The impact of feedback on perceptual decision making and metacognition:

Reduction in bias but no change in sensitivity

Nadia Haddara & Dobromir Rahnev

School of Psychology, Georgia Institute of Technology, Atlanta, GA, USA

Keywords: perceptual decision making, metacognition, criterion, bias, sensitivity, confidence calibration.

Corresponding authors:

Nadia Haddara

Georgia Institute of Technology

831 Marietta St NW

Atlanta, GA 30318

nadiahaddara@gatech.edu

Dobromir Rahnev

Georgia Institute of Technology

654 Cherry St NW

Atlanta, GA 30332

rahnev@psych.gatech.edu

Abstract

It is widely believed that feedback improves behavior but the mechanisms behind this improvement remain unclear. Different theories postulate that feedback has either a direct effect on performance through automatic reinforcement mechanisms or only an indirect effect mediated by a deliberate change in strategy. To adjudicate between these competing accounts, we performed two large studies (total N = 518) with approximately half the subjects receiving trial-by-trial feedback on a perceptual task, while the other half not receiving any feedback. We found that feedback had no effect on either perceptual or metacognitive sensitivity even after seven days of training. On the other hand, feedback significantly affected subjects' response strategies by reducing response bias and improving confidence calibration. These results strongly support the view that feedback does not improve behavior through direct reinforcement mechanisms but that its beneficial effects stem from allowing people to adjust their strategies for performing the task.

Word count: 149/150

Statement of relevance

How can we help people improve their performance? It is often thought that across a variety of domains from education to work settings to sports achievements to various cognitive tasks, performance can be improved by simply providing feedback. However, it remains unclear when and why feedback helps. To address this question, we performed two large experiments where subjects made perceptual and metacognitive judgments. We found that trial-by-trial feedback had no effect on perceptual or metacognitive sensitivity but reduced both perceptual and metacognitive bias. These results demonstrate that, contrary to popular beliefs, feedback has no direct, automatic effects on performance. Instead, the beneficial effects of feedback appear to be exclusively driven by the fact that feedback allows subjects to change their strategy for performing the task. Therefore, feedback can be expected to only be effective if people are willing and able to change the way they perform a specific task.

Word count: 149/150

Introduction

The notion that feedback improves performance has been a mainstay of experimental psychology for over a century (Judd, 1905; Wright, 1906). Indeed, as one paper put it, “the positive effect of feedback on performance has become one of the most accepted principles in psychology” (Pritchard, Jones, Roth, Stuebing, & Ekeberg, 1988, p.338). However, despite the large amount of research on this topic, the mechanisms through which feedback affects performance are still unclear.

One long-standing view about the mechanism of feedback is that it allows for stimuli and responses to become associated via automatic reinforcement mechanisms. For example, an early version of this theory, dubbed the “Law of Effect” (Thorndike, 1927), proposed that feedback directly and automatically strengthens internal connections and thus leads to improvements in both animals and humans independent of any cognitive strategies adopted by the organism. Similar theories regarding automatic mechanisms of feedback remain popular to this day (Petrov, Doshier, & Lu, 2005). On the other hand, a competing view is that feedback is only effective to the extent to which it improves one’s strategy for completing the task (Vollmeyer & Rheinberg, 2005). However, these two accounts of the mechanisms of feedback have never been distinguished. In fact, adjudicating between these two theories is particularly challenging because both direct changes to observer sensitivity and indirect strategy adjustment could result in identical outcomes of improved performance.

Fortunately, these two possibilities can be differentiated in the context of simple 2-choice perceptual tasks. Behavior in such tasks can be fully captured by two different parameters derived from signal detection theory: perceptual sensitivity (d') and response

criterion (c) (Green & Swets, 1966). Perceptual sensitivity is normally determined by sensory processing outside of a subject's conscious control, whereas response criterion is under one's control and can be influenced by a deliberate change in strategy (Macmillan & Creelman, 2005). Therefore, by examining whether feedback affects perceptual sensitivity or response criterion, we can distinguish between a direct (based on automatic reinforcement) and an indirect (based on strategy change) effect of feedback on behavior.

However, previous research with perceptual tasks has been surprisingly inconsistent, with some studies reporting feedback-related performance enhancements (Ball & Sekuler, 1987; Fahle & Edelman, 1993; Herzog & Fahle, 1997; Seitz, Nanez, Holloway, Tsushima, & Watanabe, 2006), while others showing no difference between Feedback and No Feedback groups (Goldhacker, Rosengarth, Plank, & Greenlee, 2014; Petrov, Doshier, & Lu, 2006; Rouault, Dayan, & Fleming, 2019; Shibata, Yamagishi, Ishii, & Kawato, 2009). Considering the small sample sizes used in most of these studies (6-36 subjects split into two or more groups except Shiu & Pashler, 1992, N = 59), the inconsistency is not surprising since such designs are likely to lead to both false positives and false negatives, while also reflecting the prevailing bias rather than true effects (Ioannidis, 2005).

In addition to simple perceptual tasks, the mechanisms of feedback can also be investigated using tasks that assess people's metacognition (Metcalf & Shimamura, 1994). Such tasks collect confidence ratings regarding the perceived accuracy of the primary decision, and can provide a complementary test of whether feedback has a direct impact on sensitivity or an indirect strategy-mediated influence that can only be seen in one's confidence calibration (Rahnev & Denison, 2018). However, while a relatively sizeable

literature has at least examined how feedback affects perceptual sensitivity, no studies to date have examined the effects of trial-by-trial feedback on metacognitive sensitivity (though block feedback was examined by Rouault et al., 2019 and a neurofeedback paradigm was used by Cortese, Amano, Koizumi, Kawato, & Lau, 2016).

Here we investigated the impact of trial-by-trial feedback on both perceptual and metacognitive judgments in two experiments. Experiment 1 included a single day of training in a sample ($N = 443$) much larger than previous studies. On the other hand, a preregistered Experiment 2 employed a much longer training period (seven days, $N = 75$). Subjects in both experiments were randomly assigned to Feedback and No Feedback groups. To anticipate, we found that feedback had no effect on perceptual or metacognitive sensitivity, but it reduced both perceptual and metacognitive biases, indicating that it affected performance via strategy adjustments rather than direct reinforcement. These findings reveal the mechanisms of feedback in perceptual decision making and metacognition, and have strong implications about the use of feedback in applied settings.

Methods

Subjects

Subjects for both experiments were recruited on Amazon's Mechanical Turk online platform and were compensated \$7.25/hour. For Experiment 1, we chose a sample size ($N = 443$) much larger than in previous studies examining the role of feedback in perceptual decision making in order to ensure sufficient power for our analyses. For Experiment 2, we preregistered a sample size of 60 after exclusions. The preregistration can be accessed at <https://osf.io/efymg> and was completed after data from Experiment 1 were analyzed. A total of 75 new subjects completed all seven days of Experiment 2 (participants from Experiment 1 were not allowed to participate in Experiment 2) with 15 of them being excluded based on predetermined criteria. Subjects in Experiment 2 received a bonus for completing all seven days of the experiment, such that they gained one cent for every three correct trials (average bonus = \$8.7). Recruiting Mechanical Turk research participants is considered a form of convenience sampling. All subjects reported normal or corrected-to-normal vision and provided written informed consent. Experimental procedures were approved by the Georgia Institute of Technology Institutional Review Board.

Procedure

In Experiment 1, subjects were randomly assigned to a Feedback or No Feedback group and completed two separate perceptual tasks adapted from Rahnev, Koizumi, McCurdy, D'Esposito, & Lau (2015). In Task 1, subjects made a perceptual judgment via a key press indicating whether the letter X or O occurred more frequently in a 7 x 7 grid presented for 500 ms. Subjects were given an untimed response period for the perceptual decision. After indicating their response, subjects were prompted to rate their confidence

in the accuracy of their decision on a 4-point scale (1 = low confidence to 4 = high confidence; untimed response). Following the confidence rating, the Feedback group received feedback concerning the accuracy of their perceptual judgment ("Correct"/"Wrong"; 500 ms), whereas the No Feedback group saw a fixation cross for 500 ms, thus keeping the time between the confidence response and the onset of the next trial the same between the two groups (Fig. 1). Subjects completed 11 blocks of 30 trials each (330 total trials) and were given the opportunity to take a break at the end of each block. The first trial of each block began with a longer, 2-second fixation to give subjects time to focus their attention, and each subsequent trial began with a shorter 500-ms fixation.

Fig. 1. Experimental tasks. Experiment 1 consisted of two tasks depicted above conducted in a single session. For both tasks, subjects were required to indicate which of two stimulus classes occurred more frequently in a 7 x 7 grid (Task 1: X/O, Task 2: red/blue). After providing their response, subjects were prompted to give confidence ratings on a scale from 1 to 4. In Task 1, approximately half of the subjects received trial-by-trial feedback (Feedback group), whereas the rest of the subjects (No Feedback group) received no feedback at all. In Task 2, neither group received feedback; this second task was used to investigate which effects of feedback would generalize to a new task. Experiment 2 consisted of only Task 1 but was conducted over seven days.

Task 2 was designed to be very similar to Task 1 and the only difference was the perceptual dimension being discriminated. Unlike Task 1, where subjects discriminated between X's and O's, Task 2 featured a discrimination between red and blue circles and consisted of 5 blocks of 30 trials (150 total trials with opportunities for breaks after each block). Critically, in Task 2 no subject received feedback on their performance. This was done in order to determine whether any putative feedback-related improvements in Task 1 would generalize to a new task.

The dominant stimulus class (X or O in Task 1, red or blue in Task 2) was randomly determined on each trial, with the proportion of the dominant stimulus within the 7 x 7 grid fixed at 30/49 for Task 1 and 27/49 for Task 2. The proportions were different for the two tasks to ensure that the two tasks had similar difficulty levels, since our pilot data suggested that the red/blue discrimination was easier for subjects. The average accuracy was 80% for Task 1 and 78% for Task 2.

Experiment 2 employed Task 1 from Experiment 1 (see Fig. 1) with subjects performing the task over seven days. Each day of the experiment included 20 blocks of 25 trials each (for a total of 500 trials per day and 3,500 in total). The first six blocks of Day 1 were used to adjust the task difficulty for each subject. We ran a standard 2-up-1-down staircase that started with 30 (of 49) characters from the dominant stimulus. The final

difficulty for each subject was determined as the average of the number of dominant items at the time of staircase reversals (ignoring the first four reversals; Bang, Milton, Sasaki, Watanabe, & Rahnev, 2019; Bang, Sasaki, Watanabe, & Rahnev, 2018), rounded to the nearest integer. The resulting difficulty value was then fixed for the remainder of the experiment (14 blocks on Day 1 and all trials in subsequent days). After completing Day 1, subjects who passed quality checks (see Analyses section) were invited to complete the remainder of the experiment. The average accuracy across all days of the experiment was 75%.

The experiments were designed using jsPsych 5.0.3 library and grid stimuli were created using in-house JavaScript code. To ensure that the stimulus size was similar for different subjects who took the experiment on different screens, we used a procedure previously established in our lab (Bang, Shekhar, & Rahnev, 2019) where subjects were asked to adjust the stimuli on the screen to match the size of real-life objects (a credit card or a quarter) to on-screen objects.

Sensitivity and bias measures

We computed the signal detection theoretic (Green & Swets, 1966; Macmillan & Creelman, 2005) parameters stimulus sensitivity (d') and response criterion (c) to determine subjects' performance and degree of response bias on the tasks. These measures were calculated based on the observed hit rate (HR) and false alarm rate (FAR) such that:

$$d' = \phi^{-1}(HR) - \phi^{-1}(FAR)$$

and

$$c = -\frac{1}{2} \times (\phi^{-1}(HR) + \phi^{-1}(FAR))$$

where ϕ^{-1} is the inverse of the cumulative standard normal distribution that transforms HR and FAR into z-scores. HR and FAR were defined by treating the letter X in Task 1 and blue circles in Task 2 as the target. Therefore, negative c values indicate a bias for the letter X (Task 1) or the color blue (Task 2), whereas positive c values indicate a bias for the letter O (Task 1) or the color red (Task 2).

We further computed metacognitive sensitivity, which is a measure of one's ability to distinguish between one's own correct and incorrect judgments (Fleming & Lau, 2014). Similar to perceptual sensitivity, which reflects the strength of the relationship between one's choices and the stimulus, metacognitive sensitivity reflects the strength of the relationship between one's confidence ratings and the accuracy of the perceptual decisions. Higher metacognitive sensitivity implies that one's confidence ratings are more informative regarding the accuracy of one's judgments. We quantified metacognitive sensitivity using the measure *meta-d'* developed by Maniscalco & Lau (2012), which is derived based on a signal detection model and is expressed in the same units as the measure d' . Finally, we computed metacognitive bias as one's tendency to have confidence ratings that are too low or too high relative to one's level of perceptual sensitivity (Fleming & Lau, 2014). This was assessed as the across-subject Pearson's correlation between d' and average confidence for each subject. This measure was used to indicate the degree of confidence bias across subjects with higher correlation coefficients indicative of better confidence calibration.

Analyses

All statistical analyses for both experiments were conducted in MATLAB, and figures were generated in MATLAB and R software.

Experiment 1

Subjects who performed at less than 55% correct or greater than 95% correct were excluded from the analyses. These exclusions were done separately for each of the two perceptual tasks, with 48 subjects excluded from Task 1 and 47 from Task 2 (~11% exclusion rate in each task). Additionally, trials with reaction times that were either too fast (<200 ms) or too slow (>2,000 ms) were excluded from the analyses.

To determine the effect of feedback on perceptual sensitivity, we performed independent sample t-tests comparing d' of the Feedback and No Feedback groups. Response bias was assessed by comparing the across-subject variability of the criterion c for each of the two groups using a two-tailed F-test for equality of variances. (Note that because the equality of variance F-test is two-tailed while a standard ANOVA F-test is one-tailed, the p value associated with an F statistic from an equality of variances test with $F > 1$ is two times higher than the p value for the same F statistic obtained from an ANOVA.) Because strong bias results in very large positive or negative values of the criterion c , a distribution with low variability in the criterion values indicates the presence of smaller bias in the group, whereas a distribution with high variability in the criterion values indicates the presence of larger bias.

To determine the effect of feedback on metacognitive sensitivity, the Feedback and No Feedback groups were compared using independent-sample t-tests to assess differences in $meta-d'$ (i.e., metacognitive sensitivity). On the other hand, to compare the size of metacognitive bias between the two groups, the d' -confidence correlation coefficient

for each group was computed and the two values were compared using a standard z-test for comparing two independent r-values such that:

$$z = \frac{z_1 - z_2}{\sqrt{\frac{1}{N_1 - 3} + \frac{1}{N_2 - 3}}}$$

where $z_i = \frac{1}{2} \log \left(\frac{1+r_i}{1-r_i} \right)$ is the z-transformed correlation coefficient (r-value) and N_i is the sample size for the corresponding group ($i = 1, 2$).

We further investigated whether there were any feedback-related changes in decision and confidence reaction times (RT) by conducting independent-samples t-tests. In the analysis of confidence RTs, we additionally excluded trials with confidence RTs outside the range of 200-2,000 ms. In a separate analysis, we checked whether the difference in decision and confidence RTs for the Feedback and No Feedback groups changed from Task 1 to Task 2; for that analysis, we only included subjects with >55% and <95% accuracy for both tasks and performed an independent samples t-test.

In addition, we performed Bayesian analyses using the JASP software Version 0.9.2 (JASP Team, 2020) and MATLAB. When comparing groups, we report either BF_{01} , which denotes the support for the null over the alternative hypothesis, or BF_{10} , which denotes the support for the alternative over the null hypothesis with higher values indicating stronger evidence in both cases (Masson, 2011). For ANOVAs, we report $BF_{Inclusion}$, which indicates whether the observed data is more probable under models that include a particular factor. For all Bayesian analyses, the default priors in JASP were used.

As an alternative to our SDT-based analyses, we computed the parameters drift rate, boundary separation, and non-decision time of the drift diffusion model (Ratcliff, 1978) separately for the Feedback and No Feedback groups. We used a simple version of the drift diffusion model without variability parameters (Wagenmakers, van der Maas, & Grasman, 2007) that has been previously found to better recover parameter changes in the drift diffusion model (van Ravenzwaaij, Donkin, & Vandekerckhove, 2017).

Finally, we computed d' and $meta-d'$ in a hierarchical estimation framework that is more robust to small sample sizes and edge effects. We employed the HMeta-d toolbox (Fleming, 2017), which uses Markov chain Monte Carlo (MCMC) sampling to estimate posterior distribution over model parameters. To assess the existence of group differences in d' and $meta-d'$, we calculated the 95% highest density intervals (HDIs) of the posteriors.

Experiment 2

All analyses for Experiment 2 were conducted in accordance with our preregistration, and mostly followed the analysis steps from Experiment 1. We first performed subject and trial exclusions based on our pre-registration criteria. We removed the data from the first six blocks of Day 1 for all subjects because these blocks contained the trials from our staircase procedure. Upon completion of Day 1, data quality for each subject was assessed using the remaining data from that day and we excluded subjects from participating in days 2-7 who had (1) accuracy lower than 60% correct, (2) accuracy higher than 85% correct, or (3) more than 10% of decision RTs that were under 200 ms or above 2,000 ms. A total of 75 subjects successfully completed all seven days of the experiment and 15 subjects were excluded from analyses based on the pre-determined criteria that included (1) performing at less than 55% correct across all seven days, or (2)

having more than 15% of decision RTs below 200 ms or above 2,000 ms. Lastly, as in Experiment 1, individual trials with decision RTs considered too short (below 200 ms) or too long (above 2,000 ms) were excluded from the analyses.

To assess changes in perceptual and metacognitive sensitivity between groups over the seven days of the experiment, we conducted mixed ANOVAs with a between-subjects factor of Group (Feedback and No Feedback) and a within-subjects factor of Session (1–7) on d' and $meta-d'$. Similar analyses on decision and confidence RTs are reported in the Supplementary Results and Supplementary Fig. S1. Additionally, we performed a linear regression on task accuracy (percent correct) as a function of block over the time course of the experiment for each subject, followed by an independent sample t-test comparing the slope of the regression between the Feedback and No Feedback groups. Lastly, Bayesian analyses were implemented to determine the strength of the evidence in support of the null-hypothesis for all relevant non-significant effects.

Data & Code

All data and codes for the analyses are freely available online at <https://osf.io/94r87/>. In addition, the complete datasets for both experiments have been uploaded to the Confidence Database (Rahnev et al., 2020).

Results

Experiment 1

To uncover whether feedback acts via automatic reinforcement mechanisms or deliberate changes in strategy, we examined the effect of trial-by-trial feedback on subjects' sensitivity and bias in both perceptual and metacognitive judgments. Subjects completed a perceptual task that required them to indicate whether more X's or O's were presented in a 7 x 7 grid (Fig. 1). Each subject was randomly assigned to a Feedback or No Feedback group with the former group receiving trial-by-trial feedback, and the latter not receiving such feedback. To investigate how the effects of feedback generalize to a novel task, all subjects also completed a second perceptual task where no feedback was given, in which they judged whether more red or blue circles appeared in a 7 x 7 grid (Fig. 1). We first examine the effects of feedback in Task 1, and later explore which of these effects generalized to Task 2.

Trial-by-trial feedback has no effect on perceptual or metacognitive sensitivity

We first sought to determine whether trial-by-trial feedback can affect behavior via automatic reinforcement mechanisms, which would be manifested in feedback increasing perceptual sensitivity. To do so, we compared d' values for the Feedback and No Feedback groups. We found that perceptual sensitivity computed over all trials was virtually identical across the two groups ($d'_{\text{Feedback}} = 1.79$, $d'_{\text{No Feedback}} = 1.80$; $t(392) = -.16$, $p = .87$, Cohen's $d = -.02$; Fig. 2A). We further investigated whether this null effect was due to a lack of power or whether our data can provide positive evidence for a lack of difference between the two groups. We therefore conducted a Bayesian independent-samples t-test, which showed that the data strongly supported the null hypothesis of no difference between the two groups

($BF_{01} = 8.9$). We also replicated the lack of d' difference in the two groups using a hierarchical estimation framework ($d'_{\text{Feedback}} = 1.78$; $d'_{\text{No Feedback}} = 1.72$; confidence interval for the difference = $[-0.33, 1.03]$).

Fig. 2. Feedback has no effect on either perceptual or metacognitive sensitivity. (A) Perceptual sensitivity (d') densities and box plots with individual subject data for each of the two feedback groups. There was no difference in d' between Feedback and No Feedback groups as can be seen in the overlap of the two distributions. (B) Metacognitive sensitivity ($meta-d'$) densities and box plots with individual subject data for Feedback and No Feedback groups. Similar to the d' results, we found no difference in $meta-d'$ between the groups. These results demonstrate that trial-by-trial feedback has no effect on either perceptual or metacognitive sensitivity. Box plots show the median and the interquartile (25-75%) range with the whiskers indicating the 2-98% range. Raincloud plots were adapted from Allen et al. (2019).

Despite the absence of overall differences in perceptual sensitivity between the Feedback and No Feedback groups, it could be argued that trial-by-trial feedback may still have had a beneficial effect on perceptual accuracy. Specifically, it is possible that the effects of trial-by-trial feedback take time to manifest themselves, and therefore analyzing all 330 task trials together may mask the beneficial effects of feedback on perceptual

sensitivity that may emerge only later in the task. To investigate this possibility, we only considered the data from the second half of the task and again found no difference between the two groups ($d'_{\text{Feedback}} = 1.88$, $d'_{\text{No Feedback}} = 1.84$; $t(388) = .59$, $p = .56$, Cohen's $d = .06$, $BF_{01} = 7.6$). Similar results were further obtained if d' for the last N blocks of 30 trials were analyzed for N between 1 and 10 (all p 's $> .72$; all BF_{01} 's > 8.4). We also considered the rate of learning by performing a linear regression using trial number to predict task accuracy (percent correct; Fig. 3) and found a significantly positive slope for both the Feedback (slope = 1.6×10^{-3} , $t(199) = 6.65$, $p = 2.8 \times 10^{-10}$, Cohen's $d = .47$) and No Feedback (slope = 1.1×10^{-3} , $t(193) = 4.49$, $p = 1.2 \times 10^{-5}$, Cohen's $d = .33$) groups, with no difference between the two groups ($t(392) = 1.45$, $p = .15$, Cohen's $d = .15$; $BF_{01} = 3.3$). Additionally, the rate of learning correlated with both perceptual and metacognitive sensitivity (Supplementary, Fig. S2).

Fig. 3. Feedback has no effect on learning rate. We conducted a linear regression using trial number to predict percent correct and found that accuracy increased over time but that the rate of increase was the same for the two feedback groups.

Finally, it is possible that a gain in sensitivity for the Feedback group is obscured due to faster responding. We therefore fit the drift diffusion model to the data and found no difference between the Feedback and No Feedback groups in drift rate v ($v_{\text{Feedback}} = .110$, $v_{\text{No Feedback}} = .106$, $t(392) = .89$, $p = .37$, Cohen's $d = .09$, $\text{BF}_{01} = 6.1$). Overall, our results show that trial-by-trial feedback had virtually no effect on perceptual sensitivity regardless of how sensitivity was assessed, thus suggesting that feedback does not affect choice behavior via direct reinforcement mechanisms.

However, even if feedback does not affect the perceptual decisions via automatic reinforcement mechanisms, it is possible that it affects higher-level metacognitive judgments via such mechanisms. Such an effect would be manifested in trial-by-trial feedback improving subjects' metacognitive sensitivity. To investigate this possibility, we computed the measure of metacognitive sensitivity $meta-d'$ (Maniscalco & Lau, 2012) and compared it between the Feedback and No Feedback groups. We again found no effect of feedback ($meta-d'_{\text{Feedback}} = 1.23$, $meta-d'_{\text{No Feedback}} = 1.33$; $t(392) = -1.29$, $p = .20$, Cohen's $d = -.13$, $\text{BF}_{01} = 4.0$; Fig. 2B). Similar effects were again obtained for the second half of the task ($meta-d'_{\text{Feedback}} = 1.21$, $meta-d'_{\text{No Feedback}} = 1.30$; $t(388) = -1.12$, $p = .26$, Cohen's $d = -.11$; $\text{BF}_{01} = 4.9$), as well as when using a hierarchical estimation framework to estimate metacognitive efficiency ($M_{\text{ratio}} = meta-d'/d'$; $M_{\text{ratio}}_{\text{Feedback}} = .73$; $M_{\text{ratio}}_{\text{No Feedback}} = .75$; confidence interval for the difference = $[-0.09, 0.06]$). Taken together, these findings demonstrate that trial-by-trial feedback does not affect either perceptual or metacognitive sensitivity in our task, and therefore does not act via direct reinforcement mechanisms either at the level of the perceptual or metacognitive judgments.

Trial-by-trial feedback reduces bias in perceptual and metacognitive judgments

The results so far demonstrate that feedback does not act via automatic reinforcement mechanisms but do not elucidate whether feedback affects subjects' strategies. To examine how feedback affects subjects' strategic behavior, we analyzed how the presence of trial-by-trial feedback influenced bias in perceptual and metacognitive judgments.

We quantified the bias in the perceptual task using the signal detection theoretic measure c , which indicates the location of the response criterion. Large negative values of c indicate a strong bias towards the letter 'X,' large positive values indicate a strong bias towards the letter 'O' and values close to zero indicate a lack of bias. A reduction in bias would manifest as criterion values becoming less extreme and therefore having smaller variance. This is exactly what we found: the Feedback group (standard deviation = .21) had smaller variability of criterion scores than the No Feedback group (standard deviation = .25) with the difference being statistically significant ($F(199, 193) = .74, p = .035$, F-test of equality of variances; Fig. 4). Both the Feedback and No Feedback groups had a mean value of c close to zero ($c_{\text{Feedback}} = -.03, c_{\text{No Feedback}} = -.02; t(392) = -.50, p = .62$, Cohen's $d = -.05$, $BF_{01} = 8.0$) indicating that neither group as a whole had a bias towards a particular stimulus category (neither the O's nor the X'es were preferred in the whole sample of subjects). However, the distribution of criterion values for the No Feedback group had longer tails corresponding to individual subjects with more extreme response biases. Control analyses where bias was analyzed in mini blocks of 10 trials suggested that this reduction in bias emerged very quickly and could already be detected by the end of the first

block of 30 trials (Supplementary, Fig. S3). These results are consistent with the notion that trial-by-trial feedback allowed subjects to adjust their response strategy thus reducing bias.

Fig. 4. Feedback reduces bias in the perceptual task. Response criterion (c) densities and box plots with individual subject data for each feedback group. The Feedback group showed a significant reduction in bias compared to the No Feedback group as demonstrated by the less extreme criterion values (smaller variance), which can be seen at the tails of the distribution. Box plots show the median and the interquartile (25-75%) range with the whiskers indicating the 2-98% range.

Given that trial-by-trial feedback induced a strategic change in bias in the perceptual judgment, we investigated whether a similar strategic change would also occur for the metacognitive judgment. Conceptually, metacognitive bias is the propensity of the overall confidence ratings of a given subject to deviate from the overall sensitivity of that same subject (Fleming & Lau, 2014). For example, subjects with a high metacognitive bias may have low perceptual sensitivity but a high average confidence or vice versa. Note that because confidence was collected on a 1-4 scale where the different options on the confidence scale were not directly associated with different accuracy levels, it is impossible

to quantify an individual subject's metacognitive bias in isolation. However, it is possible to determine whether the Feedback and No Feedback groups differ in how strongly average perceptual sensitivity and average confidence are related across subjects with a weaker relationship indicating the presence of larger metacognitive biases.

We therefore investigated the presence of metacognitive bias in the Feedback and No Feedback groups by correlating the perceptual sensitivity (d') and the average confidence across subjects for each group. We found no significant correlation between d' and average confidence in the No Feedback group ($r = .006$; $p = .94$; Fig. 5A) but a significant positive correlation in the Feedback group ($r = .27$; $p = .0001$; Fig. 5B). Critically, the difference between the r -values of the two feedback groups was statistically significant ($p = .008$, z test for comparing r -values; Fig. 5C). These results are again consistent with the notion that subjects in the Feedback group were able to strategically reduce their metacognitive biases, which led to a stronger across-subject association between sensitivity and confidence in that group.

Fig. 5. Feedback reduces metacognitive bias. (A) There was no significant across-subject correlation between d' and average confidence ratings in the No Feedback group. (B) There was a significant positive correlation between d' and average confidence in the Feedback group. (C) The difference between the correlation coefficients (r -values) in the two feedback groups was statistically significant. In (A) and (B), each point represents a single subject; the dark gray line and the shaded gray area indicate the line of best fit and its 95% confidence interval, respectively. In (C) error bars depict standard error of mean.

Trial-by-trial feedback decreases RTs

Beyond examining the effect of trial-by-trial feedback on sensitivity and bias in perceptual and metacognitive judgments, we also explored feedback's effects on RTs. We first verified that decision RTs were faster for both correct trials ($t(393) = 19.6, p = 3.5 \times 10^{-60}$, Cohen's $d = .98$) and high confidence ratings ($t(393) = -21.8, p = 1.5 \times 10^{-69}$, Cohen's $d = -1.097$), thus confirming that we can replicate known effects from the literature. Critically, we found that subjects in the Feedback group exhibited faster decision RTs (Mean RT = 607 ms) than subjects in the No Feedback group (Mean RT = 699 ms) with the difference being highly significant ($t(392) = -5.57, p = 4.7 \times 10^{-8}$, Cohen's $d = -.56$, $BF_{10} = 2.2 \times 10^5$; Fig. 6A). Additional analyses with the drift diffusion model (which found no difference in drift rate between the groups, see above) suggested that this effect was driven primarily by a feedback-induced decrease in non-decision time Ter ($Ter_{\text{Feedback}} = .23$, $Ter_{\text{No Feedback}} = .30$, $t(392) = -5.05, p = 6.7 \times 10^{-7}$, Cohen's $d = -.51$, $BF_{10} = 1.8 \times 10^4$), though there was also a small effect on boundary separation a ($a_{\text{Feedback}} = .133$, $a_{\text{No Feedback}} = .137$, $t(392) = -2.81, p = .005$, Cohen's $d = -.28$, $BF_{10} = 4.9$). Moreover, we found a similar effect for the RTs for the confidence response. Indeed, the Feedback group had lower confidence RTs (Mean confidence RT = 443 ms) than the No Feedback group (Mean confidence RT = 501 ms; $t(392) = -4.54, p = 7.4 \times 10^{-6}$, Cohen's $d = -.46$, $BF_{10} = 1.9 \times 10^3$; Fig. 6B).

Fig. 6. Feedback shortens both decision and confidence RTs. (A) Decision RT densities and box plots with individual subject data for each feedback group. We found a significant decrease in decision RTs for the Feedback group compared to the No Feedback group. (B) Confidence RT densities and box plots with individual subject data for Feedback and No Feedback groups. As with the decision RTs, confidence RTs were significantly faster for the Feedback group. Box plots show the median and the interquartile (25-75%) range with the whiskers indicating the 2-98% range.

Generalization of feedback's effects to a second task

Beyond establishing the effects of trial-by-trial feedback on different variables of interest, we wanted to explore if any of these effects generalize to a new task. For this reason, we included a second task that was designed to be very similar to the first one while tapping onto a different perceptual dimension. We therefore chose a task that requires subjects to discriminate color rather than shape. Specifically, subjects indicated whether more circles were red or blue, with other aspects of the task being identical to the first (Fig. 1). None of the subjects received feedback in Task 2. All analyses tested whether the effects of feedback in Task 1 extended to Task 2 in the absence of task-specific feedback.

Given that trial-by-trial feedback had no effect on perceptual or metacognitive sensitivity in the first task, predictably the Feedback and No Feedback groups from Task 1 did not differ in their perceptual and metacognitive sensitivities in Task 2 (Perceptual sensitivity: $d'_{\text{Feedback}} = 1.74$, $d'_{\text{No Feedback}} = 1.67$; $t(393) = 1.31$, $p = .19$, Cohen's $d = .13$, $BF_{01} = 4.0$; Metacognitive sensitivity: $meta-d'_{\text{Feedback}} = 1.15$, $meta-d'_{\text{No Feedback}} = 1.20$; $t(393) = -.57$, $p = .57$, Cohen's $d = -.06$, $BF_{01} = 7.7$; Fig. 7A-B). Unlike for Task 1, we did not observe significant learning over the course of the 150 trials in Task 2, potentially due to the smaller number of trials in Task 2 (Supplementary, Fig. S4). Thus, trial-by-trial feedback in Task 1 had no effect on perceptual or metacognitive sensitivity in either Task 1 or Task 2.

Fig. 7. Only the RT effects of feedback generalize to a new task. We examined the behavior of the Feedback and No Feedback groups from Task 1 on Task 2 where both groups received no feedback. We found no difference between the two groups in (A) Perceptual sensitivity d' , (B) Metacognitive sensitivity $meta-d'$, (C) Response bias (variability of criterion c), or (D) Metacognitive bias (d' -confidence correlation). On the other hand, the Feedback group from Task 1 exhibited (E) lower decision RTs and (F) lower confidence RTs compared to the No Feedback group from Task 1. Box plots show the median and the interquartile (25-75%) range with the whiskers indicating the 2-98% range. Each point in panel D represents a single subject; the dark gray line and the shaded gray area indicate the line of best fit and its 95% confidence interval, respectively.

More importantly, the presence of trial-by-trial feedback in Task 1 also had no effect on bias in the perceptual and metacognitive judgments in Task 2. Specifically, we found no difference between the variability of response criterion scores in Task 2 ($SD_{\text{Feedback}} = .37$, $SD_{\text{No Feedback}} = .36$; $F(200, 193) = 1.03$, $p = .81$, F-test of equality of variances; Fig. 7C). Similarly, the d' -confidence correlation that was modulated by feedback in Task 1 did not differ between the two groups in Task 2 ($r_{\text{Feedback}} = .002$, $p = .98$; $r_{\text{No Feedback}} = -.08$, $p = .30$; difference between the two r values: $p = .45$; Fig. 7D). Therefore, trial-by-trial feedback appears to allow subjects to reduce their bias in the specific task where such feedback is provided but does not allow for generalization even to a very similar task as long as the new task employs a different perceptual dimension.

Despite the lack of generalization of the effects of trial-by-trial feedback on bias, we found such generalization for RT. Indeed, subjects who received feedback in Task 1 were significantly faster on Task 2 even though subjects in Task 2 did not receive feedback ($RT_{\text{Feedback}} = 610$ ms, $RT_{\text{No Feedback}} = 657$ ms; $t(393) = -2.69$, $p = .007$, Cohen's $d = -.27$, $BF_{10} = 3.6$; Fig. 7E). However, it is worth noting that the decision RT difference between the two groups was 92 ms in Task 1 and decreased by almost half to 47 ms in Task 2 with the difference in the RT effect in the two tasks being significant ($t(376) = -4.93$, $p = 1.2 \times 10^{-6}$, Cohen's $d = -.51$). The effects of feedback on decision RTs in Task 2 also extended to confidence RTs ($RT_{\text{Feedback}} = 392$ ms, $RT_{\text{No Feedback}} = 434$ ms; $t(393) = -3.49$, $p = .0005$, Cohen's $d = -.35$, $BF_{10} = 37.2$; Fig. 7F). Thus, receiving trial-by-trial feedback allows subjects to respond faster not only on the task where feedback is given but also on a new task where the perceptual dimension changes.

Experiment 2

Experiment 1 demonstrated that trial-by-trial feedback provided over the course of 330 trials did not improve perceptual and metacognitive sensitivity. Nevertheless, it is possible that the effects of feedback take longer to manifest. Therefore, in a preregistered Experiment 2, we examined the effects of feedback over the course of seven days of training with 500 trials per day.

To determine whether feedback enhanced perceptual and metacognitive sensitivity over the course of training, we conducted a mixed ANOVA with a between-subjects factor of Group (Feedback and No Feedback) and a within-subjects factor of Day (1–7). We found no significant difference in d' between Feedback and No Feedback groups ($F(1, 58) = .43, p = .51, \eta_p^2 = .007, BF_{01} = 1.9$) with the Feedback group exhibiting slightly lower d' values (Fig. 8A). Perceptual sensitivity d' increased over the course of the seven days of training (main effect of Day: $F(6, 348) = 8.31, p < .0001, \eta_p^2 = .13, BF_{10} = 4.3 \times 10^5$), but this increase was similar for both groups (interaction between Group and Day: $F(6, 348) = 1.50, p = .18, \eta_p^2 = .03, BF_{\text{Inclusion}} = .21$). We specifically examined the rate of increase and again found no difference between the two groups ($\text{slope}_{\text{Feedback}} = .026, \text{slope}_{\text{No Feedback}} = .044, t(58) = -1.14, p = .26, \text{Cohen's } d = .30; B_{01} = 2.2$), with the slope for the No Feedback group being, in fact, numerically higher. Together these findings indicate that trial-by-trial feedback, even when provided over the course of multiple days and a total of 3,500 trials, does not improve sensitivity on a perceptual task.

Fig. 8. No effect of feedback on perceptual and metacognitive sensitivity in Experiment 2 where feedback was provided over seven days. There was no significant difference between Feedback and No Feedback groups in (A) perceptual sensitivity (d') or (B) metacognitive sensitivity ($meta-d'$). Error bars depict standard error of mean.

In addition to perceptual sensitivity (d'), we examined whether feedback had an effect on metacognitive sensitivity ($meta-d'$). We found no evidence for such a difference (main effect of Group: $F(1, 58) = 2.29$, $p = .14$, $\eta_p^2 = .04$, $BF_{01} = 1.2$, nor for an interaction between Group and Day: $F(6, 348) = .43$, $p = .86$, $\eta_p^2 = .007$, $BF_{Inclusion} = .02$) with the No Feedback group showing slightly higher $meta-d'$ scores (Fig. 8B). As with perceptual sensitivity, the rate of increase in $meta-d'$ was also similar between the two groups (slope_{Feedback} = .015, slope_{No Feedback} = .023, $t(58) = -.58$, $p = .56$, Cohen's $d = -.18$; $BF_{01} = 3.3$). Taken together, these results show that even a much longer period of trial-by-trial feedback does not affect either perceptual or metacognitive sensitivity, thus providing additional evidence against the notion that feedback acts via direct reinforcement mechanisms.

Discussion

Feedback is one of the most universal and powerful ways in which behavior can be improved across domains as diverse as educational, work, and sports settings, as well as a myriad of cognitive tasks (Hattie & Timperley, 2007). However, the mechanisms through which feedback affects behavior remain unclear. Specifically, two very different accounts co-exist in the literature: feedback could have a direct effect on performance through automatic reinforcement mechanisms (Petrov et al., 2005; Thorndike, 1927), or alternatively, feedback may only have an indirect effect mediated by a deliberate change in strategy (Vollmeyer & Rheinberg, 2005). Perceptual decision making and metacognition provide ideal testbeds for adjudicating between these two possibilities. In both cases, a direct effect via reinforcement mechanisms would predict an increase in perceptual and metacognitive sensitivity, whereas an indirect effect mediated by altered strategies would predict an exclusive effect on perceptual and metacognitive bias. Our results demonstrate that feedback does not affect perceptual or metacognitive sensitivity but reduces both perceptual and metacognitive bias. These findings strongly suggest that explicit feedback affects behavior only by allowing subjects to deliberately adjust their response strategy.

It is hardly surprising that trial-by-trial feedback allows people to alter their strategies. For example, in the absence of trial-by-trial feedback, a subject with a response bias for the letter 'X' may simply assume that there are more trials where the X's are dominant. However, even a few trials of receiving feedback about being incorrect when choosing 'X' can lead to a change in strategy where the subject avoids choosing the X's as frequently, thus reducing her original bias. Similarly, when an overconfident subject

receives feedback about being wrong, this feedback allows him to lower his confidence ratings, thus improving his confidence calibration.

A more controversial question is whether feedback also has a direct effect on behavior based on automatic reinforcement mechanisms. For example, much of the behaviorist tradition was able to achieve substantial success in training animals under the assumption that feedback exclusively acts to strengthen desirable stimulus-response associations outside of any need to posit cognitive constructs such as strategy (Skinner, 1938; Thorndike, 1927). Nevertheless, it is important to note that whether or not the experimenter poses the existence of cognitive strategies does not preclude their existence, both for humans and animals. In fact, our results strongly question the theory that feedback has a direct, automatic influence on behavior outside of the strategies adopted by the subject.

The current findings also help resolve a longstanding controversy in the field of perceptual decision making about the effects of feedback on sensitivity. A close examination of this literature reveals several issues that preclude drawing strong conclusions from previous studies. First and foremost, with total sample sizes ranging from 6 to 36 subjects further split into groups (except Shiu & Pashler, 1992), no study was sufficiently powered. Second, papers claiming to find a positive effect of feedback on sensitivity sometimes did not conduct a direct statistical comparison between the Feedback and No Feedback groups (Seitz et al., 2006). Finally, in the rare cases where a statistically significant difference was observed, the difference occurred for some but not other conditions and only for certain time points during training (Ball & Sekuler, 1987). These factors suggest that previous claims that feedback improves sensitivity may have

been based more on the seeming plausibility of this claim rather than the strength of the evidence. The converging evidence from our two well-powered experiments (consisting of over 200,000 trials each) imply that previous positive results could have been statistical aberrations caused by a combination of small sample sizes and pre-existing assumptions that feedback must improve perceptual sensitivity.

The present work is also the first to demonstrate that feedback reduces response bias in perceptual decision making. Note that our design allowed us to examine a true, non-directional reduction in response bias by examining whether the absolute value of the response criterion (that is, $|c|$) was smaller in the Feedback group (for a similar analysis, see (Hu & Rahnev, 2019). Instead, previous studies only examined directional changes in bias by exploring whether the response criterion had an overall increase or decrease (Aberg & Herzog, 2012; Goldhacker et al., 2014; Jones, Moore, Shub, & Amitay, 2015; Petrov et al., 2006; Wenger, Copeland, Bittner, & Thomas, 2008). However, a change in bias in a specific direction only signifies that subjects have started to choose one response option over the other; such an effect cannot in and of itself be interpreted as either a reduction or an increase in bias.

Our results fit well with the findings in fields outside of perceptual decision making and metacognition. For example, Kantner & Lindsay (2010) showed that trial-by-trial feedback does not affect memory sensitivity across different experimental conditions chosen to create a feedback advantage. Even more tellingly, Kluger & DeNisi (1996) conducted an extensive review and argued that feedback in education can have either a positive or a negative effect on subsequent performance. They showed that the widespread assumption in education research that feedback always leads to improvement is false and

suggested that the effect of feedback is determined by a number of mediating factors. Similar to the situation in perceptual learning, Kluger & DeNisi argued that weak and inconsistent results that would normally be strongly questioned were being widely cited in support of the notion that feedback is always beneficial, presumably due to the intuitive nature of this claim. Subsequent research has yielded more support to the notion that the effects of feedback in education are indeed mediated by other factors (Hattie & Timperley, 2007; Vollmeyer & Rheinberg, 2005). Together with these previous findings, our results strongly suggest that feedback does not automatically improve behavior; it only does so if the people being provided with feedback are willing and able to adjust their strategy for performing the task.

The current study also sheds light on the question of what components of feedback-induced learning generalize to a new task. We found that the effects of feedback on perceptual and metacognitive bias were task-specific and did not generalize to a new task. This lack of generalization may be driven by the fact that different kinds of processes lead to the response bias observed in different tasks. For example, the bias in the X/O task is likely independent from the bias in the red/blue task and thus reducing the former has no effect on the latter. Similarly, feedback indicating a tendency for over- or under-confidence in a given perceptual task may be considered irrelevant in a different task. Nevertheless, in our Experiment 1, feedback helped subjects reduce their RTs and this effect generalized to a new task. It is likely that this generalization was at least in part driven by the fact that the second task had the exact same timing structure as the first one. Thus, the strategy for responding faster in Task 1 could also be applied successfully in Task 2. Future studies

should examine whether this feedback-related RT advantage generalizes to tasks that differ in other aspects such as the timing and structure of the task.

In conclusion, we found that trial-by-trial feedback reduces response bias and improves confidence calibration but does not affect perceptual or metacognitive sensitivity. These results strongly suggest that feedback acts via modifying subjects' strategies for completing the task but has no direct, mediator-free effect on learning.

Transparency

Author Contributions

D.R. and N.H. designed the studies, and N.H. collected the data. N.H. analyzed the data. N.H. and D.R. worked together to interpret the results. N.H. drafted the manuscript and D.R. provided critical revisions. Both authors approved the final version of the manuscript for submission.

Declaration of Conflicting Interests

The author(s) declared no conflicts of interest with respect to the authorship or the publication of this article.

Acknowledgements

We thank Samuel Weiss-Cowie for help with data collection and Sriraj Aiyer, Matthew Davidson, Matthew Jaquiere and Nicholas Yeung for helpful feedback on our manuscript. This work was funded by National Institute of Health (grant R01MH119189) and Office of Naval Research (grant N00014-20-1-2622).

Open Practices

Experiment 1 was not formally preregistered; the preregistration for Experiment 2 can be accessed at <https://osf.io/efymg>. De-identified data for both experiments along with the code for data analysis and figure scripts are posted at <https://osf.io/94r87/>.

References

- Aberg, K. C., & Herzog, M. H. (2012). Different types of feedback change decision criterion and sensitivity differently in perceptual learning. *Journal of Vision*, 12(3), 1–11.
<https://doi.org/10.1167/12.3.3>
- Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., & Kievit, R. A. (2019). Raincloud plots: A multi-platform tool for robust data visualization. *Wellcome Open Research*, 4, 1–40.
<https://doi.org/10.12688/wellcomeopenres.15191.1>
- Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research*, 27(6), 953–965. <https://doi.org/10.22226/2410-3535-2017-3-209-217>
- Bang, J. W., Milton, D., Sasaki, Y., Watanabe, T., & Rahnev, D. (2019). Post-training TMS abolishes performance improvement and releases future learning from interference. *Communications Biology*, 2(1), 1–7. <https://doi.org/10.1038/s42003-019-0566-4>
- Bang, J. W., Sasaki, Y., Watanabe, T., & Rahnev, D. (2018). Feature-specific awake reactivation in human v1 after visual training. *Journal of Neuroscience*, 38(45), 9648–9657. <https://doi.org/10.1523/JNEUROSCI.0884-18.2018>
- Bang, J. W., Shekhar, M., & Rahnev, D. (2019). Sensory noise increases metacognitive efficiency. *Journal of Experimental Psychology: General*, 148(3), 437–452.
<https://doi.org/10.1037/xge0000511>
- Cortese, A., Amano, K., Koizumi, A., Kawato, M., & Lau, H. (2016). Multivoxel neurofeedback selectively modulates confidence without changing perceptual performance. *Nature Communications*, 7, 1–18. <https://doi.org/10.1038/ncomms13669>
- Fahle, M., & Edelman, S. (1993). Long-term learning in vernier acuity: Effects of stimulus

orientation, range and of feedback. *Vision Research*, 33(3), 397–412.

[https://doi.org/10.1016/0042-6989\(93\)90094-D](https://doi.org/10.1016/0042-6989(93)90094-D)

Fleming, S. M. (2017). HMeta-d: hierarchical Bayesian estimation of metacognitive efficiency from confidence ratings. *Neuroscience of Consciousness*, 2017(1), 1–14.

<https://doi.org/10.1093/nc/nix007>

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, 8, 1–9. <https://doi.org/10.3389/fnhum.2014.00443>

Goldhacker, M., Rosengarth, K., Plank, T., & Greenlee, M. W. (2014). The effect of feedback on performance and brain activation during perceptual learning. *Vision Research*, 99, 99–110. <https://doi.org/10.1016/j.visres.2013.11.010>

Green, D. M., & Swets, J. A. (1966). *Signal Detection Theory and Psychophysics*. New York: John Wiley & Sons Ltd. <https://doi.org/10.1177/001316446702700437>

Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of Educational Research*, 77(1), 81–112. <https://doi.org/10.3102/003465430298487>

Herzog, M. H., & Fahle, M. (1997). The role of feedback in learning a vernier discrimination task. *Vision Research*, 37(15), 2133–2141. [https://doi.org/10.1016/S0042-6989\(97\)00043-6](https://doi.org/10.1016/S0042-6989(97)00043-6)

Hu, M., & Rahnev, D. (2019). Predictive cues reduce but do not eliminate intrinsic response bias. *Cognition*, 192, 104004. <https://doi.org/10.1016/j.cognition.2019.06.016>

Ioannidis, J. P. A. (2005). Why most published research findings are false. *Getting to Good: Research Integrity in the Biomedical Sciences*, 2(8), 2–8.

<https://doi.org/10.1371/journal.pmed.0020124>

JASP Team. (2020). JASP (Version 0.9.2)[Computer software].

- Jones, P. R., Moore, D. R., Shub, D. E., & Amitay, S. (2015). The role of response bias in perceptual learning. *Journal of Experimental Psychology: Learning Memory and Cognition*, 41(5), 1456–1470. <https://doi.org/10.1037/xlm0000111>
- Judd, C. H. (1905). Practice without knowledge of results. *Psychological Review Monographs*, 2(7,Suppl.), 185–198.
- Kantner, J., & Lindsay, D. S. (2010). Can corrective feedback improve recognition memory? *Memory and Cognition*, 38(4), 389–406. <https://doi.org/10.3758/MC.38.4.389>
- Kluger, A. N., & DeNisi, A. (1996). Effects of feedback intervention on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin*, 119(2), 254–284. <https://doi.org/10.1037//0033-2909.119.2.254>
- Macmillan, N., & Creelman, C. (2005). *Detection Theory*. LEA Publishing.
- Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, 21(1), 422–430. <https://doi.org/10.1016/j.concog.2011.09.021>
- Masson, M. E. J. (2011). A tutorial on a practical Bayesian alternative to null-hypothesis significance testing. *Behavior Research Methods*, 43(3), 679–690. <https://doi.org/10.3758/s13428-010-0049-5>
- Metcalf, J. E., & Shimamura, A. (1994). *Metacognition: knowing about knowing*. Cambridge, MA: MIT Press.
- Petrov, A. A., Doshier, B. A., & Lu, Z. L. (2005). The dynamics of perceptual learning: An incremental reweighting model. *Psychological Review*, 112(4), 715–743. <https://doi.org/10.1037/0033-295X.112.4.715>

- Petrov, A. A., Doshier, B. A., & Lu, Z. L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Research*, 46(19), 3177–3197.
<https://doi.org/10.1016/j.visres.2006.03.022>
- Pritchard, R. D., Jones, S. D., Roth, P. L., Stuebing, K. K., & Ekeberg, S. E. (1988). Effects of group feedback, goal setting, and incentives on organizational productivity. *Journal of Applied Psychology*, 73(2), 337–358.
- Rahnev, D., & Denison, R. N. (2018). Suboptimality in Perceptual Decision Making. *Behavioral and Brain Sciences*, 1–107.
- Rahnev, D., Desender, K., Lee, A. L. F., Adler, W. T., Aguilar-Lleyda, D., Akdoğan, B., ... Zylberberg, A. (2020). The Confidence Database. *Nature Human Behaviour*.
<https://doi.org/https://doi.org/10.1038/s41562-019-0813-1>
- Rahnev, Dobromir, Koizumi, A., McCurdy, L. Y., D'Esposito, M., & Lau, H. (2015). Confidence Leak in Perceptual Decision Making. *Psychological Science*, 26(11), 1664–1680.
<https://doi.org/10.1177/0956797615595037>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108.
<https://doi.org/10.1037/0033-295X.85.2.59>
- Rouault, M., Dayan, P., & Fleming, S. M. (2019). Forming global estimates of self-performance from local confidence. *Nature Communications*, 10(1), 1–11.
<https://doi.org/10.1038/s41467-019-09075-3>
- Seitz, A. R., Nanez, J. E., Holloway, S., Tsushima, Y., & Watanabe, T. (2006). Two cases requiring external reinforcement in perceptual learning. *Journal of Vision*, 6(9), 966–973. <https://doi.org/10.1167/6.9.9>
- Shibata, K., Yamagishi, N., Ishii, S., & Kawato, M. (2009). Boosting perceptual learning by

fake feedback. *Vision Research*, 49(21), 2574–2585.

<https://doi.org/10.1016/j.visres.2009.06.009>

Shiu, L. P., & Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Perception & Psychophysics*, 52(5), 582–588.

<https://doi.org/10.3758/BF03206720>

Skinner, B. F. (1938). *The behavior of organisms: an experimental analysis*. Appleton-Century.

Thorndike, E. L. (1927). The Law of Effect. *The American Journal of Psychology*, 39(1/4), 212–222.

van Ravenzwaaij, D., Donkin, C., & Vandekerckhove, J. (2017). The EZ diffusion model provides a powerful test of simple empirical effects. *Psychonomic Bulletin and Review*, 24(2), 547–556. <https://doi.org/10.3758/s13423-016-1081-y>

Vollmeyer, R., & Rheinberg, F. (2005). A surprising effect of feedback on learning. *Learning and Instruction*, 15(6), 589–602. <https://doi.org/10.1016/j.learninstruc.2005.08.001>

Wagenmakers, E.-J., van der Maas, H. L. J., & Grasman, R. P. P. P. (2007). An EZ-diffusion model for response time. *Psychonomic Bulletin & Review*, 14(1), 3–22.

Wenger, M. J., Copeland, A. M., Bittner, J. L., & Thomas, R. D. (2008). Evidence for criterion shifts in visual perceptual learning: Data and implications. *Perception and Psychophysics*, 70(7), 1248–1273. <https://doi.org/10.3758/PP.70.7.1248>

Wright, W. R. (1906). Some effects of incentive on work and fatigue. *Psychological Review*, 13, 23–34.