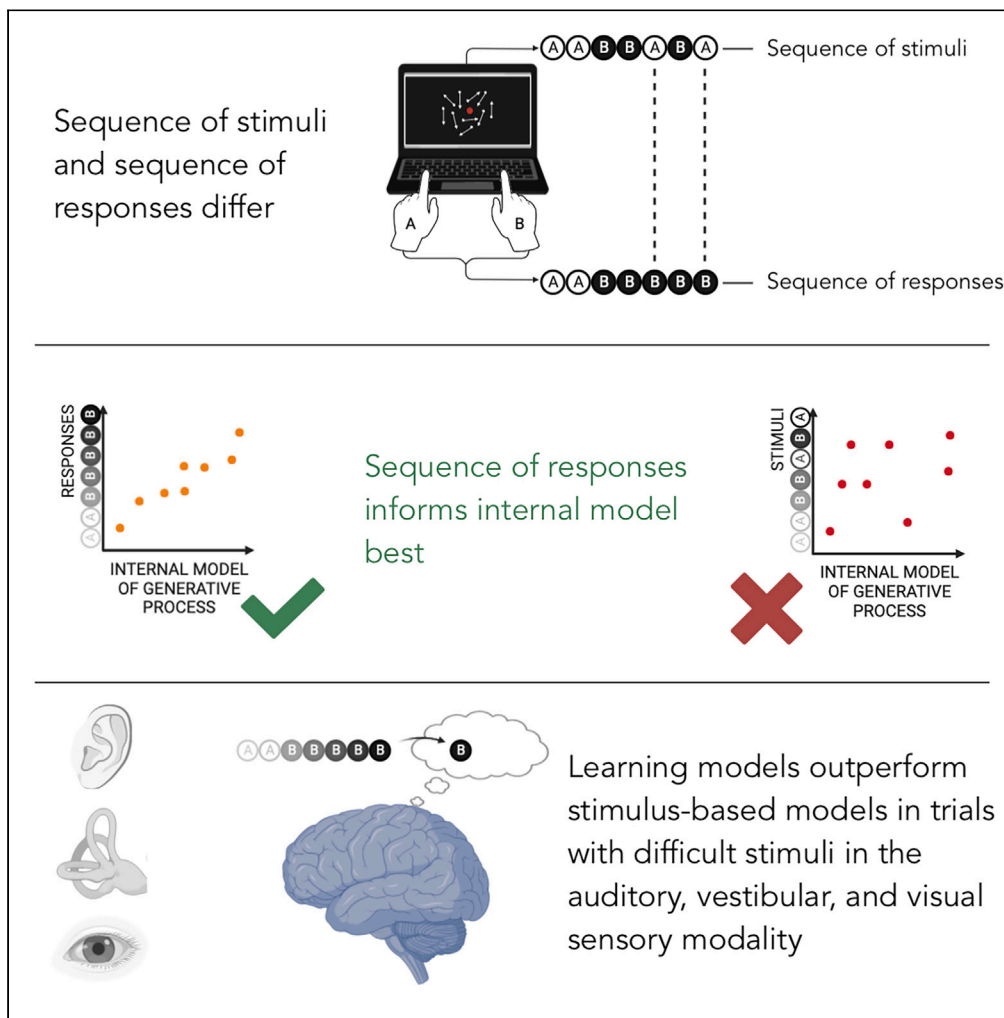


Article

Probabilistic integration of preceding responses explains response bias in perceptual decision making



Daniel Schlunegger, Fred W. Mast

daniel.schlunegger@unibe.ch

Highlights

Recent decisions instead of stimuli drive sequential choice effects

Serial biases are based on statistical regularities of the decision variable

Learning models outperform stimulus-based models in trials with difficult stimuli

Results are consistent across the auditory, vestibular, and visual sensory modality



Article

Probabilistic integration of preceding responses explains response bias in perceptual decision making

Daniel Schlunegger^{1,2,*} and Fred W. Mast¹

SUMMARY

Expectations of sensory information change not only how well but also what we perceive. Even in an unpredictable environment, the brain is by default constantly engaged in computing probabilities between sensory events. These estimates are used to generate predictions about future sensory events. Here, we investigated the predictability of behavioral responses using three different learning models in three different one-interval two-alternative forced choice experiments with either auditory, vestibular, or visual stimuli. Results indicate that recent decisions, instead of the sequence of generative stimuli, cause serial dependence. By bridging the gap between sequence learning and perceptual decision making, we provide a novel perspective on sequential choice effects. We propose that serial biases reflect the tracking of statistical regularities of the decision variable, offering a broader understanding of this phenomenon.

INTRODUCTION

From traffic lights to weather phenomena to speech perception, many events in real life follow a highly structured temporal sequence. Statistical regularities commonly determine what follows next. By extracting these regularities from the environment, the brain builds expectations regarding incoming sensory information.^{1–4} Numerous studies on sequence learning provided compelling evidence that the brain is capable of detecting, using, and leveraging statistical regularities. Findings include serial dependencies in psychophysical tasks,^{5–18} variations in reaction times as a function of statistical properties of a sequence,^{19–22} the biased perception of randomness^{23–25} and surprise-like signals in electrophysiological (EEG) or magnetoencephalographic (MEG) data.^{26–29} The local transition probability (TP) model³⁰ unifies these seemingly unrelated findings. It suggests that the brain's constant engagement in trying to infer a non-stationary transition probability matrix must represent a core computation of human sequence knowledge. The local integration of the most recent transition probabilities about sensory events causes the brain to predict incoming sensory information even when the stimuli are embedded in a completely unpredictable sequence. In psychophysical experiments, the observer is usually confronted with an unpredictable environment that supposedly ensures that stimuli cannot be predicted on the basis of statistical regularities.⁴ Many real-life events, however, are highly structured in space and time, so it is plausible that the brain makes predictions by default rather than distinguishing when it can and cannot predict.

We hypothesize that the findings from Meyniel et al.³⁰ provide a previously unexplored framework for investigating serial dependencies in perceptual decision making. We examine how the brain's default predictive mode unfolds and translates into choice behavior. Relying on the principles of Bayesian inference, we use derivatives of the models previously employed in Meyniel et al.³⁰ and Maheu et al.²⁷ to capture serial dependencies in psychophysical tasks.^{8,12} We suggest that sequential choice effects arise from human observers possessing an internal model of the generative process underlying the sequences of observations they encounter.^{23,24} These observers combine their prior beliefs with stimulus information in a way that approximates Bayes' rule.^{12,30–33} When accurately performing a psychophysical task, this internal model is most accurately reflected by the sequence of responses rather than the potentially unknown sequence of generative stimuli. In addition, these responses are weighted according to their recency, leading to leaky integration.^{13,30} The use of a leak factor approximates the Bayes-optimal solution and allows for flexibility in case the statistical underpinnings of the environment change.

¹Department of Psychology, University of Bern, 3012 Bern, Switzerland

²Lead contact

*Correspondence: daniel.schlunegger@unibe.ch

<https://doi.org/10.1016/j.isci.2023.107123>






Figure 1. Behavioral tasks from three independent experiments

(A) Auditory two-alternative forced-choice (2AFC) task: Participants ($n = 75$) had to decide whether the syllable /ka/ or the syllable /to/ was presented. Stimuli were masked with varying degrees of white noise resulting in six different signal-to-noise ratios per stimulus. The task was performed on a standard laptop and participants were wearing over-ear headphones. After each trial, they indicated the perceived syllable by pressing a key without receiving feedback.

(B) Vestibular 2AFC task from Klaus et al.³⁹: Participants ($n = 20$) were blindfolded to eliminate visual cues and equipped with headphones delivering white noise to eliminate auditory cues. Stimuli consisted of roll, pitch, and yaw rotations delivered by means of a motion platform. Participants had to indicate the motion direction by button press during each trial without receiving feedback.

(C) Visual 2AFC task from Braun et al.⁸: Participants ($n = 22$) were presented with a dynamic random dot pattern of variable direction and coherence. One beep indicated the start of the evidence interval, which contained some degree of coherent motion. A second beep indicated the evidence offset and the start of the response interval in which participants had to indicate the perceived net motion (up versus down). Figure (C) adapted from Braun et al.⁸

A large body of work has provided different explanations for the modulation of response biases, such as context and attention,¹¹ adaptation,⁵ arousal in humans^{17,34} and arousal in animals,³⁵ strength of sensory evidence,⁶ fluctuations between modes of sensory processing,³⁶ and confidence.^{8,17,37} Most likely, response biases are modulated by simultaneous mechanisms.⁷ However, it is still unclear how the probabilistic integration of preceding responses can help explain serial dependencies. To this end, we bridge the gap between two different lines of research: sequence learning on the one hand and perceptual decision making and psychophysics on the other hand.

We used adapted versions of previously used observer models (i.e., learning models)^{27,30} to predict participants' responses in three psychophysical experiments involving either auditory, vestibular, or visual stimuli (Figures 1, and S1). These observer models estimate different statistics that can be used to describe a particular sequence. Sequences can be described at different levels of abstraction.³⁸ Following the findings from Meyniel et al.³⁰ and Maheu et al.,²⁷ we focus on the most basic statistics that can be used to describe sequences: The item frequency (IF), the alternation frequency (AF), and transition probabilities (TP). The following example sequence is used to illustrate the characteristics of each statistic: BBBBABBBAAAA. IF is the simple count of each observation $p(B) = 1 - p(A) = 7/12$, whereas AF counts whether the successive observation is the same or different $p(\text{repetition}) = p(1 - \text{alternation}) = 8/11$. IF is ignorant of the number of

repetitions and alternations, and AF is ignorant of the number of A's and B's. TP captures both these aspects. Estimating TP requires tracking two statistics simultaneously and applying one of them depending on the context $p(B|A) = 1 - p(A|A) = 1/4$ and $p(B|B) = 1 - p(A|B) = 5/7$ (see [STAR Methods](#) for details and see Meyniel et al.³⁰ for a formal description of these and additional models). These estimates are used to predict the next response trial-by-trial. More specifically, IF, AF, and TP are used to compute the predictive likelihood of the next response. In the case of AF and TP, the predictor needs to be transformed so that response = 1 is predicted on each trial (see [STAR Methods](#) for details). Note that we refer to the predictive likelihoods computed by the observer models as IF, AF, and TP. For all three observer models we can change the input to examine whether the observer model is better informed by participants' responses or by the generative sequence of stimuli. Another aspect to consider when using models from the sequence learning literature is that participants are often presented with stimulus intensities at threshold level, which is typical for psychophysical research. Indeed, we will show that the maximally diagnostic conditions to arbitrate between observer models and inputs for observer models are the difficult, and not the easy trials; thereby showing a more informative model comparison strategy than the abrupt comparison on all trials altogether.

Taken together, the current study examines three learning models (TP, IF, and AF) in three sensory modalities (auditory, vestibular, and visual) for two trial difficulties (easy, difficult) by using two different model inputs (responses, stimuli).

RESULTS

In this study, we employed two complementary analyses to investigate our research question. In the main text, we report the log-likelihood analysis and in the supplemental information, we report the GLM analysis (probit regressions). We predict participants' responses on a trial-by-trial basis for each participant, type of learning model, input of learning model, and timescale of integration using a first- and second-level analysis approach (i.e., subject-level followed by group-level analysis). To examine the effect of stimulus intensity in the log-likelihood analysis, we categorized data into easy and difficult trials based on participants' categorization accuracy (see [STAR Methods](#) and [Figure S1](#) for details).

Model comparison

Advantages for model comparison using a log-likelihood approach are its robustness and clear interpretability. Models reach high likelihood when participants' choice behavior positively aligns with the predictive likelihood of the choice obtained from the different learning models. To reduce any ambiguity in terminology we clarify the following terms: In (1) response-based models, participants' responses are predicted by the predictive likelihood obtained from learning models that had participants' preceding responses as input. In (2) stimulus-based models, participants' responses are predicted by the predictive likelihood obtained from learning models that had preceding generative stimulus identities as input. Note that these predictions do not contain any actual stimulus information. Lastly, in (3) stimulus-only models, responses are predicted by the actual stimulus identity (i.e., type of syllable (i.e., /ka/ or /to/), direction of dot motion (i.e., up or down), and direction of passive motion stimulus (i.e., left or right and forward or backward).

The initial stage in evaluating the validity of the predictions obtained from the response-based learning models is comparing them to a stimulus-only model. As anticipated, the stimulus-only model exhibits higher prediction accuracy in easy trials where perceptual processing is unambiguous and clear. Conversely, in difficult trials where sensory information is ambiguous and uncertain, the learning models outperform the stimulus-only model in conforming to participants' choice behavior ([Figure 2](#), [Table S1](#)).

Following the confirmation of superior performance by the stimulus-only model in easy trials and the learning models in difficult trials, the subsequent stage involves contrasting the response-based model with a stimulus-based model for further analysis. This comparison aims to shed light on the question if the learning models are better informed by participants' preceding responses or by the preceding sequence of generative stimuli.

Improvement in model fit by using participants' responses as input

In active tasks where participants are required to make a conscious behavioral response, such as pressing a button, the question is whether the performance of the learning model is improved by using the generative sequence of stimuli or participants' responses as input. Informing the learning model with participants'

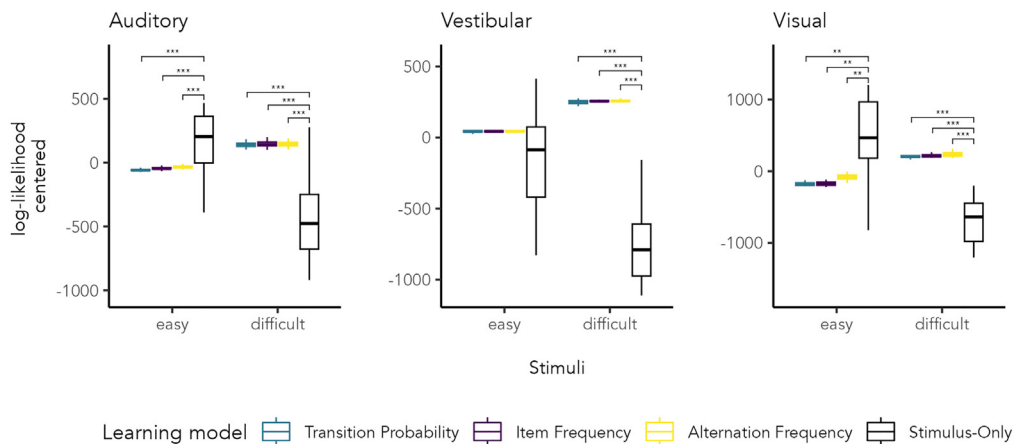


Figure 2. Learning models versus stimulus-only model

Comparison of log-likelihood values between response-based learning models (TP, IF, and AF) versus a stimulus-only model. Absolute log-likelihood values are centered and outliers are hidden for plotting purposes. Higher values indicate better performance. In easy trials, the stimulus-only model performs better, whereas in difficult trials, the learning models based on participants' preceding responses outperform the stimulus-only model, except for easy trials in the vestibular dataset where some participants exhibited inaccurate performance. Despite this, we retained the data as is without removing participants based on their performance. Corresponding statistics are reported in Table S1. Significance: *** = $p < 0.001$; ** = $p < 0.01$.

responses can be interpreted as a meta-observer making future predictions about its own actions based on observations of those actions, without access to the actual stimuli. We investigated this question for all three learning models. To ensure a fair comparison, for each participant and model separately, we tested different values for the leak factor (ω) to find the timescale of integration that best fits the data. Figure 3A shows that across all modalities and learning models, using participants' responses as input provides a better fit of the data. To quantify this effect, we compute the differences in log-likelihoods between models that differ in their input and obtain $\Delta\log\text{-likelihood}$ where a lower value indicates that the response-based model performs better. In difficult trials, this difference ($\Delta\log\text{-likelihood}$) was significant for all modalities and learning models ($p < 0.05$) except for the TP model in the vestibular dataset ($p > 0.05$), complete statistics are provided in Table S2. In easy trials, this difference ($\Delta\log\text{-likelihood}$) is either not significant or the stimulus-based model performs better. Note that this comparison is not informative of which learning model performs best. Figure 4A provides an example; the effect emerges when the models using different inputs make orthogonal predictions, and the behavior of participants more often matches the predictions of the response-based model. An exhaustive characterization of orthogonal predictions between the response-based model and the stimulus-based model is provided in Figure 4B.

Differences between learning models

Having confirmed that informing the models with the sequence generated by participants' responses notably enhances model fits compared to using the generative sequence of stimuli, we proceeded to analyze the absolute log-likelihood values for each learning model and sensory modality, separately for easy and difficult trials (Figure 3B). Although we have already determined that a stimulus-only model is the preferred model for explaining participants' choice behavior in easy trials, our attention in this analysis is directed toward the difficult trials. Remarkably, across all modalities, the same pattern emerges, with both IF and AF models significantly outperforming the TP model (according to Wilcoxon signed rank tests) when ω is chosen to best fit the data. This is intriguing, as both statistics are subsumed in the two-dimensional space of transition probabilities.^{27,30} IF and AF both capture aspects of serial influence.

Timescales of integration between learning models

Following the results from Meyniel et al.³⁰ and Maheu et al.,²⁷ we examined if a local or a global integration based on preceding observations does provide the best fit of the data. The leak factor ω determines the strength of decay and thus integration over time. A smaller value for ω corresponds to a more local integration, which means that recent responses are weighted more strongly in the inference process. A truly global integration corresponds to an ideal observer with perfect memory (i.e., $\omega = \infty$). We compared

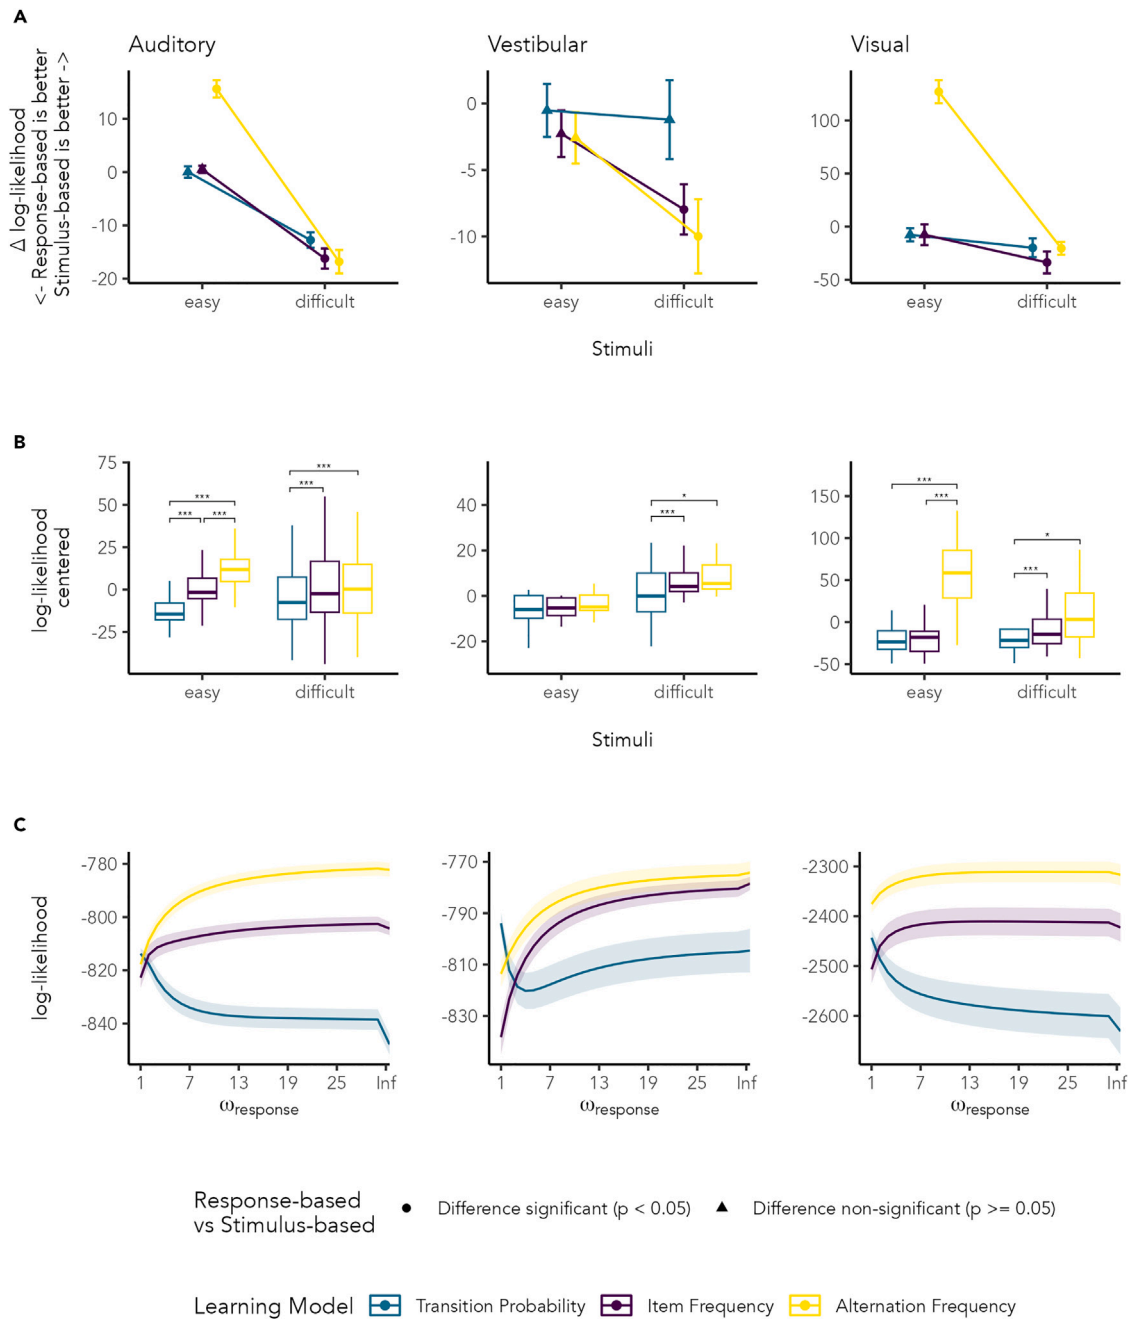


Figure 3. Learning models

(A) Comparison of model input (responses versus generative sequence of stimuli) based on accuracy (Figure S1). The data were categorized into easy and difficult trials, and the mean difference in log-likelihood (Δ log-likelihood) between the response-based model and the stimulus-based model was computed across participants. To ensure a fair comparison, the timescale of integration (ω) for both models was chosen to best fit the data. Results show that for difficult trials across modalities and learning models, the response-based model achieved a better fit of the data. Note that this analysis does not provide information about which learning model performed best. Error bars indicate the standard error of the mean (\pm SEM). Corresponding statistics are reported in Table S2. (B) Comparison of absolute log-likelihood values for each response-based (learning) model and each sensory modality, separately for easy and difficult trials. Here, the focus is on difficult trials, in which the IF and AF model perform best across modalities. Note that variance explained by stimulus identity is not considered in this analysis, which is especially important when comparing learning models in easy trials. Outliers are hidden for plotting purposes. Corresponding statistics are reported in Table S3.

(C) Fixing the timescale of integration (ω) across participants: Consistent with previous studies,^{27,30} TP favors a local timescale of integration, whereas IF and AF favor global timescales of integration. Shadings indicate the standard error of the mean (\pm SEM).

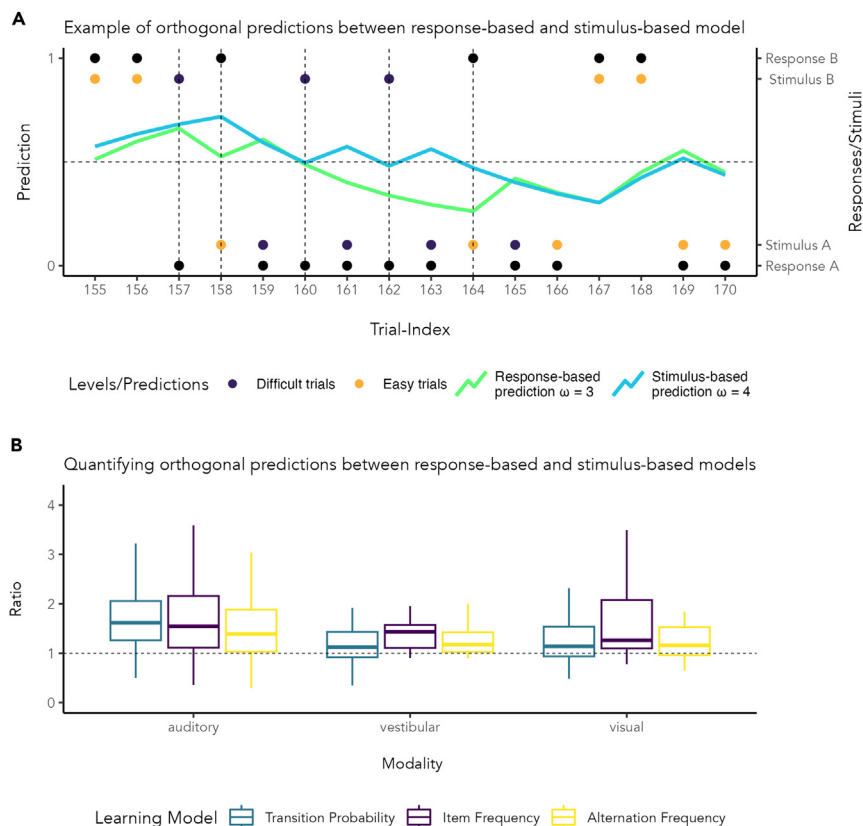


Figure 4. Orthogonal predictions between models that differ in their input

(A) Illustration of orthogonal predictions between the response-based model and stimulus-based model using the item frequency model in the auditory dataset in the fully stochastic conditions (see STAR Methods for full description of experimental conditions). The x axis shows the trial index from a random subject, whereas the y axis (left-hand side) displays the prediction. Predictions below 0.5 indicate prediction of response A, whereas predictions above 0.5 indicate prediction of response B. The y axis (right-hand side) shows the true sequence of stimuli (Stimulus A/ka/and Stimulus B/to/) as well as participants' responses (A and B). Vertical lines indicate incorrect responses. Orthogonal predictions can occur after incorrect responses, as evident in trial 161 and trial 163 where the response-based model outperforms the stimulus-based model. The best fitting value for ω was determined based on log-likelihood analysis to ensure a fair comparison.

(B) Summary of the effect described in (A). Instances of orthogonal predictions between the best fitting response-based model and the best fitting stimulus-based model were counted for all participants, modalities, and learning models. The ratio of how often one model outperforms the other was computed. A ratio higher than 1 indicates a better performance of the response-based model, whereas a ratio smaller than 1 indicates a better performance of the stimulus-based model. The finding that preceding responses, rather than the generative sequence of stimuli, are informative of participants' responses is observed across different datasets and learning models. Outliers are hidden for plotting purposes.

the log-likelihood averaged over participants as a function of the temporal integration of preceding responses for each modality and learning model separately (with fixed ω across participants). The results from these analyses suggest that local integration of TP and a global integration of IF provides the best results, adding further evidence to the findings from Meyniel et al.³⁰ and Maheu et al.²⁷ We further show that the AF model also favors longer time scales of integration and albeit not significant, even better than the IF model. Best fitting values for ω are provided in Figure S2.

Quantifying conditional probabilities of participants' preceding responses in different experimental conditions

The auditory and visual experiments entailed sequences of stimuli with different underlying statistical regularities. Experimental conditions consisted of orthogonal variations in the two-dimensional space of transition probabilities.²⁷ For example, a frequency-biased condition in which one of the two stimuli was presented more frequently than the other, or a repetition-biased condition in which there were more repetitions than

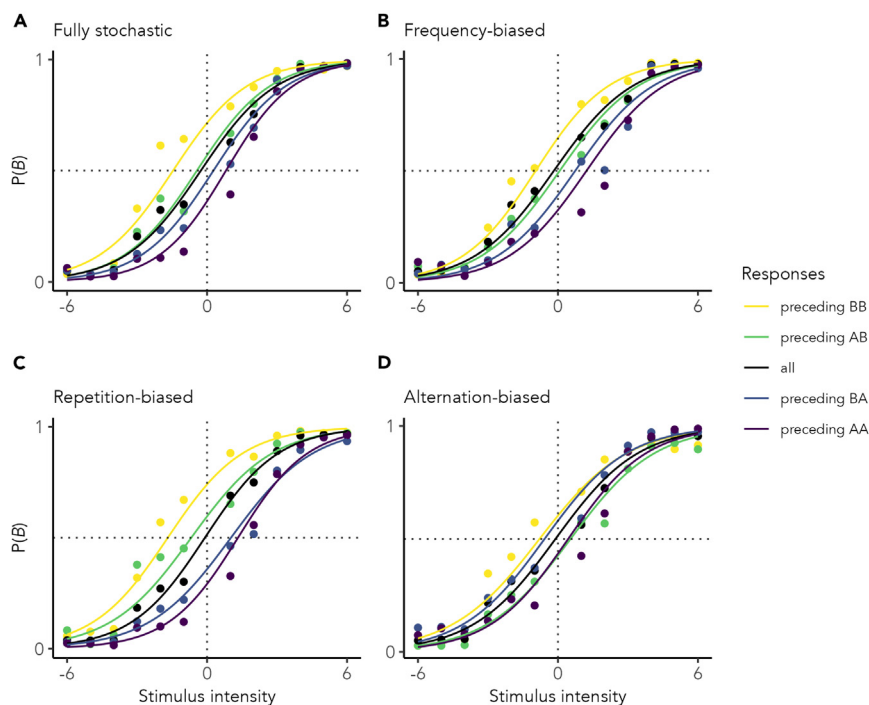


Figure 5. Psychometric functions conditioned on the two preceding responses

The pairs in the legend refer to $[rt-2, rt-1]$ where r corresponds to the response and t corresponds to the trial number. Corresponding statistics are reported in [Table S4](#).

(A) Fully stochastic, same number of A's and B's, and the same number of repetitions and alternations: Participants show a response bias even in the complete absence of any statistical bias, in this case the tendency to repeat their responses. Importantly, this effect is driven by the two preceding responses.

(B) Frequency-biased, one of the two stimuli occurred more often: Counterbalanced across participants, but data were transformed so that stimulus B appeared more often for all participants. Like in the fully stochastic condition, participants tend to repeat their responses in a streak-like manner.

(C) Repetition-biased, more repetitions than alternations: Because the generative sequence matches the tendency to repeat the preceding response, we observe the same effect as in the fully stochastic and the frequency-biased condition, but stronger.

(D) Alternation-biased, more alternations than repetitions: Here, the results well illustrate that the same analysis, but conditioned only on the preceding response (cf.^{8,41}), does not reveal that the effect is driven by the two preceding responses. For example, if the preceding responses AA and BA are combined, the effect disappears.

alternations. We complement and highlight further aspects of choice behavior by quantifying the conditional probabilities of participants' responses for each experimental condition separately. We report the results for the auditory dataset in the main text and summarize the results for the visual dataset in [Figure S5](#) and [Table S5](#). This analysis was not performed for the vestibular dataset, because conditions did not have different underlying generative stimulus probabilities. We computed and visualized psychometric functions conditional on the two preceding responses, as shown in [Figure 5](#). Psychometric functions were computed using a Generalized Linear Model (GLM) with logit link function.⁴⁰ This analysis shows various aspects of human choice behavior. In terms of group-level analysis, when considering all responses together (solid black lines), participants do not exhibit a bias, as all psychometric curves intersect the midpoint of the coordinate system ([Figure 5](#)). However, when examining responses conditioned on the preceding responses, participants tend to repeat their responses more frequently in the fully stochastic, frequency-biased, and repetition-biased conditions. Notably, this tendency to repeat responses is primarily driven by instances where a repetition of two responses occurred just before the current response, resembling a streak of repetitions. In contrast, in the alternation-biased condition, the response bias depends on whether the preceding transition was a repetition or an alternation. The same analysis, when conditioned only on the preceding response as done in previous studies,^{8,41} does not reveal the effect being driven by the two preceding responses. For instance, if the preceding responses AA and BA are combined, the effect disappears, suggesting that the bias is dependent on the transition type from the preceding responses. Regardless of the experimental condition, the differences in the

horizontal shifts in the psychometric functions between the pair of preceding responses AA versus BA and the pair of preceding responses BB versus AB were significant in six out of eight cases. Results from significance tests are provided in [Table S4](#).

DISCUSSION

This study aimed to investigate the effects of statistical learning on perceptual decision making in 2AFC tasks. Using Bayesian inference, we found that probabilistic integration of preceding responses underlies response biases in perceptual decisions, as shown through computational modeling of behavioral data across modalities and participants. Results showed an improvement in model fit when using the predictive likelihood computed based on participants' responses as input for different learning models, consistent with previous research indicating that observers are strongly influenced by their previous responses rather than by previous stimuli.^{8,12} That is, decision outcomes, instead of generative stimuli, had a serial influence on human participants' choice. By bridging the gap between sequence learning and perceptual decision making, we provide a potential new perspective on interpreting serial influence. According to this view, serial bias tracks statistical regularities of the decision variable, instead of considering the serial effect as a smoothing process to the latent environment variables, as environments are stationary or change gradually but slowly. Previous understanding is extended through this broader perspective.

We showed that the stimulus-only model performed better than the learning models in trials with easy stimuli, where sensory information was unambiguous and participants accurately performed the task with minimal to nonexistent response bias. In trials with difficult stimuli, where sensory information is insufficient, the learning models showed better performance, meaning that they are able to capture serial dependence. Our results offer a potential firsthand explanation from our analysis on why response-models outperform stimulus-based models. The internal model about the generative sequence of observations is better informed by the sequence of responses, as participants produced the response sequence through their motor actions (cf.¹⁶). We conclude that the active (re-)production by means of motor responses is creating a more accessible representation that may overwrite any potentially false representation of the generative sequence of stimuli. Furthermore, the differences between the sequence of responses and the sequence of stimuli are only detectable when post-decision metacognitive capabilities, such as decision confidence^{42,43} and post-error detection,^{44–46} are perfectly accurate. Findings were further supported by a complementary GLM analysis that included stimulus identity, stimulus intensity, type of learning model and the interaction between stimulus intensity and learning model as predictors ([Figure S4](#)). To investigate the possibility whether participants had a simple preference for repetitions as may have been indicative from visual inspection of [Figure 5](#), we compared all three learning models to a model that, instead of the predictions obtained from the learning models, included a simple preference for repetition. Probabilistic learning models do not differ significantly from each other when using a GLM approach, but they all differ significantly from the model that includes a preference for repetition. On closer examination of the psychometric curves, it becomes evident that the effect is not solely driven by a repetition bias. For example, in the alternation-biased condition, it depends on whether the preceding transition was a repetition or an alternation (i.e., depending on the specific preceding transition type).

Following Meyniel et al.³⁰ and Maheu et al.,²⁷ we systematically compared learning models that estimated item frequency, alternation frequency, and transition probabilities. The results extend previous findings on passive processing of unambiguous auditory stimuli by Maheu et al.²⁷ They found that late post-stimulus brainwaves were best explained by local integration of transition probabilities and that early post-stimulus brainwaves were best explained by a global integration of item frequency, consistent with our results as shown in [Figure 3C](#). When we allowed the timescale of integration to vary between participants and aimed for optimal performance of the models, we observed that the IF model and the AF model were most effective in predicting participants' choice behavior in difficult trials. When participants are presented with difficult perceptual decisions, their sensitivity to both frequency of items and the number of repetitions and alternations over longer periods of trials (longer timescales of integration) translates into specific choice behavior that aligns with the predictions of the item frequency and alternation frequency models. The log-likelihood approach ensures that the models achieve high likelihood only when participants' choice behavior positively aligns with the model's predictions. Of interest, our results suggest that the concept of transition probabilities may not be necessary, as both item frequency and alternation frequency capture aspects of serial influence. This finding is intriguing because transition probabilities capture the same information as the two simpler competitors, which is the frequency of items and their co-occurrence. In addition, transition probabilities also capture the serial order, whereas IF and AF do not take into account these dependencies between items in terms of their order or context.

Previous research has examined the level of information processing at which sequential dependencies emerge. Several studies suggest that the serial effect occurs no earlier than the perceptual level,^{47,48} and may require the involvement of working memory or the conscious processing level.^{49,50} Our finding that recent decision outcomes, rather than generative stimuli, generate serial dependencies underpins current knowledge in the field of sequential dependencies. Although there is a plethora of important work,^{5–9,11,12,17,18,37,51–55} we argue that the parsimonious model presented here will inspire future work on sequential dependence. The studies cited here either find that participants are more likely to repeat or alternate their responses, or they find that participants employ a win-stay, win-switch, loose-stay, or loose-switch strategy, all of which can be quantitatively captured by the analyses employed in this study.

Looking through the lens of sequence learning provides a useful framework for investigating serial choice effects in perceptual decision making. The present work sets the stage for incorporating models of sequence learning and the concept of probabilistic integration of preceding responses into other models of perceptual decision making (e.g., drift-diffusion models) and the parsimonious characters make them potential candidates for unifying the results of previous studies.

Limitations of the study

We modeled the non-stationarity of the transition matrix via leaky integration. This could also be modeled hierarchically by a dynamic belief in the possibility of a sudden change in the generative process (i.e., volatility). In other studies, both possibilities yield similar results, and both can be implemented in a biologically plausible manner.^{30,56,57} How exactly decision confidence,^{42,43} post-error detection,^{44–46} and metacognitive capabilities in general relate to the findings presented here remains to be investigated. The results presented here provide evidence that the concept of transition probabilities might not be required. However, to conclusively address this issue, future research will be necessary and incorporate experimental designs that are able to specifically test this hypothesis. Moreover, alternative measures to assess the style of response bias could be helpful. The GLM analysis also revealed that some participants responded systematically but opposite to the predictions of the learning model, resulting in significant negative regression coefficients (Figure S4). It is possible that participants for whom we obtain negative regression weights for learning models oppose the local statistics of their responses and thereby attempt to counteract local biases, so that the local sequence of their choices appears locally more random. Similar results have been observed in the literature on the perception of randomness (cf.²⁴). More advanced statistical analysis techniques may be required to fully disentangle the different learning models. A GLM approach that includes two predictions obtained from different learning models, faces problems of multicollinearity. This is because IF and TP are highly correlated, which is especially true when ω is low. For example, when $\omega = 1$, then the correlation between IF and TP ≈ 0.97 . As ω increases, the correlation between TP and IF decreases monotonically.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- [KEY RESOURCES TABLE](#)
- [RESOURCE AVAILABILITY](#)
 - Lead contact
 - Materials availability
 - Data and code availability
- [EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS](#)
 - Participants of experiment 1
 - Participants of experiment 2
 - Participants of experiment 3
- [METHOD DETAILS](#)
 - Observer models
 - Method details of experiment 1
 - Method details of experiment 2
 - Method details of experiment 3
- [QUANTIFICATION AND STATISTICAL ANALYSIS](#)
 - Observer models
 - Log-likelihood analysis
 - Easy vs. difficult trials

- Generalized Linear Models
- Estimating psychometric curves
- Statistical tests

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2023.107123>.

ACKNOWLEDGMENTS

We thank Ariel Zylberberg, Vincent Schönfelder, Andrew Ellis, Stefan Thoma, Lilla Gurtner, and Sebastian Strahm for useful discussions.

AUTHOR CONTRIBUTIONS

Conceptualization, D.S. and F.W.M.; Methodology, D.S.; Validation, D.S.; Formal Analysis, D.S.; Investigation, D.S.; Writing – Original Draft, D.S.; Writing – Review and Editing, D.S. and F.W.M.; Funding Acquisition, F.W.M.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: August 21, 2022

Revised: January 12, 2023

Accepted: June 9, 2023

Published: June 14, 2023

REFERENCES

1. de Lange, F.P., Heilbron, M., and Kok, P. (2018). How Do Expectations Shape Perception? *Trends Cognit. Sci.* 22, 764–779. <https://doi.org/10.1016/j.tics.2018.06.002>.
2. Funamizu, A. (2021). Integration of sensory evidence and reward expectation in mouse perceptual decision-making task with various sensory uncertainties. *iScience* 24, 102826. <https://doi.org/10.1016/j.isci.2021.102826>.
3. Rungratsameetaweemana, N., Itthipuripat, S., Salazar, A., and Serences, J.T. (2018). Expectations do not alter early sensory processing during perceptual decision-making. *J. Neurosci.* 38, 5632–5648. <https://doi.org/10.1523/JNEUROSCI.3638-17.2018>.
4. Summerfield, C., and de Lange, F.P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nat. Rev. Neurosci.* 15, 745–756. <https://doi.org/10.1038/nrn3838>.
5. Abrahamyan, A., Silva, L.L., Dakin, S.C., Carandini, M., and Gardner, J.L. (2016). Adaptable history biases in human perceptual decisions. *Proc. Natl. Acad. Sci. USA* 113, E3548–E3557. <https://doi.org/10.1073/pnas.1518786113>.
6. Akaishi, R., Umeda, K., Nagase, A., and Sakai, K. (2014). Autonomous Mechanism of Internal Choice Estimate Underlies Decision Inertia. *Neuron* 81, 195–206. <https://doi.org/10.1016/j.neuron.2013.10.018>.
7. Bosch, E., Fritsche, M., Ehinger, B.V., and de Lange, F.P. (2020). Opposite effects of choice history and evidence history resolve a paradox of sequential choice bias. *J. Vis.* 20, 9. <https://doi.org/10.1167/jov.20.12.9>.
8. Braun, A., Urai, A.E., and Donner, T.H. (2018). Adaptive history biases result from confidence-weighted accumulation of past choices. *J. Neurosci.* 38, 2418–2429. <https://doi.org/10.1523/JNEUROSCI.2189-17.2017>.
9. Cicchini, G.M., Mikellidou, K., and Burr, D.C. (2018). The functional role of serial dependence. *Proc. Biol. Sci.* 285, 20181722. <https://doi.org/10.1098/rspb.2018.1722>.
10. Fernberger, S.W. (1920). Interdependence of judgments within the series for the method of constant stimuli. *J. Exp. Psychol.* 3, 126–150. <https://doi.org/10.1037/h0065212>.
11. Fischer, J., and Whitney, D. (2014). Serial dependence in visual perception. *Nat. Neurosci.* 17, 738–743. <https://doi.org/10.1038/nn.3689>.
12. Fründ, I., Wichmann, F.A., and Macke, J.H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *J. Vis.* 14, 9. <https://doi.org/10.1167/14.7.9>.
13. Glaze, C.M., Kable, J.W., and Gold, J.I. (2015). Normative evidence accumulation in unpredictable environments. *Elife* 4, e08825–e08827. <https://doi.org/10.7554/eLife.08825>.
14. Gold, J.I., Law, C.T., Connolly, P., and Bennur, S. (2008). The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. *J. Neurophysiol.* 100, 2653–2668. <https://doi.org/10.1152/jn.90629.2008>.
15. Howarth, C.I., and Bulmer, M.G. (1956). Non-Random Sequences in Visual Threshold Experiments. *Q. J. Exp. Psychol.* 8, 163–171. <https://doi.org/10.1080/17470215608416816>.
16. Pape, A.A., and Siegel, M. (2016). Motor cortex activity predicts response alternation during sensorimotor decisions. *Nat. Commun.* 7, 13098–13110. <https://doi.org/10.1038/ncomms13098>.
17. Urai, A.E., Braun, A., and Donner, T.H. (2017). Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nat. Commun.* 8, 14637. <https://doi.org/10.1038/ncomms14637>.
18. Urai, A.E., de Gee, J.W., Tsetsos, K., and Donner, T.H. (2019). Choice History Biases Subsequent Evidence Accumulation. *Elife* 1–34. <https://doi.org/10.1101/251595>.
19. Bertelson, P. (1961). Sequential redundancy and speed in a serial two-choice responding task. *Q. J. Exp. Psychol.* 13, 90–102. <https://doi.org/10.1080/17470216108416478>.
20. Fozard, J.L., Thomas, J.C., and Waugh, N.C. (1976). Effects of age and frequency of stimulus repetitions on two choice reaction time. *J. Gerontol.* 31, 556–563. <https://doi.org/10.1093/geronj/31.5.556>.
21. Moss, S.M., Engel, S., and Faberman, D. (1967). Alternation and repetition reaction times under three schedules of event sequencing. *Psychonomic Sci.* 9, 557–558. <https://doi.org/10.3758/BF03327887>.
22. Yu, A.J., and Cohen, J.D. (2008). Sequential effects: Superstition or rational behavior? In *Advances in Neural Information Processing Systems 21 - Proceedings of the 2008 Conference*, pp. 1873–1880.
23. Bar-Hillel, M., and Wagenaar, W.A. (1991). The perception of randomness. *Adv. Appl.*

- Math. 12, 428–454. [https://doi.org/10.1016/0196-8858\(91\)90029-1](https://doi.org/10.1016/0196-8858(91)90029-1).
24. Falk, R., and Konold, C. (1997). Making Sense of Randomness: Implicit Encoding as a Basis for Judgment. *Psychol. Rev.* 104, 301–318. <https://doi.org/10.1037/0033-295X.104.2.301>.
 25. Nickerson, R.S. (2002). The production and perception of randomness. *Psychol. Rev.* 109, 330–357. <https://doi.org/10.1037/0033-295X.109.2.330>.
 26. Garrido, M.I., Sahani, M., and Dolan, R.J. (2013). Outlier Responses Reflect Sensitivity to Statistical Structure in the Human Brain. *PLoS Comput. Biol.* 9, e1002999. <https://doi.org/10.1371/journal.pcbi.1002999>.
 27. Maheu, M., Dehaene, S., and Meyniel, F. (2019). Brain signatures of a multiscale process of sequence learning in humans. *Elife* 8, e41541. <https://doi.org/10.7554/elife.41541>.
 28. Squires, K.C., Wickens, C., Squires, N.K., and Donchin, E. (1976). The Effect of Stimulus Sequence on the Waveform of Cortical Event-Related Potential. *Science* 193, 1142–1146.
 29. Wacongne, C., Changeux, J.-P., and Dehaene, S. (2012). A Neuronal Model of Predictive Coding Accounting for the Mismatch Negativity. *J. Neurosci.* 32, 3665–3678. <https://doi.org/10.1523/jneurosci.5003-11.2012>.
 30. Meyniel, F., Maheu, M., and Dehaene, S. (2016). Human Inferences about Sequences: A Minimal Transition Probability Model. *PLoS Comput. Biol.* 12, 1005260. <https://doi.org/10.1371/journal.pcbi.1005260>.
 31. Körding, K., and Wolpert, D.M. (2004). Bayesian integration in sensorimotor learning. *Nature* 427, 1–4.
 32. Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438. <https://doi.org/10.1038/nn1790>.
 33. Pouget, A., Beck, J.M., Ma, W.J., and Latham, P.E. (2013). Probabilistic brains: Knowns and unknowns. *Nat. Neurosci.* 16, 1170–1178. <https://doi.org/10.1038/nn.3495>.
 34. Krishnamurthy, K., Nassar, M.R., Sarode, S., and Gold, J.I. (2017). Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nat. Human Behav.* 1, 0107. <https://doi.org/10.1038/s41562-017-0107>.
 35. Schriver, B.J., Bagdasarov, S., and Wang, Q. (2018). Pupil-linked arousal modulates behavior in rats performing a whisker deflection direction discrimination task. *J. Neurophysiol.* 120, 1655–1670. <https://doi.org/10.1152/jn.00290.2018>.
 36. Weilhammer, V., Chikermane, M., and Sterzer, P. (2021). Bistable perception alternates between internal and external modes of sensory processing. *iScience* 24, 102234. <https://doi.org/10.1016/j.isci.2021.102234>.
 37. Samaha, J., Switzky, M., and Postle, B.R. (2019). Confidence boosts serial dependence in orientation estimation. *J. Vis.* 19, 25. <https://doi.org/10.1167/19.4.25>.
 38. Dehaene, S., Meyniel, F., Wacongne, C., Wang, L., and Pallier, C. (2015). The Neural Representation of Sequences: From Transition Probabilities to Algebraic Patterns and Linguistic Trees. *Neuron* 88, 2–19. <https://doi.org/10.1016/j.neuron.2015.09.019>.
 39. Klaus, M.P., Schöne, C.G., Hartmann, M., Merfeld, D.M., Schubert, M.C., and Mast, F.W. (2020). Roll tilt self-motion direction discrimination training: First evidence for perceptual learning. *Atten. Percept. Psychophys.* 82, 1987–1999. <https://doi.org/10.3758/s13414-019-01967-2>.
 40. Knoblauch, K., and Maloney, L.T. (2012). *Modeling Psychophysical Data in R* (Springer).
 41. Lueckmann, J.M., Macke, J.H., and Nienborg, H. (2018). Can serial dependencies in choices and neural activity explain choice probabilities? *J. Neurosci.* 38, 3495–3506. <https://doi.org/10.1523/JNEUROSCI.2225-17.2018>.
 42. Meyniel, F., Sigman, M., and Mainen, Z.F. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron* 88, 78–92. <https://doi.org/10.1016/j.neuron.2015.09.039>.
 43. Pouget, A., Drugowitsch, J., and Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nat. Neurosci.* 19, 366–374. <https://doi.org/10.1038/nn.4240>.
 44. Fleming, S.M., and Daw, N.D. (2017). Self-evaluation of decision-making: A general bayesian framework for metacognitive computation. *Psychol. Rev.* 124, 91–114. <https://doi.org/10.1037/rev0000045>.
 45. Yeung, N., Botvinick, M.M., and Cohen, J.D. (2004). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychol. Rev.* 111, 931–959. <https://doi.org/10.1037/0033-295X.111.4.931>.
 46. Yeung, N., and Summerfield, C. (2012). Metacognition in human decision-making: Confidence and error monitoring. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 367, 1310–1321. <https://doi.org/10.1098/rstb.2011.0416>.
 47. Ceylan, G., Herzog, M.H., and Pascucci, D. (2021). Serial dependence does not originate from low-level visual processing. *Cognition* 212, 104709. <https://doi.org/10.1016/j.cognition.2021.104709>.
 48. Zhang, H., and Alais, D. (2020). Individual difference in serial dependence results from opposite influences of perceptual choices and motor responses. *J. Vis.* 20, 2. <https://doi.org/10.1167/jov.20.8.2>.
 49. Bliss, D.P., Sun, J.J., and D’Esposito, M. (2017). Serial dependence is absent at the time of perception but increases in visual working memory. *Sci. Rep.* 7, 14739. <https://doi.org/10.1038/s41598-017-15199-7>.
 50. Kim, S., Burr, D., Cicchini, G.M., and Alais, D. (2020). Serial dependence in perception requires conscious awareness. *Curr. Biol.* 30, R257–R258. <https://doi.org/10.1016/j.cub.2020.02.008>.
 51. Berlemont, K., and Nadal, J.P. (2019). Perceptual decision-making: Biases in post-error reaction times explained by attractor network dynamics. *J. Neurosci.* 39, 833–853. <https://doi.org/10.1523/JNEUROSCI.1015-18.2018>.
 52. Fritsche, M., Mostert, P., and de Lange, F.P. (2017). Opposite Effects of Recent History on Perception and Decision. *Curr. Biol.* 27, 590–595. <https://doi.org/10.1016/j.cub.2017.01.006>.
 53. St. John-Saaltink, E., Kok, P., Lau, H.C., and De Lange, F.P. (2016). Serial dependence in perceptual decisions is reflected in activity patterns in primary visual cortex. *J. Neurosci.* 36, 6186–6192. <https://doi.org/10.1523/JNEUROSCI.4390-15.2016>.
 54. Suárez-Pinilla, M., Seth, A.K., and Roseboom, W. (2018). Serial dependence in the perception of visual variance. *J. Vis.* 18, 1–24. <https://doi.org/10.1167/18.7.4>.
 55. Lak, A., Hueske, E., Hirokawa, J., Masset, P., Ott, T., Urai, A.E., Donner, T.H., Carandini, M., Tonegawa, S., Uchida, N., and Kepecs, A. (2020). Reinforcement biases subsequent perceptual decisions when confidence is low: A widespread behavioral phenomenon. *Elife* 9, e49834. <https://doi.org/10.7554/eLife.49834>.
 56. Heilbron, M., and Meyniel, F. (2019). Confidence resets reveal hierarchical adaptive learning in humans. *PLoS Comput. Biol.* 15, e1006972. <https://doi.org/10.1371/journal.pcbi.1006972>.
 57. Norton, E.H., Acerbi, L., Ma, W.J., and Landy, M.S. (2019). Human online adaptation to changes in prior probability. *PLoS Comput. Biol.* 15, 1–26. <https://doi.org/10.1101/483842>.
 58. Christensen, R., Wesley, J., Branscum, A., and Hanson, T.E. (2011). Bayesian Ideas and Data Analysis—An Introduction for Scientists and Statisticians. https://doi.org/10.1111/j.1467-985X.2011.00725_2.x.
 59. Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., and Rubin, D. (2013). *Bayesian Data Analysis Third edition* (CRC Press).
 60. Audacity Team (2019). Audacity: Free Audio Editor and Recorder.
 61. Peirce, J., Gray, J.R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., and Lindeløv, J.K. (2019). PsychoPy2: Experiments in behavior made easy. *Behav. Res. Methods* 51, 195–203. <https://doi.org/10.3758/s13428-018-01193-y>.
 62. Python Core Team (2019). Python: A Dynamic, Open Source Programming Language.
 63. R Core Team (2023). R: A Language and Environment for Statistical Computing.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Raw and analyzed behavioral data	This paper	https://doi.org/10.5281/zenodo.8007314
Original vestibular data	Klaus et al. ³⁹	https://osf.io/dhtq8/
Original visual data	Braun et al. ⁸	https://doi.org/10.6084/m9.figshare.5726845.v1
Software and algorithms		
R 4.2.3	R Core Team	https://cran.r-project.org/
RStudio	RStudio Team	https://posit.co/
Python 3.7.3	Python Core Team	https://python.org/
PsychoPy	The PsychoPy Team	https://psychopy.org
Audacity	Audacity Team	https://audacityteam.org

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact Daniel Schlunegger (daniel.schlunegger@unibe.ch).

Materials availability

Besides data and R codes, this study did not generate any new reagents or materials.

Data and code availability

- All data have been deposited at GitHub and are publicly available. The DOI is listed in the [key resources table](#). In addition, this paper analyzes existing, publicly available data. Accession number and the DOI for the datasets are listed in the [key resources table](#).
- All original code has been deposited at GitHub and is publicly available. The DOI is listed in the [key resources table](#).
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Participants of experiment 1

Eighty participants (46 women, 34 men, $M_{\text{age}} = 27.4$ years, $SD_{\text{age}} = 10.1$ years, age range: 18 - 58 years) took part in the experiment. All participants reported normal hearing and none of them reported any history of dysfunctions or ear infections. They were naive to the purpose of the study. The study was approved by the local ethics committee and participants provided written informed consent prior to participating. Five participants were excluded from the analysis due to incomplete data sets.

Participants of experiment 2

In the study from Klaus et al.³⁹ 30 participants (18 women, 12 men, (age range: 21–38 years) took part in the experiment. In the present study, data from 20 participants were analyzed.

Participants of experiment 3

In Experiment 2 from Braun et al.,⁸ 26 participants (15 women, 11 men, $M_{\text{age}} = 26$ years, age range: 20 - 36 years) took part in the experiment. In both the original and present study, data from 22 participants were analyzed.

METHOD DETAILS

Observer models

We used learning models (i.e. observer models) similar to those used by Meyniel et al.³⁰ and Maheu et al.²⁷ These models, known as fixed belief models, assume that there is no change point in the transition probabilities of the generative process. Instead, the non-stationarity of the transition matrix is implemented through leaky integration, which weights observations according to their recency. Binary responses are labeled as either A=0 or B=1. Here, we provide details for the TP model. The IF and AF models follow the same logic but are simpler in nature. To estimate transition probabilities, we track two statistics simultaneously and apply one of them depending on the preceding response (for example, $\theta^{B|A} = 1 - \theta^{A|A}$). This is done by counting the transition types and updating the model after each observation using Bayes' rule:

$$p(\theta^{B|A}|u_{1:k}) = \frac{p(u_{1:k}|\theta^{B|A}) \cdot p(\theta^{B|A})}{p(u_{1:k}|\theta^{B|A}) \cdot p(\theta^{B|A}) + p(u_{1:k}|\theta^{A|A}) \cdot p(\theta^{A|A})}$$

where u represents the binary sequence and k the k -th trial in the sequence.

For simplicity, we assume that the first response is random. This allows us to compute analytical posterior values. In the case of binary responses and the conjugate prior, the posterior distribution is proportional to a beta distribution^{58,59}:

$$p(\theta^{B|A}|u_{1:k}) \propto \text{Beta}(N_k^{B|A} + 1, N_k^{A|A} + 1)$$

The mean of the distribution corresponds to the predictive likelihood that the next response will be B given that the previous response was A^{58,59}:

$$E\left(\theta_k^{B|A} = \frac{N_k^{B|A} + 1}{N_k^{B|A} + N_k^{A|A} + 2} = p(u_{k+1} = (B|A)|u_{1:k})\right)$$

To enable the model to be flexible and adapt to the recent history of events, we implemented a leak factor. Observations are weighted according to their recency. The leak factor ω is the only free parameter in this model and determines the strength of decay and thus integration over time. A smaller value for ω corresponds to a more local integration, which means that the more recent responses are weighted more strongly in the inference process.

$$N_{\omega}^{B|A} = \sum_{k=1}^n u_{n-k} \cdot \exp\left(\frac{-k}{\omega}\right)$$

Note that this inference must also be performed for $\theta^{B|B} = 1 - \theta^{A|B}$. The final step in creating the predictor for the analysis is to extract the value that applies to the current response, which depends exclusively on the previous response.

Method details of experiment 1

Apparatus and stimuli

Stimuli consisted of two syllables: /ka/ (labelled as A) and /to/ (labelled as B). The syllables were normalized and adjusted to 300 milliseconds length using Audacity (2.3.2).⁶⁰ For each syllable, six stimuli were created by adding white noise, resulting in six different signal-to-noise ratios (SNR): -18, -16, -12, 0, 10, 14 dB which are hereafter labeled as stimulus intensities: 1, 2, 3, 4, 5, and 6. After re-normalizing all stimuli, linear fade-in, and fade-out of 30 milliseconds was applied. Prior to the main experiment, stimuli were tested at different SNRs to find the appropriate range of psychophysical performance, with the intention that stimulus intensities 4, 5, and 6 would result in perfect or near-perfect performance. Stimuli were presented in stereo sound via over-ear headphones with a sampling rate of 44100 Hz via Psychopy (3.1.2)⁶¹ running on Python (3.7.3).⁶²

Experimental procedure

Participants performed a one-interval 2AFC auditory discrimination task (Figure 1A) on a standard laptop. The experiment was set up as a standard auditory discrimination task and participants were asked to respond as accurately as possible. In each trial, participants were to indicate their response for stimulus A via button "F" and their response for stimulus B via button "J" on the keyboard. Upon stimulus onset,

a fixation cross appeared in the middle of the screen. The fixation cross and the auditory stimulus were presented for 300 milliseconds. The next trial started 550 milliseconds after the previous response. Importantly, participants received no feedback about the correctness of their choice. The experiment was divided into four blocks of 300 stimuli each; each block corresponded to one of four different experimental conditions. As suggested in Maheu et al.²⁷, conditions consisted of orthogonal variations in the two-dimensional space of transition probabilities: 1) fully stochastic (i.e., the same number of A's and B's and the same number of repetitions and alternations), 2) frequency-biased (i.e., one of the stimuli was presented more frequently, 3/4 vs. 1/4), 3) repetition-biased (i.e., 3/4 repetitions and 1/4 alternations), and 4) alternation-biased (i.e., 3/4 alternations and 1/4 repetitions). We validated each sequence by fitting a Markov chain to the first 50 stimuli and to the remaining 250 stimuli of each sequence, ensuring that the sequences truly represented the transition probabilities they were drawn from. Our goal is to understand how the brain integrates sequence knowledge when sensory information is ambiguous. Therefore, we took the following measures: During the first 10 trials, only stimulus intensities from 4 to 6 occurred; during trials 11 to 50, only stimulus intensities from 3 to 6 occurred; during trials 51 to 100, only stimulus intensities from 2 to 6 occurred; and during trials 101 to 300, all stimulus intensities occurred. In addition, stimulus intensities 1 and 2 were not allowed to occur consecutively. Because participants were not informed of the differences in the underlying generative transition probabilities, these measures were taken to implicitly provide participants with information about this experimental manipulation and to ensure that participants generated response sequences that actually reflected the underlying bias. For each condition, we created five sequences that were counterbalanced across participants; the order of conditions was fully randomized. Each participant completed 1200 trials, resulting in 90000 trials in total for the final analysis.

Method details of experiment 2

Complete details about Experiment 2 are available in the original article from Klaus et al.³⁹ In short, 30 participants participated in a self-motion discrimination experiment. The training group ($n=10$) received a self-motion discrimination training, as well as pre- and post-test measurements to determine their self-motion perception thresholds. The control group ($n=20$) received no training but participated in the same pre- and post-test measurements. All participants were screened for vestibular disorders. The motion stimuli were applied using a 6-degree-of-freedom motion platform and consisted of single cycles of sinusoidal acceleration motion with a frequency of 0.2 Hz or 1 Hz. During the training, the participants in the training group experienced roll tilts about an earth horizontal axis with a frequency of 0.2 Hz. The peak velocity of the stimuli was determined individually for each participant based on their performance in the pre-test measurement. The pre- and post-test measurements for the experimental group also included roll motion at 1 Hz, pitch motion at 0.2 Hz, and y-translation motion at 0.2 Hz. The control group completed the pre- and post-test measurements in the roll 0.2 Hz and roll 1 Hz conditions, with half of the group also completing pitch and y-translation measurements at 0.2 Hz and the other half completing pitch measurements at 1 Hz. The post-test took place on the ninth day after the pre-test session and the participants received 24 practice trials before the measurement started. The motion stimuli were presented to the participants in random order and the participants were asked to indicate the direction of the motion. The threshold for each motion condition was determined using a 2AFC task and a maximum likelihood estimate.

We had to choose a subset of the original data set because not all participants completed all experimental conditions. For reasons of comparability, we chose the subset of the data in such a way, that the overlap of completed conditions and the number of trials to analyze is maximized. This resulted in a data set of 20 participants and a total of 22280 trials, all from the pre- and post and not from the training conditions. Conditions included roll motion at 1 Hz, roll motion at 0.2 Hz, pitch motion at 0.2 Hz, and y-translation motion at 0.2 Hz.

Method details of experiment 3

Complete details about Experiment 3 are available in the original article from Braun et al.⁸ In short, in Experiment 2 of Braun et al.,⁸ 26 participants were presented with a random-dot motion discrimination task in blocks of 60 trials. Four participants were excluded from the analysis. Data from 22 participants was analyzed, where each participant completed six sessions. Experimental conditions were similar to the ones used in the auditory data set: Neutral (i.e., fully stochastic), repetitive (repetition-biased), alternating (alternation-biased), with a slightly stronger bias in the biased conditions (i.e., 0.8 vs. 0.75). Crucially

and in contrast to our experiments, participants were informed that generative transition probabilities changed from session to session but stayed constant within each session.

QUANTIFICATION AND STATISTICAL ANALYSIS

Observer models

Data analysis was implemented in R 4.2.3.⁶³ In order to create the predictors (predictive likelihood of IF, AF, tdeand TP) for our analyses, we ran the observer models on all trials. In the case of IF, the observer starts counting with a uniform prior on θ^B . In the case of AF, the observer starts counting with a uniform prior on $\theta^{\text{repetition}}$, and in the case of TP, the observer starts counting from zero with a uniform prior on $\theta^{B|A}$ and $\theta^{B|B}$ for each participant and condition separately (i.e. sequences of 300 responses/stimuli in the auditory data set, sequences of 140 responses/stimuli in the vestibular data set, and sequences of 60 responses/stimuli in the visual data set). Note that in the case of AF and TP, the predictor has to be transformed before it enters the analyses. The predictor has to be transformed so that response = 1 is predicted on each trial.

Log-likelihood analysis

We assessed the quality of fit of the probabilistic models using a log-likelihood approach:

$$\log p(\text{response}|\text{model}) = \sum_{k=1}^n (\text{response}_k \cdot \log p(\text{response}_k|\text{model})) + (1 - \text{response}_k) \cdot \log(1 - p(\text{response}_k|\text{model}))$$

where response_k is the binary choice (either 0 or 1) for the k -th trial in a sequence of length n , and $p(\text{response}_k|\text{model})$ is the inferred probability of response k given the predictions of the learning model. The first term in the summation corresponds to the log-likelihood when $\text{response}_k = 1$, and the second term corresponds to the log-likelihood when $\text{response}_k = 0$. To implement the stimulus-only model using the log-likelihood approach, we replaced the binary stimulus identities (coded as 0 and 1) with extremely low and high probabilities (i.e., 0.001 and 0.999, respectively).

Easy vs. difficult trials

The log-likelihood analysis was performed for easy and difficult trials separately. We choose the cutoff between easy and difficult trials at the stimulus intensity level. This means that in terms of number of trials, all participants contributed equally to the easy and difficult data subsets. To determine which stimulus intensities were to classify as easy or difficult trials, we first computed the accuracy per participant, stimulus, and stimulus intensity. In each modality, a cutoff for easy vs. difficult trials was chosen at 0.75, which refers to the median accuracy per stimulus and stimulus intensity across participants (Figure S1).

Generalized Linear Models

We used Generalized Linear Models (GLMs) with probit link function to complement the log-likelihood analysis and report the results in the supplemental information (Figures S3 and S4). We predict participants' binary responses in all trials (except the first trial at the beginning of a block). For all modalities and learning models, this was done using the predictive likelihood that was based on the responses (i.e., response-based models).

$$P(\text{response} = 1) = \Phi[\alpha + \beta(\text{stimulus}) + \sum_{i=1}^{n-1} \beta_{i+1}(\text{intensity}_i) + \beta_{n+1}(\text{model}) + \sum_{i=1}^{n-1} \beta_{i+n+1}(\text{intensity}_i \times \text{model})]$$

where n represents the number of intensity levels (levels of the ordered factor). Including an ordered factor in the analysis will fit polynomial functions up to the $n-1$ -th order, which means linear, quadratic, cubic, 4th order polynomials for example, depending on the number of levels. In this equation, model refers to either one of the three learning models or additionally to a model that includes a simple preference for repetition. The probability of repeating the preceding response was set to 0.6. Like AF and TP, the predictor needs to be transformed so that response = 1 is predicted on each trial.

Estimating psychometric curves

To estimate psychometric curves for the auditory and visual data set, we employed Generalized Linear Models with logit link function. They were applied to each participant and experimental condition separately.

$$\log\left(\frac{p(N^B)}{p(N^A)}\right) = \alpha + \beta_i(\text{intensity})$$

where N^B and N^A represent the number of times response B and A were pressed, respectively. Note that N^A is implicitly the complement of N^B (i.e., $N^A = 1 - N^B$). We separated the trials from all experimental conditions into four subsets based on the two preceding choices from the two preceding trials. Psychometric curves were fitted separately to the observed proportion of response B in each of these four subsets.

Statistical tests

We used nonparametric two-sided Wilcoxon signed rank tests for all statistical tests against zero or pairwise comparisons. This choice was motivated by the fact that, in many instances, the Shapiro- Wilk normality test yielded statistically significant results ($p < 0.05$), indicating that the assumption of normality was not met.