the group-level parameters, that constrain the parameters to a plausible range given past experiments^{61,62}. As a result, constraining the parameter estimates for individual subjects by group-level inference leads to a better recovery of the true parameters, especially with few trials per subject⁶¹.

Our analysis required three minor modifications to the code: (1) In the unbiased and 'IC bias' DDMs, we posit that $A = 0$, which corresponds to an unbiased drift rate. This is because in the bisection task and the recurrent network simulations we only analyzed the case of $\Delta L = 0$. This constraint was lifted in the 'drift bias' and 'IC+drift bias' DDMs, in which A was a free parameter. (2) In the HDDM fitting procedure, the estimation of each of the model parameters is constrained by the informative priors relevant for the group level statistics of the sample's parameter. Specifically, it is assumed that the mean drift rate is drawn from a normal distribution with a positive mean, $m = 2$, conceivably because behavior is typically studied in possible trials, in which performance is above chance. Because in the impossible trials there is no *a-priori* reason to assume that one action is more likely than the other, we modified the code such that $m = 0$. Notably, comparable posteriors are obtained also when using $m = 2$ (not shown). (3) The assumptions regarding the width of the prior distribution of initial conditions can constrain the values of the estimated initial conditions in the HDDM fitting procedure, thus limiting the extent to which the initial conditions can capture the ICBs in the DDM. Therefore, we considered a wider prior distribution of initial conditions, by increasing the standard deviation of σ_z in⁶¹ from 0.05 to 5. Notably, when keeping the standard deviation of σ_z at 0.05, the contribution of initial conditions to the ICBs in the resultant fitted DDMs is even smaller (not shown). As is standard in the HDDM fitting procedure, we allowed 5% of responses to be considered 'contaminants'⁵⁹, *i.e.*, trials which do not follow the DDM dynamics (*e.g.* due to attentional lapses).

In order to estimate the posteriors, we ran 12 separate Markov chains with 40,000 samples each. Of those, the first half was discarded as burn-in and to reduce sample autocorrelations, 4/5 of the remaining samples were discarded for thinning. This left

4,000 samples per chain. We computed the $\hat{R}$ Gelman-Rubin statistic, to assess model convergence by comparing between-chain and within-chain variance of each posterior distribution. For all datasets and all models, the $\hat{R}$ of all group-level posteriors (0.9998-1.01) and that of the observer-level posterior (0.9998-1.035) indicated a proper convergence^{63,64}. All chains were concatenated for further analyses, resulting in 48,000 samples per model, from which each posterior was estimated.

Model comparison using the DIC

Using the HDDM Python toolbox⁶¹, we also computed the Deviance Information Criterion (DIC⁶⁵) and used it to compare the different variants of the DDM. The DIC compares models by the goodness of fit, while penalizing for model complexity. The lower the DIC the better the model (see⁶⁵). Because of the nondeterministic nature of hierarchical modeling, we also computed confidence intervals of the Δ DIC (DIC of the variant of the DDM relative to the DIC of the baseline, unbiased DDM). For the vertical bisection task and the numerical simulations of the recurrent network, the SEM of the Δ DIC was estimated by repeating the fitting procedure and DIC analysis 3 times. For the motor datasets, the Δ DIC of each biased DDM variant was obtained separately for each pair of dots, and the SEM was then evaluated over all 10 pairs.

Posterior-based simulations

The quality of the DDM models can also be evaluated by comparing the behavior of the decision-maker to the behavior predicted by the estimated posteriors. Specifically, we simulated responses (choices and reaction time) using the posteriors obtained from the HDDM procedure for each dataset separately. For all datasets, the simulated probabilities of choice well-matched the observed ones for the ‘drift bias’ and for the ‘IC+drift bias’ DDM variants but not for the ‘IC bias’ DDM (Fig. S8a). All models provided reasonable fits of the normalized distribution of reaction times to the data (Fig. S8b).

Relative contributions of the IC bias and drift bias to the ICBs in the ‘IC+drift bias’ DDM

Here we describe the procedures underlying Figs. 1d, 2d and S7. It is well-known that in the DDM^{17,31}

$$\Pr('Up') = \frac{1}{1+e^{-Aa}} + \frac{1-e^{Aa(1-2z)}}{e^{Aa}-e^{-A}} \quad (6)$$

To dissect the relative contributions of the IC and drift biases to the ICBs, we computed the average parameters A , a and z from the estimated posteriors of each observer in the 'IC+drift bias' DDM. We then computed the predicted $\Pr('Up')$ in three conditions: all estimated parameters (black Xs), estimated initial conditions $z \cdot a$ and $A = 0$ (purple squares) and estimated product, Aa , of drift with the threshold, while assuming $z = 0.5$ (green circles).

Relative contribution of idiosyncratic thresholds to the ICBs

According to Eq. (6), the drift bias A and the threshold a contribute to the ICB via their product Aa . Importantly, while the drift parameter, A , can be positive or negative, the threshold parameter, a , is strictly positive. Therefore, the direction of the bias is necessarily determined by the drift A . Nevertheless, idiosyncrasies in a can also contribute to the heterogeneity in the bias between the decision makers. We studied the relative contributions of these two parameters in the framework of the 'drift bias' DDM, in which the product Aa is the sole contributor to the ICBs. For each decision maker we computed the posterior-averaged values of a and A . We then used Eq. (6) to predict the ICBs assuming that all decision-makers are characterized by the same (average) threshold or the same (average) drift ($|A|$). In all datasets we found that the contribution to the ICBs of heterogeneity in the thresholds is small relative to that of the drift (Fig. S9).

The Poisson model is equivalent to the drift bias DDM

Comparing equations (2) and (6), we note that for $a = 2\sqrt{N} \cdot \tilde{\theta}$, $A = \Delta v / \Sigma v$ and $z = 0.5$, Eq. (6) is equivalent to equation (2).

Data availability

The data and the relevant codes can be downloaded from TBD link.

Acknowledgements

We thank Thomas Boraud, Gianluigi Mongillo, Haim Sompolinsky and Talia Tron for discussions. This work was conducted within the scope of the France-Israel Laboratory of Neuroscience. D. H. thanks the Department of Neurobiology at the Hebrew University for its warm hospitality. This work was supported by the Israel Science Foundation (Y. LO., Grant No. 757/16), the DFG (Y. LO.), the Gatsby Charitable Foundation (Y. LO.), ANR-09-SYSC-002-01 (D. H.) and the France-Israel High Council for Science and Technology (D. H. and Y. LO.).

Author contributions

L. L., Y. LA., D. H. and Y. LO. conceived and planned the experiments; L. L., R. D., D. H. and Y. LO. developed the models; L. L., D. H. and Y. LO. wrote the manuscript.

References

1. Poulton, E. C. *Bias in quantifying judgements*. (Lawrence Erlbaum, 1989).
2. Green, D. M. & Swets, A. J. *Signal detection theory and psychophysics*. (Peninsula Press, 1966).
3. Klein, S. A. Measuring , estimating , and understanding the psychometric function : A commentary. *Attention, Perception, Psychophys.* **63**, 1421–1455 (2001).
4. Daniel, L. & López-Moliner, J. Humans do not evidence choice biases in simple discrimination tasks. *bioRxiv* 62380 (2016).
5. Baum, W. M. On two types of deviation from the matching law: bias and undermatching. *J. Exp. Anal. Behav.* **22**, 231–242 (1974).
6. Laquitaine, S., Piron, C., Abellanas, D., Loewenstein, Y. & Boraud, T. Complex population response of dorsal putamen neurons predicts the ability to learn. *PLoS One* **8**, e80683 (2013).
7. Barraclough, D. J., Conroy, M. L. & Lee, D. Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* **7**, 404–410 (2004).
8. Ashourian, P. & Loewenstein, Y. Bayesian Inference Underlies the Contraction Bias in Delayed Comparison Tasks. *PLoS One* **6**, (2011).
9. Raviv, O., Ahissar, M. & Loewenstein, Y. How recent history affects perception : the normative approach and its heuristic approximation. *PLoS Comput. Biol.* **8**, e1002731 (2012).
10. Yeshurun, Y., Carrasco, M. & Maloney, L. T. Bias and sensitivity in two-interval forced choice procedures: tests of the difference model. *Vision Res.* **48**, 1837–1851 (2008).
11. Skinner, B. F. *The behavior of organisms: an experimental analysis*. (Appleton Century Crofts, 1938).
12. Thorndike, E. L. *Animal Intelligence*. (Macmillan, 1911).
13. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).

14. Ratcliff, R. & McKoon, G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922 (2008).
15. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci* **30**, 535–574 (2007).
16. Kira, S., Yang, T. & Shadlen, M. N. A neural implementation of Wald’s sequential probability ratio test. *Neuron* **85**, 861–873 (2015).
17. Ratcliff, R. A theory of memory retrieval. *Psychol. Rev.* **85**, 59–108 (1978).
18. Ratcliff, R. & Smith, P. L. A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* **111**, 333–367 (2004).
19. White, C. N. & Poldrack, R. A. Decomposing bias in different types of simple decisions. *J. Exp. Psychol. Learn. Mem. Cogn.* **40**, 385–398 (2014).
20. Ratcliff, R. Theoretical Interpretations of the Speed and Accuracy of Positive and Negative Responses. *Psychol. Rev.* **92**, 212–225 (1985).
21. Leite, F. P. & Ratcliff, R. What cognitive processes drive response biases ? A diffusion model analysis. *Judgm. Decis. Mak.* **6**, 651–687 (2011).
22. Mulder, M. J., Wagenmakers, E., Ratcliff, R., Boekel, W. & Forstmann, B. U. Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *J. Neurosci.* **32**, 2335–2343 (2012).
23. Shinomoto, S., Shima, K. & Tanji, J. Differences in spiking patterns among cortical neurons. *Neural Comput.* **15**, 2823–2842 (2003).
24. Hromádka, T., Deweese, M. R. & Zador, A. M. Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol.* **6**, e16 (2008).
25. Buzsáki, G. & Mizuseki, K. The log-dynamic brain: how skewed distributions affect network operations. *Nat. Rev. Neurosci.* **15**, 264–278 (2014).
26. Van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science (80-.).* **274**, 1724–1726 (1996).
27. Roxin, A., Brunel, N., Hansel, D., Mongillo, G. & Vreeswijk, C. Van. On the distribution of firing rates in networks of cortical neurons. *J. Neurosci.* **31**, 16217–16226 (2011).

28. Mongillo, G., Rumpel, S. & Loewenstein, Y. Inhibitory connectivity defines the realm of excitatory plasticity. *Nat. Neurosci.* **21**, 1463–1467 (2018).
29. Amari, S. I. & Arbib, M. A. in *Systems neuroscience* 119–165 (1977).
30. Wang, X. J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
31. Feller, W. *An introduction to probability theory and its applications*. (Wiley, 1968).
32. Smith, P. L. A note on the distribution of response times for a random walk with Gaussian increments. *J. Math. Psychol.* **34**, 445–459 (1990).
33. Broderick, T., Wong-lin, K. F. & Holmes, P. Closed-form approximations of first-passage distributions for a stochastic decision-making model. *Appl. Math. Res. eXpress* **2009**, 123–141 (2009).
34. Urai, A. E., Gee, J. W. De, Tsetsos, K. & Donner, T. H. Choice history biases subsequent evidence accumulation. *bioRxiv* 251595 (2018).
35. De Gee, J. W. *et al.* Dynamic modulation of decision biases by brainstem arousal systems. *Elife* **6**, e23232 (2017).
36. Rorie, A. E., Gao, J., McClelland, J. L. & Newsome, W. T. Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS One* **5**, e9308 (2010).
37. Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* **31**, 6339–6352 (2011).
38. Bernacchia, A., Seo, H., Lee, D. & Wang, X. A reservoir of time constants for memory traces in cortical neurons. *Nat. Neurosci.* **14**, 366–372 (2011).
39. Van Vreeswijk, C. & Sompolinsky, H. Chaotic balanced state in a model of cortical circuits. *Neural Comput.* **10**, 1321–1371 (1998).
40. Renart, A. *et al.* The asynchronous state in cortical circuits. *Science* (80-.). **327**, 587–590 (2010).
41. Darshan, R., Wood, W. E., Peters, S., Leblois, A. & Hansel, D. A canonical neural mechanism for behavioral variability. *Nat. Commun.* **8**, 1–13 (2017).

42. Doiron, B., Litwin-kumar, A., Rosenbaum, R., Ocker, G. K. & Josić, K. The mechanics of state-dependent neural correlations. *Nat. Neurosci.* **19**, 383–393 (2016).
43. Ferster, C. B. & Skinner, B. F. *Schedules of reinforcement*. (Appleton-Century-Crofts, 1957).
44. Mongillo, G., Shteingart, H. & Loewenstein, Y. The misbehavior of reinforcement learning. *Proc. IEEE* **102**, 528–541 (2014).
45. Shteingart, H. & Loewenstein, Y. Reinforcement learning and human behavior. *Curr. Opin. Neurobiol.* **25**, 93–98 (2014).
46. Körding, K. Decision theory: what ‘should’ the nervous system do? *Science* (80-.). **318**, 606–610 (2007).
47. Buchanan, S. M., Kain, J. S. & de Bivort, B. L. Neuronal control of locomotor handedness in *Drosophila*. *Proc. Natl. Acad. Sci.* **112**, 6700–6705 (2015).
48. Horton, J. J., Rand, D. G. & Zeckhauser, R. J. The online laboratory: conducting experiments in a real labor market. *Exp. Econ.* **14**, 399–425 (2011).
49. Jewell, G. & Mccourt, M. E. Pseudoneglect: a review and meta-analysis of performance factors in line bisection tasks. *Neuropsychologia* **38**, 93–110 (2000).
50. Gentet, L. J., Avermann, M., Matyas, F., Staiger, J. F. & Petersen, C. C. H. Membrane potential dynamics of GABAergic neurons in the barrel cortex of behaving mice. *Neuron* **65**, 422–435 (2010).
51. Hansel, D., Mato, G., Meunier, C. & Neltner, L. On numerical simulations of integrate-and-fire neural networks. *Neural Comput.* **10**, 467–483 (1998).
52. Wimmer, K. *et al.* Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat. Commun.* **6**, 6177 (2015).
53. Ratcliff, R., Cherian, A. & Segraves, M. A Comparison of Macaque Behavior and Superior Colliculus Neuronal Activity to Predictions From Models of Two-Choice Decisions. *J. Neurophysiol.* **90**, 1392–1407 (2003).
54. Philiastides, M. G., Ratcliff, R. & Sajda, P. Neural Representation of Task Difficulty and Decision Making during Perceptual Categorization : A Timing Diagram. *J. Neurosci.* **26**, 8965–8975 (2006).

55. Heekeren, H. R., Marrett, S. & Ungerleider, L. G. The neural systems that mediate human perceptual decision making. *Nat. Rev. Neurosci.* **9**, 467–479 (2008).
56. Ratcliff, R. Modeling response signal and response time data. *Cogn. Psychol.* **53**, 195–237 (2006).
57. Ratcliff, R. Continuous versus discrete information processing: Modeling accumulation of partial information. *Psychol. Rev.* **95**, 238–255 (1988).
58. Kiani, R., Hanks, T. D. & Shadlen, M. N. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* **28**, 3017–3029 (2008).
59. Ratcliff, R. & Tuerlinckx, F. Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychon. Bull. Rev.* **9**, 438–481 (2002).
60. Ratcliff, R. & Rouder, J. N. Modeling response times for two-choice decisions. *Psychol. Sci.* **9**, 347–356 (1998).
61. Wiecki, T. V., Sofer, I. & Frank, M. J. HDDM: hierarchical bayesian estimation of the drift-diffusion model in python. *Front. Neuroinform.* **7**, 14 (2013).
62. Matzke, D. & Wagenmakers, E.-J. Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychon. Bull. Rev.* **16**, 798–817 (2009).
63. Gelman, A. & Rubin, D. B. Inference from iterative simulation using multiple sequences. *Stat. Sci.* **7**, 457–472 (1992).
64. Brooks, S. P. & Gelman, A. General methods for monitoring convergence of iterative simulations. *J. Comput. Graph. Stat.* **7**, 434–455 (1998).
65. Spiegelhalter, D. J., Best, N. G., Carlin, B. P. & Van Der Linde, A. Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B (Statistical Methodol.)* **64**, 583–639 (2002).