

<https://doi.org/10.1038/s44271-025-00263-0>

Stability and robustness of idiosyncratic choice bias



Lior Lebovich^{1,2,3}✉, Lea Kaplan¹, David Hansel⁴ & Yonatan Loewenstein^{1,5}

A well-known observation in repeated-choice experiments is that a tendency to prefer one response over the others emerges if the feedback consistently favors that response. Choice bias, a tendency to prefer one response over the others, however, is not restricted to biased-feedback settings and is also observed when the feedback is unbiased. In fact, participant-specific choice bias, known as idiosyncratic choice bias (ICB), is common even in symmetrical experimental settings in which feedback is completely absent. Here we ask whether feedback-induced bias and ICB share a common mechanism. Specifically, we ask whether ICBs reflect idiosyncrasies in choice-feedback associations *prior* to the measurement of the ICB. To address this question, we compared the long-term dynamics of ICBs with feedback-induced biases in two longitudinal experiments involving 319 participants. We show that while feedback effectively induced choice preferences, its effect was transient and diminished within several weeks. By contrast, we show that ICBs remained stable for at least 22 months. These results indicate that different mechanisms underlie the idiosyncratic and feedback-induced biases.

Perceptual decision making is often studied using the 2-alternative forced choice (2AFC) paradigm, in which the experimentalist presents two stimuli that vary along one physical dimension, and the participant is instructed to choose the 'larger' one. In many of these experiments, participants receive feedback after every trial, which can bias their choices. For example, if the feedback indicates that choosing one of the alternatives was the correct response, participants are more likely to choose it again in the following trial. This tendency can be viewed as a form of operant learning, in which participants bias their choices in a direction that they deem more likely to be rewarded. If over many trials, the feedback is asymmetric, consistently favoring one alternative over the other, participants will develop a substantial preference in favor of that alternative^{1–3}.

Somewhat surprisingly, however, participants exhibit preferences in 2AFC experiments, even when the feedback is unbiased. These preferences are idiosyncratic: different participants exhibit a preference in favor of different alternatives and the magnitudes of these preferences vary across participants^{4–13}. These idiosyncratic choice biases (ICBs) are observed even in well-controlled experiments and in absence of any direct feedback¹⁴. Finally, ICBs are not restricted to humans. They are also observed in animals ranging from flies¹⁵ to rodents^{16–20} to monkeys^{21,22}.

Why do participants exhibit these ICBs? One possibility is that ICBs arise from idiosyncratic feedback experienced by participants prior to the experiment. Participants are not tabula rasa at the beginning of the experiment. Their life experiences include interactions with the environment that are followed by feedback. According to this hypothesis, ICBs reflect these past experiences and their influence on decision-making processes. In a previous study¹⁴ we put forward an alternative hypothesis: ICBs may arise from microscopic heterogeneities in the neural circuits involved in decision-making. This hypothesis is supported by our theoretical results showing that ICBs emerge when we simulate decision-making in a completely symmetric setting by a cortical-like network, in which the two alternatives are represented by two, statistically identical populations of neurons. These findings suggest that ICBs are an almost inevitable feature of neural computation, unless they are actively suppressed.

Testing these hypotheses directly is challenging. Accessing the complete life-history of participants is impossible, and even if such data were available, identifying the specific life experiences relevant to a given two-alternative forced choice (2AFC) task would be extraordinarily difficult. Similarly, we lack the means to observe the microscopic organization of the neural circuits underlying decision-making.

¹The Edmond and Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem, Jerusalem, Israel. ²Department of Collective Behavior, Max Planck Institute of Animal Behavior, Konstanz, Germany. ³Centre for the Advanced Study of Collective Behaviour, University of Konstanz, Universitätsstraße 10, 78464 Konstanz, Germany. ⁴Cerebral Dynamics, Plasticity and Learning Laboratory, CNRS, 45 Rue des Saints Pères, 75270 Paris, France. ⁵The Alexander Silberman Institute of Life Sciences, Dept. of Cognitive and Brain Sciences and The Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem, Israel. ✉e-mail: lebovich.lior@gmail.com; yonatan@huji.ac.il

Here we adopt an indirect approach and compare the long-term stability of ICBs with those of experimentally-induced choice biases. If both reflect the same underlying mechanism, we expect them to exhibit similar stability. If, instead, their stability differs, this would suggest that different mechanisms underlie idiosyncratic and feedback-induced biases.

Methods

The experiments were approved by the Hebrew University Committee for the Use of Human Subjects in Research and all participants provided informed consent. Recruitment was based on the online labor market Amazon Mechanical Turk. Each of the studies was described as a longitudinal academic survey of visual acuity. Participants received a link to the experimental website. Response data in each trial was stored in a secured SQL database. These studies were not preregistered. Demographic details, including participants' sex, were determined based on self-report.

Study 1: Stability of ICB in absence of feedback

Human participants were invited to take part in a two-session experiment. Upon the completion of the first session (session 1), they received a link indicating the time-window of the second session. Of a total of 245 participants, 62 who failed to start the second session (session 2) within their designated time window were excluded. This resulted in a final sample size of 183 participants, as described in Table 1. In addition, 88 of those who completed session 1 and session 2 in 8 days or less were invited to participate in a third session (session 3), 22 months later; 29 of them completed this session (see: Table 1, bottom row). None of the participants that started session 2 on time were excluded from the analysis.

A base monetary compensation was given to all participants, separately in each session. To encourage good performance, an additional bonus fee was given for every correct response, and another bonus was guaranteed to 10% of participants with highest scores in each session. All participants were in the United States of America and all reported normal or corrected to normal vision.

Experimental procedure. All sessions followed the same procedure. In each trial, participants were instructed to make a perceptual discrimination decision. There were two types of perceptual tasks, a vertical bisection task, analyzed in the Results section and a horizontal, size-comparison task, which is mentioned in the Discussion section and the Supplementary Information (Supplementary Figs. S4–S10 and Supplementary Tables S1, S2). Participants were instructed to make their decision as quickly and as accurately as possible.

In a bisection trial, a 200 pixel-long vertical white line, transected by a horizontal 20 pixel-long white line was presented on a black screen and participants were instructed to indicate which of the two segments is longer (Fig. 1A, inset). In a size-comparison trial, two white Gaussian blur circles with an average radius of 75 pixels were presented on a black screen and participants were instructed to indicate which circle out of the two is bigger (Supplementary Fig. S4A, inset). In both cases, participants were instructed to press the spacebar key once coming to a decision. Upon the spacebar

press, the stimulus was replaced by a decision screen composed of two arrows buttons, appearing on opposite sides of the screen, and a middle 4-squares submit button. In the bisection task, the two arrows were either 'up' or 'down', whereas in the size-comparison task, they were either 'left' or 'right'. The participants indicated their decision by moving the initially centered cursor to one of the arrow buttons, pressing it, and finalizing their decision by pressing the 'submit' button. Upon the 'submit' button press, the decision screen was replaced by a black screen and the next trial began after 500 milliseconds.

The stimuli were limited to a 400-pixel × 400-pixel square at the center of the screen. The horizontal location of all vertical bisection lines and the vertical location of the center of all circles were centered. 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.

No feedback was given regarding the correct response. The participants were, however, informed about the accumulated bonus fee every 30 trials.

Each session consisted of 480 trials, 240 vertical bisection (B) and 240 horizontal size-comparison (C) trials. Trials in each session were ordered in 160 alternating blocks of 3 bisection and 3 size-comparison trials (...CCCBBCCCBBB...). Unknown to the participants, there were 40 impossible bisection and 40 impossible size-comparison trials (1/6 of the trials) in each session. To minimize sequential effects in the impossible trials, each impossible bisection trial was preceded by three size-comparison trials. Similarly, each impossible size-comparison trial was preceded by three bisection trials. For the possible bisection trials, the deviation from the veridical midpoint was uniformly distributed between 2, 4, 6, 8 and 10 pixels. The quantity L^U and L^D denoting the lengths of the 'up' and 'down' segments of the vertical line, this corresponds to $\frac{\Delta L}{\Sigma L} = \frac{(L^U - L^D)}{(L^U + L^D)}$ between 0.02 and 0.1, with an equal number of offsets in each direction. For the possible size-comparison trials, the difference in radii between the two circles was uniformly distributed between 4, 8, 12, 16 and 20 pixels. Defining $\frac{\Delta R}{\Sigma R} = \frac{(R^R - R^L)}{(R^R + R^L)}$ with R^R and R^L the radii of the right and left circles, this corresponds to $\frac{|\Delta R|}{\Sigma R}$ between 0.0267 and 0.1333, with an equal number of deviations in each direction. To minimize external causes of heterogeneity, the order of the trials was identical for all participants.

To verify that the participants understood the instructions, they were required to successfully complete a bisection practice session and a size-comparison practice session prior to the experiment. This session consisted of blocks of 4 easy trials with feedback and balanced polarity of ΔL or ΔR . In each session, the main experiment started after the participant completed one bisection and one size-comparison block successfully. Responses in this practice session were not included in the analysis.

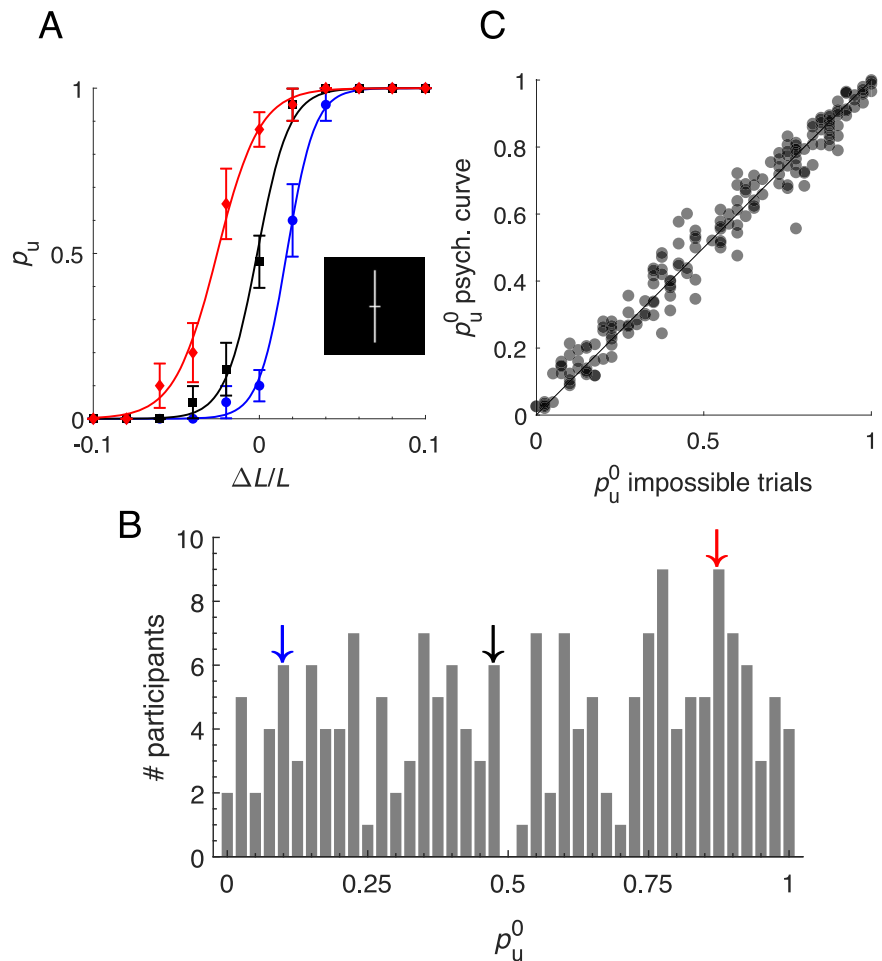
Mean performance in the possible bisection trials was $92.1\% \pm 5.3\%$ (standard deviation), range 62% – 99.5% in the first session and $92.6\% \pm 5.4\%$ (standard deviation), range 66% – 99.5% in the second session. Mean performance in the possible size-comparison trials was $98.6\% \pm 1.8\%$ (standard deviation), range 89.5% – 100% in the first

Table 1 | Study 1 – Participants included in the analysis

Time between sessions		n (#female participants + #male participants)	# sinistral	Age on session 1 (years)		Timing
Mean	Range			Mean	Range	
1.7 h	0.3–2.6 h	29 (19 f + 10 m)	3	39.9	24–67	July 2019
26.3 h	23.1–42.9 h	32 (24 f + 8 m)	4	41.0	28–69	July–August 2019
6.7 days	6.0–7.9 days	27 (13 f + 14 m)	3	40.4	24–66	August 2019
28 days	27–31.9 days	33 (18 f + 15 m)	1	38.8	22–61	August–September 2019
88.5 days	85.9–97.1 days	30 (18 f + 12 m)	3	39.6	25–65	August–November 2019
240.1 days	232.2–247.2 days	32 (15 f + 17 m)	2	40.9	24–69	October 2019–August 2020
665.7 days	656.7–672.1 days	29 (22 f + 7 m)	2	42.3	28–66	July 2019– May 2021

Fig. 1 | ICBs in the vertical bisection task.

A Psychometric curves of three participants. The observed fraction, p_u , of responding ‘up’ is plotted as a function of the sensory offset, $\frac{\Delta L}{L}$. Error bars denote the estimated standard error (SE) of the proportion. Solid lines are best-fit logistic functions. One participant did not exhibit any significant bias (black), one was biased in favor of responding ‘up’ (red) and one (blue) against it. Inset: the stimulus in a single trial. **B** The distribution across the participants ($n = 183$) of p_u^0 , the value of p_u , in trials in which $\Delta L = 0$. Arrows denote p_u^0 for the three participants in panel A, color coded. **C** p_u^0 , estimated using a ML fit of a logistic function to the psychometric curve, vs. p_u^0 , measured in the impossible trials only. The correlation coefficient is 0.98. All measurements are from Session 1 of Study 1.



session and $98.3\% \pm 3.6\%$ (standard deviation), range 58% – 100% in the second session.

Study 2: Feedback-induced bias and its stability

This study was like study 1 in terms of tasks, stimuli and instructions, with one major exception: it consisted of intermediate feedback third session (session 3). Participants were told in advance that the strict time window for starting session 2 is 24–54 h after completion of session 1. In addition, participants were told in advance that they will be informed about the time window for starting session 3 only during session 2. The time window could either be 24–72 h or 28–35 days after completion of session 2. This allowed us to match participants into groups based on their ICBs in session 1 (see below) to create groups that are as possible similar in their initial biases. Overall, data in session 1 were collected from 148 participants. We excluded from the analysis those participants who failed to start session 2 and session 3 within their designated time windows (see below, 12 participants, 8.1% dropout). This resulted in a final sample of 136 participants.

After completing session 1 and before starting session 2, all participants were divided into groups that later defined both their feedback manipulations in session 2 and the time window between session 2 and session 3. The feedback manipulations in session 2 were either ‘enhance up’ or ‘enhance down’ for the bisection task and either ‘enhance right’ or ‘enhance left’ for the size-comparison task. Hence, each participant was associated with one bisection manipulation group and with one size-comparison manipulation. Further, half of the participants waited 1–3 days between session 2 and session 3 whereas the other half waited 28–34 days. Thus, there were 8 groups of participants. Because our studies primarily focused on the bisection task, participants were first divided into the bisection

manipulation and delay time groups according to their p_u^0 , as measured from the first session, such that each of the four groups (‘enhance up’, ‘enhance down’ \times 1-day, 1-month) had as similar distribution of initial p_u^0 biases as possible. Specifically, we first sorted in ascending according to p_u^0 of all the 148 participants in session 1. These were further divided into 37 quartets, each quartet consisting of 4 participants with consecutive p_u^0 s. Following, the participants in each quartet were randomly divided between the 4 groups of bisection ‘enhance’ polarities and delay times (1 participant per group). This routine was repeated until all participants were assigned to a bisection ‘enhance’ and a time delay group. Only then were the participants in each of the 4 groups above further divided into size-comparison manipulation groups, based on matching their p_r^0 , defined as the fraction of responses that the circle to the right was larger, in session 1. Table 2 provides a summary for each group.

As in study 1, participants received a link with a countdown for the time to start the next session and were informed about the time window. Once the countdown was over, participants were automatically invited to take part in the next session. Base monetary compensation and bonuses were given to all applied participants in each session (separately). In addition, to encourage participants to take part in all three sessions, an additional bonus was given to those who completed session 3. As in study 1, all participants were in the United States of America. They reported normal or corrected to normal vision.

Experimental procedure. The sessions in this study consisted of the bisection and size-comparison tasks described above for study 1, with the same number of impossible and possible bisection and size-comparison with one exception. For the possible size-comparison trials, the difference in radii between the two circles was smaller than in study 1 (trials were

Table 2 | Study 2 – Participants included in the analysis

Feedback manipulation in session 2		Time between sessions 1 and 2		<i>n</i> (#female participants + #male participants)	# sinistral	Age on first session (years)		Timing
		Mean (days)	Range (days)			Mean	Range	
'Enhance up'	'Enhance right'	1.2	1–1.9	14 (6 f + 8 m)	2	40.1	25–70	June 2021
		29.3	28–33.9	18 (6 f + 12 m)	3	36.5	19–54	June–July 2021
	'Enhance left'	1.1	1–1.8	20 (8 f + 12 m)	3	40.8	26–60	June 2021
		28.5	28–29.9	17 (7 f + 10 m)	2	42.2	23–74	June–July 2021
'Enhance down'	'Enhance right'	1.2	1–2.1	18 (6 f + 12 m)	5	38.3	22–68	June 2021
		28.5	28–31.3	17 (7 f + 10 m)	6	44.4	27–72	June–July 2021
	'Enhance left'	1.2	1–2.9	16 (7 f + 9 m)	2	41.5	25–60	June 2021
		29.5	28.2–32.9	16 (8 f + 8 m)	0	39.7	28–55	June–July 2021

more difficult): uniformly distributed between 2, 4, 6, 8 and 10 pixels ($\frac{\Delta R}{\Sigma R}$ between 0.0133 and 0.0667).

Session 1 and session 3 followed the same procedure as in study 1. Session 2 differed in that after each trial, participants received feedback about their response, a smiley or a sad emoji indicating whether their choice was correct or incorrect. The feedback was accurate in all possible trials ($\Delta L \neq 0$ and $\Delta R \neq 0$, 5/6 of the trials). By contrast, in the impossible trials, different feedback schedules were used. For participants in the 'enhance up' group, the 'up' response was considered 'correct' in 95% of the impossible trials, whereas for the 'enhance down' group, the 'down' response was considered 'correct' in 95% of the impossible trials. For participants in the 'enhance left' group, the 'left' response was considered 'correct' in 95% of the impossible trials, whereas for the 'enhance right' group, the 'right' response was considered 'correct' in 95% of the impossible trials (see Table 2).

All the data used for the statistical analysis was binary choices. The assumption in all statistical analyses is that these choices follow a binomial distribution. This was not formally tested here, but in a previous study that used the same experimental design, we tested for a repetition or alternation tendencies but did not find any, consistent with this assumption (Supplementary Fig. S2 in ref. 14).

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Results

Characterization of the ICBs

We used the vertical bisection discrimination task, depicted in Fig. 1A (inset) to study ICBs (Methods). In each trial, a vertical transected line was presented on the screen, and participants were instructed to indicate which of the two vertical segments was longer. No trial-by-trial feedback was given to the participants ($n = 183$) about their answers. We computed the fraction of responses in which the participants reported that the upper segment of the transected line was longer than the bottom segment, a quantity which we denote by p_u , separately for each participant and offset. Figure 1A depicts p_u for three participants as a function of the offset. As expected, p_u increased with the offset, $\frac{\Delta L}{\Sigma L} \equiv \frac{(L^u - L^d)}{(L^u + L^d)}$, where L^u and L^d denote the lengths of the up and down segments of the vertical line. Notably, the participant denoted by a red line was biased in favor of responding 'up' even in the trials in which ΔL was negative (but small). By contrast, the psychometric curve of the participant denoted by the blue curve is shifted to the right, indicating a bias in favor of responding 'down'. To quantify the idiosyncratic preferences towards and against responding 'up', we focused on the choices of the participants in 'impossible' trials, in which the line was transected at its veridical midpoint ($\Delta L = 0$, 40/240 trials). We denote by p_u^0 the fraction of 'up' responses of a participant in these trials. While the participant whose

psychometric curve is in black in Fig. 1A did not exhibit any significant bias ($p_u^0 = 19/40$ trials, $p = 0.87$, two-sided Binomial test, CI 0.32–0.64, Clopper–Pearson), the participants whose psychometric curves are plotted in red and blue were significantly biased ($p_u^0 = 35/40$ and $p_u^0 = 4/40$, respectively; in both cases $p < 0.001$, two-sided Binomial tests, CI 0.73–0.96 and 0.03–0.24, respectively, Clopper–Pearson).

The distribution of p_u^0 is depicted in Fig. 1B. Across the participants, 69% exhibited a significant bias (39% with $p_u^0 > 0.5$ and 30% with $p_u^0 < 0.5$; $p < 0.05$, two-sided Binomial tests, not corrected for multiple comparisons). The mean of the distribution, 0.53, is not significantly different from chance (95% CI 0.49–0.58, bootstrap). The variance of the distribution was significantly larger than expected by chance ($p < 0.001$, two-sided bootstrap test, Bernoulli process; $n = 183$ participants, 40 trials per participant) further establishing the existence of ICBs in the task. Finally, the ICBs were not restricted to the impossible trials. Rather, they manifested as idiosyncratic lateral shifts of the psychometric curve. Across the participants, the probability of choosing 'up' in 'possible trials' was highly correlated with that probability in the impossible trials (Supplementary Fig. S1; two-sided Pearson's $\rho = 0.82$, $p < 0.001$, degrees of freedom = 181, 95% CI: 0.77–0.87). We also estimated the bias by fitting a logistic function to the psychometric curves. It was correlated with the bias estimated from the impossible trials (Fig. 1C, Correlation coefficient 0.98).

ICBs are stable over many months

To study how ICBs change over time, the ICB of each of the participants depicted in Fig. 1B was measured again later (session 2 in Fig. 2A; also see: Methods). One subgroup of 29 participants was also tested 22 months after the second session (session 3).

Figure 2B, plots p_u^0 in session 1 vs. its value 8 months later (32 participants, Supplementary Fig. S2 for all other groups). It shows that p_u^0 hardly changed during that delay. Quantifying this change, the correlation between the values of p_u^0 in the two sessions 8 months apart, was $\rho = 0.93$ (Pearson, degrees of freedom = 30, 95% CI: 0.86–0.96). Figure 2C depicts the correlation between the values of p_u^0 in the two sessions also 1 h apart ($\rho = 0.81$, Pearson, degrees of freedom = 30, 95% CI: 0.63–0.91), 1 day apart ($\rho = 0.76$ Pearson, degrees of freedom = 30, 95% CI: 0.59–0.89), 1 week apart ($\rho = 0.82$ Pearson, degrees of freedom = 30, 95% CI: 0.64–0.91), 1 month apart ($\rho = 0.84$ Pearson, degrees of freedom = 30, 95% CI: 0.71–0.92), 3 months apart ($\rho = 0.73$ Pearson, degrees of freedom = 30, 95% CI: 0.51–0.86) and 22 months apart ($\rho = 0.67$ Pearson, degrees of freedom = 30, 95% CI: 0.40–0.83). To estimate the dynamics of ICB over time, we fitted the correlations ρ with delay times using an exponential function, $\rho = \rho_0 e^{-\frac{\Delta t}{\tau}}$. The estimated decay time τ was $\tau = 170$ months, (95% CI: $42 - 1.3 \cdot 10^6$, bootstrapping). We therefore concluded that the participants retained their ICB even after many months.

Fig. 2 | ICB is stable over months. **A** The experimental design. The ICBs of $n = 183$ participants were measured in the first session (session 1). Subsets of the participants were tested again in an identical session (session 2) 1 h ($n = 29$), 1 day ($n = 32$), 1 week ($n = 27$), 1 month ($n = 33$), 3 months ($n = 30$) or 8 months later ($n = 32$). A subset ($n = 29$) also participated in a third session (session 3), 22 months after session 2 (dashed lines). **B** ICB is stable over 8 months. Dots: participants tested in session 2, 8 months after session 1. The diagonal is plotted for comparison. C ICBs are stable over time. Gray bars: Pearson correlation of p_u^0 between session 1 and 2 (circles) and between session 2 and session 3 (star). Error bars: 95% confidence interval, bootstrap. Magenta area: 95% confidence interval expected under complete stability, bootstrap Bernoulli processes.

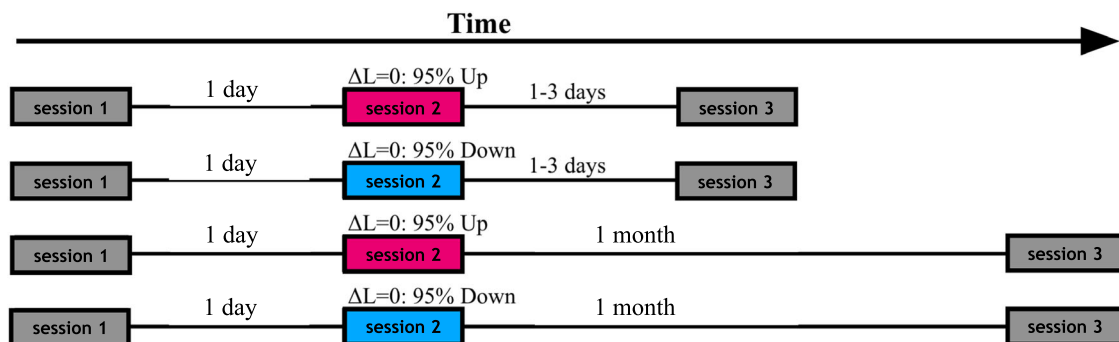
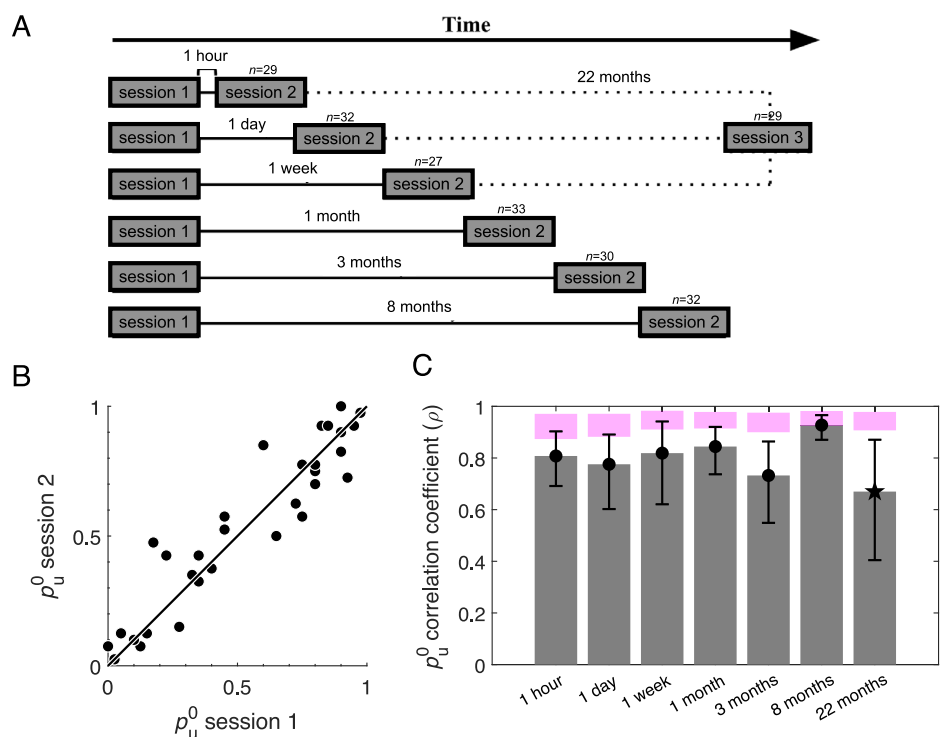


Fig. 3 | The experimental design to investigate the dynamics of feedback-induced bias. All participants ($n = 136$) participated in three sessions. In sessions 1 and 3 (gray) no trial-to-trial feedback was provided. In session 2, choices in all trials were followed by feedback. Feedback in 200 possible trials was accurate. Impossible trials (40 trials) were followed by asymmetrical feedback about the ‘correct’ response. Pink: ‘enhance up’ condition. The feedback indicated that the ‘correct’ response is ‘up’ in 38 impossible trials and ‘down’ in 2 impossible trials. Blue: ‘enhance down’

condition. The feedback indicated that the ‘correct’ response is ‘up’ in 2 impossible trials and ‘down’ in 38 impossible trials. All participants performed session 2 one day after their completion of session 1. Half of them participated in session 3 one day after completing session 2 (top) and the other half participated in session 3, 1 month after completing session 2 (bottom). Participants were divided into the 4 groups (2 feedback \times 2 delay) by matching their ICBs in session 1 (Methods and Supplementary Fig. S3B).

The dynamics of feedback-induced bias

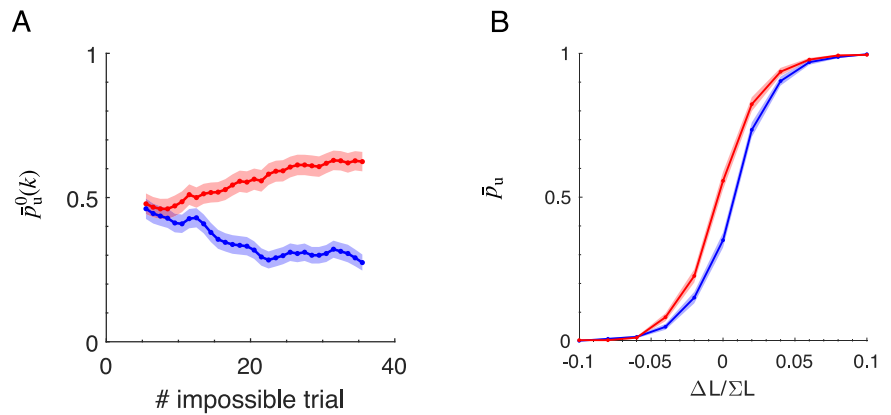
If the observed ICBs reflect idiosyncratic histories of operant learning, we expect comparable stability for choice biases induced by providing proper feedback to the participants. Therefore, we set out to study the dynamics of such feedback-induced effect. To that goal, we recruited a group of $n = 136$ new participants and tested each one in three experimental sessions (Fig. 3). In session 1 we measured their ICBs without providing trial-by-trial feedback. This was not the case in the second session (session 2), 1 day later, where feedback was provided. In session 3, either 1 day later or 1 month later, we again measured their ICBs without trial-by-trial feedback.

As expected, since session 1 was identical to the first session of Fig. 2A, the distribution of ICBs over the participants was comparable in these two situations (Supplementary Fig. S3A). While in session 2, participants were presented with the same stimuli and the same task as in session 1, they also received, in every trial, feedback indicating whether their choice was correct

(happy emoji) or incorrect (sad emoji). While the feedback was accurate in all possible trials ($\Delta L \neq 0$; 200/240 of the trials), it was biased in the impossible trials ($\Delta L = 0$; 40/240 of the trials), to systematically influence the participants’ preferences. Specifically, participants were divided into two groups with different feedback schedules. In one group, ‘up’ response was considered ‘correct’ in 95% of the impossible trials while ‘down’ response was considered ‘correct’ in 5% of the trials. We refer to this group of participants as ‘enhance up’. In the second group, ‘enhance down’, the feedback was biased in the opposite direction: ‘up’ response was considered ‘correct’ in 5% of the impossible trials and ‘down’ response was correct in the other 95%.

So far, we separately considered each participant and averaged their responses over all the trials. To study how the feedback in session 2 changed the tendency to choose ‘up’ over trials, we separately considered the impossible trials and averaged the responses over the participants. To clarify the difference between averaging over responses given a participant and

Fig. 4 | Feedback biases choices in session 2. A The participant-averaged bias in the impossible trials, $\bar{p}_u^0(k)$, smoothed with a sliding window of 10 trials. Red and Blue: ‘enhance up’ and ‘enhance down’ groups of participants ($n = 69$ and $n = 67$ participants, respectively). Shaded area: SEM. **B** The average (over participants) of the psychometric curve \bar{p}_u and SEM (shaded area) as a function of the relative difference between the length of the segments, $\frac{\Delta L}{\Sigma L}$. Colors and sample sizes as in (A).



averaging over participants given a trial, we denote by $a_{u,i}^0(k)$ the response of participant i in impossible trial k : $a_{u,i}^0(k) = 1$ if participant i chose ‘up’ in impossible trial k ($k \in \{1, \dots, 40\}$), and $a_{u,i}^0(k) = 0$ otherwise. The fraction of impossible trials in a session in which a participant i chose ‘up’ is $p_{u,i}^0 = \frac{1}{40} \sum_{k=1}^{40} a_{u,i}^0(k)$. Note that the quantity p_u^0 considered so far for each participant is in fact $p_{u,i}^0$, where the index i was omitted for simplicity of notations.

To study how feedback changed the tendency of choosing ‘up’ over trials, we averaged over participants, rather than over trials. We computed $\bar{p}_u^0(k) = \frac{1}{n} \sum_{i=1}^n a_{u,i}^0(k)$, where n denotes the number of participants in each ‘enhance’ group. The red and blue lines in Fig. 4A depict a running average of $\bar{p}_u^0(k)$ for the ‘enhance up’ and the ‘enhance down’ groups, respectively. Initially, the two groups exhibited comparable population averages. However, with trials, the effect of the feedback became more pronounced. These results are consistent with previous studies that demonstrated that feedback in even a small number of impossible trials can substantially bias participants’ choices¹.

In our experiment, the biased feedback (95%; 5%) was applied only to the impossible trials, while it was correct and unbiased (50%; 50%) in the possible trials. A previous study has argued that the effect of biased feedback is selective to the specific stimuli that were associated with that feedback². On the other hand, in a previous study that utilized auditory stimuli we showed that feedback in impossible trials also affects those which are possible¹. To test whether the feedback also affects choices in the possible trials in our bisection task, we compared the average psychometric curve of the ‘enhance up’ (red) and ‘enhance down’ (blue) participants in the second session (Fig. 4B). We found that the biased feedback shifted the entire psychometric curve, indicating that it also affected the possible trials.

If ICBs are the result of idiosyncratic feedback prior to the experiment, we predict the feedback-induced bias and ICBs to be equally stable. We therefore asked: how long would the feedback-induced bias in sessions 2 affect the bias of the participants? To address this question, half of the participants (68/136) were tested 1 day after the manipulation whereas half of the participants were tested after 1 month (session 3). As in session 1, session 3 was devoid of any trial-to-trial feedback.

To combine in our analysis the ‘enhance up’ and ‘enhance down’ groups, we now consider actions according to their congruency with the feedback. Specifically, $a_{c,i}^0(k) \equiv a_{u,i}^0(k)$ for participants in the ‘enhance up’ group and $a_{c,i}^0(k) \equiv 1 - a_{u,i}^0(k)$ for participants in the ‘enhance down’ group. To quantify the global bias of a participant in a session, we averaged this quantity over the trials in that session. Because it took several trials for the bias to develop (Fig. 4A), we used the average of $a_{c,i}^0(k)$ in the second half of the session 2, $p_{c,i}^0(\text{session } 2) \equiv \frac{1}{20} \sum_{k=21}^{40} a_{c,i}^0(k)$ as a metric of the feedback-induced bias. For the bias session 3, we included all impossible trials: $p_{c,i}^0(\text{session } 3) \equiv \frac{1}{40} \sum_{k=1}^{40} a_{c,i}^0(k)$. Omitting the index i for simplicity of notations, stability of preference would manifest as $p_c^0(\text{session } 3) =$

$p_c^0(\text{session } 2)$ whereas $p_c^0(\text{session } 3) \neq p_c^0(\text{session } 2)$ would imply that the preference of the participant has changed.

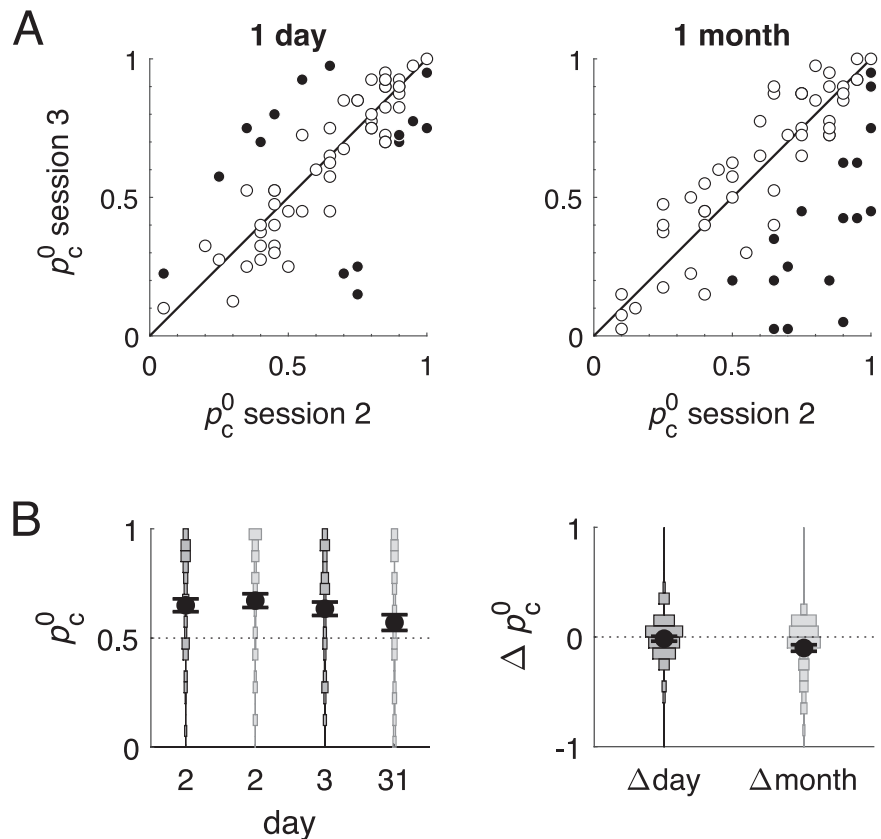
In our experiment, the delay between session 3 and session 2 was 1 day for one group of participants and 1 month for the other. Figure 5A depicts $p_c^0(\text{session } 3)$ vs. $p_c^0(\text{session } 2)$ for each participant in the two groups. After 1 day, 25% of the participants (17/68) exhibited a significant change in p_c^0 (closed circles in Fig. 5A, left; significance threshold $p < 0.05$, two-sided two proportion z-tests, not corrected for multiple comparisons). Of those, the number of participants who increased ($n=7$) and the numbers of participants that decreased their biases ($n=10$) were comparable ($p=0.63$, two-sided Binomial tests, CI 0.33–0.82, Clopper–Pearson). By contrast, in the one-month group, 29% of the participants (20/68) exhibited a significant change in their bias (closed circles significance threshold $p < 0.05$, two-sided two proportion z-tests, not corrected for multiple comparisons), all of them decreased it ($p < 0.001$, two-sided Binomial tests, CI 0.83–1, Clopper–Pearson). These results indicate that the feedback-induced bias has significantly decayed within a month.

To further quantify this decay, we considered the population-average of the bias. We define $\bar{p}_c^0(\text{session } x) \equiv \frac{1}{n} \sum_{i=1}^n p_{c,i}^0(\text{session } x)$, where $x \in \{2, 3\}$. A decay of the feedback-induced bias would manifest as $\bar{p}_c^0(\text{session } 3) - \bar{p}_c^0(\text{session } 2) < 0$ (Fig. 5B). For the 1-day group, $\bar{p}_c^0(\text{session } 3) - \bar{p}_c^0(\text{session } 2)$ was not significantly different from 0 (mean \pm SE: $\bar{p}_c^0(\text{session } 3) - \bar{p}_c^0(\text{session } 2) = -0.016 \pm 0.022$, $p = 0.43$, two-sided Wilcoxon signed-rank test, $z = 0.79$). By contrast, for the 1-month interval group, $\bar{p}_c^0(\text{session } 3) - \bar{p}_c^0(\text{session } 2)$ was significantly smaller than 0 (mean \pm SE: $\bar{p}_c^0(\text{session } 3) - \bar{p}_c^0(\text{session } 2) = -0.10 \pm 0.03$, $p = 0.013$, two-sided Wilcoxon signed-rank test, $z = 2.48$). These results confirm that the feedback-induced bias which persists over 1 day, is unstable over 1 month.

To quantify the decay rate of the feedback induced bias, we considered the change in the global bias of the 1-month group from session 2 to session 3. Assuming that the decay is exponential, we can use the metric $(\bar{p}_c^0(\text{session } 3) - 0.5) / (\bar{p}_c^0(\text{session } 2) - 0.5)$ to quantify the decay time constant. We found $(\bar{p}_c^0(\text{session } 3) - 0.5) / (\bar{p}_c^0(\text{session } 2) - 0.5) = 0.41 \pm 0.18$ (mean \pm bootstrap SE) corresponding to a decay time constant of 1.1 months (the bootstrap analysis shows that with 95% confidence, the decay time-constant is smaller than 2.6 months). Assuming this value, the feedback-induced bias is expected to diminish to 1% of its original value by 5 months. By contrast, the choice bias in the previous experiment remained stable for many more months (Fig. 2). Together, these results indicate that the feedback-induced bias is substantially less stable than the ICB observed in session 1.

To what value does the bias converge to at the single participant level? One possibility is that the feedback somehow resets the ICB. In that case, the correlation between the biases in sessions 1 and 3 will not increase with the time duration between sessions 2 and 3. Alternatively, the contribution of

Fig. 5 | The feedback-induced bias is stable over one day but not over one month. **A** The value of p_c^0 (the bias congruent with the feedback of session 2) in session 3 as a function of p_c^0 in session 2, in the one-day (left, $n = 68$ participants) or one-month (right, $n = 68$ participants) groups. Each dot corresponds to one participant. Filled circles: participants for which the change in p_c^0 was statistically significant. Open circles: participants with non-significant changes. In the 1-month group, 29% of the participants ($n = 20$ participants) exhibited a significant change in their bias with a decrease for all of them. **B** Left, the distribution over the population of p_c^0 in session 2 (two left bars) and session 3 (two right bars) for the one-day interval group (dark gray) and one-month interval group (light gray). Right, the distribution of the change in the bias congruent with the feedback of session 2, $p_c^0(\text{session 3}) - p_c^0(\text{session 2})$. For the one-month interval group (light gray), the population average $\bar{p}_c^0(\text{session 3}) - \bar{p}_c^0(\text{session 2})$ was significantly negative ($p = 0.013$, two-sided Wilcoxon signed-rank test, $z = 2.48$). Dots and error bars denote the mean and SEM.



the feedback provided in session 2 will diminish with time. In that case, the choice bias is predicted to revert with time to an ICB like the one in session 1.

Figure 6A plots the value of p_u^0 in session 3 vs. its value in session 1 for all the participants in the 1-day and 1-month groups (red and blue, ‘enhance up’ and ‘enhance down’ groups, respectively). Inspection shows that the dots are closer to the diagonal in the 1-month group suggesting that with time choice biases revert to their original value. To quantify this effect, we computed the Pearson correlation of p_u^0 in sessions 1 and 3. Figure 5B shows that this correlation is significant for both groups. However, it is smaller in the 1-day interval group compared with the 1-month interval group (Fig. 6B; $p = 0.04$, two-sided z-test for the difference in correlation after Fisher z-transformation). We thus conclude that at the single participant level the feedback-induced bias is transient, unlike the ICBs of session 1.

Discussion

We studied the long-term dynamics of human ICBs. We showed that in the absence of feedback, these ICBs are remarkably stable for many months. Moreover, we found that while they can be temporarily shaped using feedback, this effect is transient: the ICB exhibited before the feedback manipulation recovers within weeks. Together, our results challenge the hypothesis attributing ICBs to the idiosyncratic history of choices and reinforcers.

The generality of the results

The focus of this paper was the vertical bisection trials. This is because previous studies have reported global biases in horizontal comparison tasks, which were attributed to pseudoneglect²³. To test the generality of our results we interleaved the vertical bisection trials with trials of a different perceptual task. In the latter, participants were instructed to compare the sizes of two, horizontally displaced Gaussian-blurred circles (see Methods and Supplementary Fig. S4A, inset). Like the vertical-bisection task, 1/6 of the circle-comparison trials were impossible, and feedback about the correct answer followed the same schedule as in the vertical-

bisection trials. The two types of trials were interleaved such that participants were presented with alternating blocks of three vertical bisection (B) trials and three horizontal comparison (C) trials such that the impossible trials (B⁰ and C⁰, respectively) were positioned as the first trial in a block (...BBBCCCC⁰BBC⁰CC...). This design was chosen as to minimize sequential effects in the impossible trials, which could confound our estimate of the ICBs in the impossible trials. This design also allowed us to test whether ICBs in the circle-comparison task are comparable to those of the bisection task and whether they follow the same dynamics following biased feedback. We replicated the main results of the paper in the circles’ task, namely, participants exhibited substantial ICBs, which were stable over many months. Feedback temporarily influenced these ICBs, but the induced effect was short-lived and decayed over time (Supplementary Figs. S4–S10). This suggests that our main findings can be generalized to other perceptual comparison tasks.

Relation to previous studies

Stability of ICBs has been demonstrated in humans over time intervals varying from days (e.g. ref. 24) to weeks (e.g. refs. 11,13,25,26). In refs. 11,25,26, repeated measurements were used to quantify the stability of the bias, a procedure which may have a stabilizing effect on the bias¹¹. Our study demonstrates that ICBs can be stable for much longer periods, at least two years, even without multiple measurements.

Of relevance to our work is a human study that used ambiguous stimuli to measure the stability over time of idiosyncratic biases in direction-of-motion perception⁹. Like our findings, biases in favor of a particular direction, measured at an interval of one year, were significantly correlated. However, exhibiting a substantial drift over time and measurements, these biases seemed to be less stable than those of our study.

Idiosyncratic side-preferences were also demonstrated in several species, including mice^{16,17}, rats^{18–20}, zebrafish²⁷. In *Drosophila*¹⁵, non-heritable idiosyncratic choice bias is stable for at least one month, a substantial fraction of their 2–3 months lifespan.

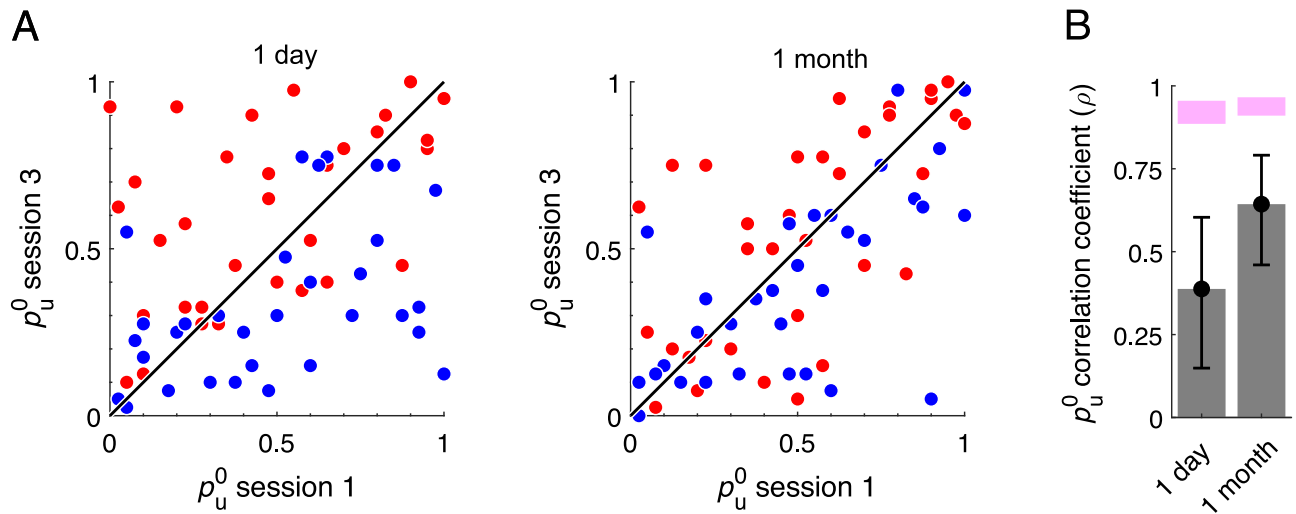


Fig. 6 | ICBs revert to their original values. **A** Value of p_u^0 in session 3 as a function of p_u^0 in session 1 for each participant in the 1-day (left, $n = 68$ participants) or 1-month (right, $n = 68$ participants) groups. Red dots: participants in the ‘enhance up’ group in session 2. (day: $n = 34$ participants, month: $n = 35$ participants). Blue dots: ‘enhance down’ participants in session 2 (day: $n = 34$ participants, month: $n = 33$ participants). Line: the diagonal. **B** Between-session correlation of p_u^0 in sessions 1 and 3 for the two delay groups ($n = 68$ participants in each group). Error

bars: 95% confidence interval of the correlation, bootstrap. Magenta: 95% confidence interval of the correlation expected under the assumption of complete stability, bootstrap Bernoulli processes. The correlation between p_u^0 is significantly larger in the 1-month interval group than in the 1-day group ($p = 0.04$, two-sided z-test for the difference in correlation after Fisher z-transformation) indicating that the feedback manipulation is ‘forgotten’, and the ICBs revert to their original values.

Our finding that participants are sensitive to feedback is in line with studies in a vernier and 2-tone discrimination tasks^{1,28}. They reported that providing incorrect feedback for a subset of high difficulty or impossible trials rapidly biases choices congruently with its asymmetry. In the vernier task, however, the effect of the feedback decays within the session (timescale of minutes), once the feedback is turned off²⁸, and is not transferred overnight³. Along the same lines, providing correct trial-by-trial feedback to participants that exhibit ICB can reduce it within the experimental session²⁹. Beyond perceptual tasks, choice preference induced by feedback is a form of operant learning^{30,31}, and can have long-lasting effects.

Limitations of the study

Our conclusion that ICB is not the result of idiosyncratic feedback-induced learning is based on the discrepancy between the relatively transient nature of the feedback-induced bias and the stability of the ICB. However, we cannot exclude the possibility that this discrepancy reflects the specific conditions in which we induced the bias. For example, it is possible that feedback-induced bias, if applied at an earlier age, would have led to a more sustained bias. It is also possible that providing biased feedback over more trials and / or for more sessions would have resulted in a longer-lasting effect^{3,28,32–36}. Moreover, it is possible that the dynamics of this decay are characterized by multiple timescales. In this case, the fact that feedback-induced bias decays within weeks does not imply that it would almost fully disappear if we waited months. Perhaps some of the timescales that characterize the dynamics of the decay are sufficiently long to account for ICB.

Another limitation of our study is its inability to disentangle the perceptual and motor contributions to the observed bias. For instance, participants in our task might prefer to report “up” in impossible trials either because they perceive the upper segment of the bar as longer than the lower segment (a perceptual bias) or because, all else being equal, they prefer the motor action associated with the “up” response (a motor bias). Previous works have shown that both biases can play a role. One study⁷ addressed this issue using a clever manipulation of the instructions to dissociate perceptual and motor biases. It reported that while both sensory and decisional (motor-related) biases play a role, the sensory component tends to be more dominant. Other studies, however, suggest a dominant motor component¹¹. Specifically, the cost of acting seems to be an important determinant of choice bias^{37,38}. If idiosyncratic, such costs can substantially contribute to the ICBs.

What underlies ICBs and their stability

Previous studies have shown that choices are strongly affected by the history of stimuli that the participant has been exposed to, a phenomenon that can be explained in the Bayesian framework³⁹. Therefore, a potential contributor to the ICB is the idiosyncratic distribution of stimuli that the participants were exposed to. However, this contribution is likely to be small. First, participants are primarily affected by the most recent stimuli⁴⁰ whereas in our experiments, a distribution of ICBs is observed even though within the experiment all participants were exposed to the same sequence of stimuli. Second, the influence of perturbing stimuli on bias has been shown to decay within tens of seconds⁹. However, we cannot rule out that habits that were formed before the experiment contributed to the bias. For example, participants may favor the ‘Down’ arrow button because they are accustomed to pressing taskbar icons located at the bottom of their computer monitor. Other participants may prefer the ‘Up’ arrow button because they are used to a taskbar at the top of the screen.

Another possible contributor are stable idiosyncratic asymmetries. For example, asymmetries in the anatomy of the eyes could give rise to asymmetries in the perception of the bar segments, which would manifest in a stable perceptual bias. While such idiosyncratic anatomical asymmetries are likely to exist, several lines of evidence suggest that this contribution is not dominant. First, we have previously reported a comparable distribution of ICBs in a sensory-motor task that does not entail any sensory ambiguity¹⁴. Second, in *drosophila*, left-right ICB and idiosyncratic object orientation control have been shown to be uncorrelated with anatomical asymmetries (e.g., between the legs length)^{15,41,42}. Finally, if the idiosyncratic asymmetry is weak relative to the width of the psychometric curve, its contribution to the bias would be negligible. Alternatively, if it is relatively strong, participants would always report ‘up’ or always report ‘down’ when the two segments are of the same length. Only a small range of idiosyncratic asymmetries is consistent with the experimentally observed wide distribution of ICBs.

In a previous study¹⁴, we showed that under general conditions, a wide distribution of ICBs emerges naturally from the dynamics of competition between neuronal networks that are *statistically identical*, that is, networks whose parameters are drawn from the same distribution. The intuition behind this result is that in such networks the microscopic structure of the connectivity varies across networks because they are only *statistically identical*. The main result there was that surprisingly, such heterogeneity is

predicted to substantially bias choices even in (infinitely) large decision-making networks. If the ICBs are the result of microscopic heterogeneities in the synaptic connections of decision-making networks, the long-term dynamics of the former can be related to one of the latter. We do not know much about the stability of the synaptic connectivity in decision-making networks, certainly not in humans. However, the question of stability of connectivity has been previously addressed, primarily by spine imaging of neurons in the mouse cortex. Several studies have shown that in some regions the connectivity is highly volatile⁴³. For example, in the mouse auditory cortex, most dendritic spines on pyramidal neurons, a proxy of the excitatory synapses that reside on them, are replaced within three weeks⁴⁴. Of the remaining spines, most change their size by at least a factor of two within that period⁴⁵. These results indicate that the fine structure of connectivity is highly dynamic⁴³, which according to our mechanism predicts a substantial volatility of ICBs in a timescale of weeks. On the other hand, in the barrel cortex, a fraction of spines is stably maintained for months in support of a stable connectome⁴⁶ and stable ICBs. The stability of the ICBs in our experiments predicts stability of the underlying connectivity.

Data availability

Raw and processed data are available in CSV and MATLAB .mat formats. All data from both studies, along with the MATLAB code (.m files) used for data processing and analysis, are available on Zenodo under the accession code <https://doi.org/10.5281/zenodo.15249456>. This includes the source data used to generate the figures, charts, and tables in both the main manuscript and Supplementary Information. The MATLAB code (.m files) is also available at <https://github.com/Lior-Lebovich/stabilityFeedbackICB>.

Code availability

All custom MATLAB code (.m files) used for data processing, statistical analyses, and generation of figures and tables (including those in the Supplementary Information) is available on Zenodo under the accession code <https://doi.org/10.5281/zenodo.15249456>, and at <https://github.com/Lior-Lebovich/stabilityFeedbackICB>. The code was written and tested using MATLAB R2023b, and all relevant parameters and instructions for reproducing the analyses are provided in the repository. In addition, the repository includes the full set of PHP-based scripts used to implement the web-based experimental interface for both studies. These scripts include embedded HTML and JavaScript and were used to present the tasks and store participant responses in SQL databases. The experiment code was originally developed and tested using PHP version 5 and may rely on outdated dependencies. Diagrams outlining the flow of these web-based experiments and their interaction with data storage are also included.

Received: 6 October 2024; Accepted: 8 May 2025;

Published online: 17 May 2025

References

- Loewenstein, Y., Raviv, O. & Ahissar, M. Dissecting the roles of supervised and unsupervised learning in perceptual discrimination judgments. *J. Neurosci.* **41**, 757–765 (2021).
- Herzog, M. H., Ewald, K. R. F., Hermens, F. & Fahle, M. Reverse feedback induces position and orientation specific changes. *Vis. Res.* **46**, 3761–3770 (2006).
- Aberg, K. C. & Herzog, M. H. Different types of feedback change decision criterion and sensitivity differently in perceptual learning. *J. Vis.* **12**, 1–11 (2012).
- Green, D. M. & Swets, J. A. *Signal Detection Theory and Psychophysics*, vol. 1 (Wiley, 1966).
- Klein, S. A. Measuring, estimating, and understanding the psychometric function: A commentary. *Percept. Psychophys.* **63**, 1421–1455 (2001).
- Baum, W. M. On two types of deviation from the matching law: bias and undermatching 1. *J. Exp. Anal. Behav.* **22**, 231–242 (1974).
- Linares, D., Aguilar-Lleyda, D. & López-Moliner, J. Decoupling sensory from decisional choice biases in perceptual decision making. *Elife* **8**, e43994 (2019).
- 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).
- Wexler, M., Duyck, M. & Mamassian, P. Persistent states in vision break universality and time invariance. *Proc. Natl. Acad. Sci.* **112**, 14990–14995 (2015).
- Rahnev, D. & Denison, R. N. Suboptimality in perceptual decision making. *Behav. Brain Sci.* **41**, e223 (2018).
- Schütz, A. C. Interindividual differences in preferred directions of perceptual and motor decisions. *J. Vis.* **14**, 16 (2014).
- Wexler, M., Mamassian, P. & Schütz, A. C. Structure of visual biases revealed by individual differences. *Vis. Res.* **195**, 108014 (2022).
- Kosovicheva, A. & Whitney, D. Stable individual signatures in object localization. *Curr. Biol.* **27**, R700–R701 (2017).
- Lebovich, L., Darshan, R., Lavi, Y., Hansel, D. & Loewenstein, Y. Idiosyncratic choice bias naturally emerges from intrinsic stochasticity in neuronal dynamics. *Nat. Hum. Behav.* **3**, 1190–1202 (2019).
- 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).
- Collins, R. L. On the inheritance of handedness I. laterality in inbred mice. *J. Hered.* **59**, 117–119 (1968).
- Collins, R. L. When left-handed mice live in right-handed worlds. *Science* **187**, 181–184 (1975).
- Kurikawa, T., Haga, T., Handa, T., Harukuni, R. & Fukai, T. Neuronal stability in medial frontal cortex sets individual variability in decision-making. *Nat. Neurosci.* **21**, 1764–1773 (2018).
- Andrade, C., Alwarshetty, M., Sudha, S. & Chandra, J. S. Effect of innate direction bias on T-maze learning in rats: implications for research. *J. Neurosci. Methods* **110**, 31–35 (2001).
- Castellano, M. A., Diaz-Palarea, M. D., Rodriguez, M. & Barroso, J. Lateralization in male rats and dopaminergic system: evidence of right-side population bias. *Physiol. Behav.* **40**, 607–612 (1987).
- Barracough, D. J., Conroy, M. L. & Lee, D. Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* **7**, 404–410 (2004).
- Regaiolli, B., Spiezio, C. & Hopkins, W. D. Asymmetries in mother-infant behaviour in Barbary macaques (*Macaca sylvanus*). *PeerJ* **6**, e4736 (2018).
- Jewell, G. & McCourt, M. E. Pseudoneglect: a review and meta-analysis of performance factors in line bisection tasks. *Neuropsychologia* **38**, 93–110 (2000).
- Learmonth, G., Gallagher, A., Gibson, J., Thut, G. & Harvey, M. Intra- and inter-task reliability of spatial attention measures in pseudoneglect. *PLoS One* **10**, e0138379 (2015).
- Mamassian, P. & Wallace, J. M. Sustained directional biases in motion transparency. *J. Vis.* **10**, 23 (2010).
- Grabot, L. & van Wassenhove, V. Time order as psychological bias. *Psychol. Sci.* **28**, 670–678 (2017).
- Horstick, E. J., Bayle, Y. & Burgess, H. A. Molecular and cellular determinants of motor asymmetry in zebrafish. *Nat. Commun.* **11**, 1170 (2020).
- Herzog, M. H. & Fahle, M. Effects of biased feedback on learning and deciding in a vernier discrimination task. *Vis. Res.* **39**, 4232–4243 (1999).
- Haddara, N. & Rahnev, D. The impact of feedback on perceptual decision-making and metacognition: reduction in bias but no change in sensitivity. *Psychol. Sci.* **33**, 259–275 (2022).
- Mongillo, G., Shteingart, H. & Loewenstein, Y. The misbehavior of reinforcement learning. *Proc. IEEE* **102**, 528–541 (2014).

31. Shteingart, H. & Loewenstein, Y. Reinforcement learning and human behavior. *Curr. Opin. Neurobiol.* **25**, 93–98 (2014).
32. Wright, B. A. & Sabin, A. T. Perceptual learning: how much daily training is enough? *Exp. Brain Res.* **180**, 727–736 (2007).
33. Aberg, K. C., Tartaglia, E. M. & Herzog, M. H. Perceptual learning with Chevrons requires a minimal number of trials, transfers to untrained directions, but does not require sleep. *Vis. Res.* **49**, 2087–2094 (2009).
34. Teichner, W. H. Experimental extinction as a function of the intertrial intervals during conditioning and extinction. *J. Exp. Psychol.* **44**, 170 (1952).
35. Terrace, H. S., Gibbon, J., Farrell, L. & Baldock, M. D. Temporal factors influencing the acquisition and maintenance of an autoshaped keypeck. *Anim. Learn. Behav.* **3**, 53–62 (1975).
36. Wimmer, G. E., Li, J. K., Gorgolewski, K. J. & Poldrack, R. A. Reward learning over weeks versus minutes increases the neural representation of value in the human brain. *J. Neurosci.* **38**, 7649–7666 (2018).
37. Marcos, E., Cos, I., Girard, B. & Verschure, P. F. M. J. Motor cost influences perceptual decisions. *PLoS One* **10**, 1–12 (2015).
38. Hagura, N., Haggard, P. & Diedrichsen, J. Perceptual decisions are biased by the cost to act. *Elife* **6**, 1–20 (2017).
39. Ashourian, P. & Loewenstein, Y. Bayesian inference underlies the contraction bias in delayed comparison tasks. *PLoS One* **6**, e19551 (2011).
40. Raviv, O., Ahissar, M. & Loewenstein, Y. How recent history affects perception: the normative approach and its heuristic approximation. *PLoS Comput. Biol.* **8**, e1002731 (2012).
41. Linneweber, G. A. et al. A neurodevelopmental origin of behavioral individuality in the Drosophila visual system. *Science* **367**, 1112–1119 (2020).
42. Skutt-Kakaria, K., Reimers, P., Currier, T. A., Werkhoven, Z. & de Bivort, B. L. A neural circuit basis for context-modulation of individual locomotor behavior. *Biorxiv*, <https://www.biorxiv.org/content/10.1101/797126v1> (2019).
43. Mongillo, G., Rumpel, S. & Loewenstein, Y. Intrinsic volatility of synaptic connections—a challenge to the synaptic trace theory of memory. *Curr. Opin. Neurobiol.* **46**, 7–13 (2017).
44. Loewenstein, Y., Yanover, U. & Rumpel, S. Predicting the dynamics of network connectivity in the neocortex. *J. Neurosci.* **35**, 12535–12544 (2015).
45. Loewenstein, Y., Kuras, A. & Rumpel, S. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *J. Neurosci.* **31**, 9481–9488 (2011).
46. Yang, G., Pan, F. & Gan, W.-B. Stably maintained dendritic spines are associated with lifelong memories. *Nature* **462**, 920–924 (2009).

Acknowledgements

This work was supported by the German Research Foundation (DFG CRC 1080) and the Gatsby charitable foundation (Y.L.) and in part by ANR-17-NEUR-005 (DH). Y.L. is the incumbent of the David and Inez Myers Chair in

Neural Computation. Work performed in the framework of the France-Israel Center for Neural Computation (CNRS/Hebrew University of Jerusalem). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript. L.L. thanks Nelly Lebovich for insightful discussions.

Author contributions

L.L., L.K., D.H. and Y.L. conceived and planned the experiments; L.L. and L.K. programmed and ran the experiments; L.L. and Y.L. analyzed the data; L.L., D.H. and Y.L. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s44271-025-00263-0>.

Correspondence and requests for materials should be addressed to Lior Lebovich.

Peer review information *Communications Psychology* thanks the anonymous reviewers for their contribution to the peer review of this work. Primary Handling Editors: Troby Ka-Yan Lui. [A peer review file is available].

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025