

# Sustained EEG responses to rapidly unfolding stochastic sounds reflect Bayesian inferred reliability tracking

Sijia Zhao<sup>a,b,\*</sup>, Benjamin Skerritt-Davis<sup>c,1</sup>, Mounya Elhilali<sup>d</sup>, Frederic Dick<sup>e</sup>, Maria Chait<sup>a,\*\*</sup>

<sup>a</sup> Ear Institute, University College London, London WC1X 8EE, UK

<sup>b</sup> Department of Experimental Psychology, University of Oxford, Oxford OX2 6GG, UK

<sup>c</sup> Applied Physics Laboratory, Johns Hopkins University, Laurel, MD 20723, United States

<sup>d</sup> Electrical & Computer Engineering, Johns Hopkins University, Baltimore, MD 21218, United States

<sup>e</sup> Department of Experimental Psychology, University College London, London WC1H 0DS, UK

## ARTICLE INFO

### Keywords:

Precision  
Statistical inference  
Auditory perception  
EEG  
Bayesian brain

## ABSTRACT

How does the brain track and process rapidly changing sensory information? Current computational accounts suggest that our sensations and decisions arise from the intricate interplay between bottom-up sensory signals and constantly changing expectations regarding the statistics of the surrounding world. A significant focus of recent research is determining which statistical properties are tracked by the brain as it monitors the rapid progression of sensory information. Here, by combining EEG (three experiments  $N \geq 22$  each) and computational modelling, we examined how the brain processes rapid and stochastic sound sequences that simulate key aspects of dynamic sensory environments. Passively listening participants were exposed to structured tone-pip arrangements that contained transitions between a range of stochastic patterns. Predictions were guided by a Bayesian predictive inference model. We demonstrate that listeners automatically track the statistics of unfolding sounds, even when these are irrelevant to behaviour. Transitions between sequence patterns drove a shift in the sustained EEG response. This was observed to a range of distributional statistics, and even in situations where behavioural detection of these transitions was at floor. These observations suggest that the modulation of the EEG sustained response reflects a process of belief updating within the brain. By establishing a connection between the outputs of the computational model and the observed brain responses, we demonstrate that the dynamics of these transition-related responses align with the tracking of “precision” – the confidence or reliability assigned to a predicted sensory signal – shedding light on the intricate interplay between the brain’s statistical tracking mechanisms and its response dynamics.

### Significant Statement

The human brain excels at processing rapidly changing sensory information, yet the specific statistical computations it performs remain a key area of enquiry. Current hypotheses suggest it monitors “precision”, the reliability of predictions about future inputs. However, evidence, mainly from decision-making tasks, is indirect. It’s unclear if the brain automatically tracks precision or if observed effects stem from input relevance. We introduce a new paradigm where participants are passively exposed to specifically structured, rapid sound patterns. By relating EEG data to a Bayesian ideal-observer model, we reveal that the dynamics of

EEG responses are closely associated with dynamically modelled Bayesian precision. This sheds light on how neural systems compute and monitor statistics across contexts, offering insights into human and non-human cognition.

## 1. Introduction

The brain is increasingly understood as a complex system that learns and adapts by constantly processing and interpreting sensory information (Soltani and Izquierdo, 2019). Our perceptions and actions emerge from the dynamic interplay between incoming sensory signals and

\* Correspondence to: Department of Experimental Psychology, New Radcliffe House, Walton Street, Oxford OX2 6GG, UK.

\*\* Correspondence to: Ear Institute, University College London, 332 Gray’s Inn Road, London WC1X 8EE, UK.

E-mail addresses: [sijia.zhao@psy.ox.ac.uk](mailto:sijia.zhao@psy.ox.ac.uk) (S. Zhao), [m.chait@ucl.ac.uk](mailto:m.chait@ucl.ac.uk) (M. Chait).

<sup>1</sup> Equal contribution.

continuously updated expectations (“model beliefs”) about our surrounding environment. A major current research effort in systems neuroscience is to unravel the mechanisms behind this “observe-predict-update” cycle. Fundamental to this goal is understanding how the statistics of the environment are being monitored and represented by perceptual systems.

Based on Bayesian ideal observer models, theoretical accounts of these inference processes suggest that in noisy and volatile environments, optimal belief updating depends on the observer’s ability to assess both the stability and the reliability of the environment (Nassar et al., 2010, 2012; Skerritt-Davis and Elhilali, 2021a; Wilson et al., 2013). Specifically, observers are hypothesized to track the probability that the environment has changed, i.e., that the generative process underlying the observations is now different (Nassar et al., 2010; Glaze et al., 2015; Boubenec et al., 2017; Skerritt-Davis and Elhilali, 2018). They must also simultaneously update a predictive model of the next observation, based on previously encountered inputs. Evidence suggests that the latter process involves not only the representation and continuous updating of the next expected input, but also tracking of the confidence, or *inferred reliability* (“precision”) of this prediction, such that incoming events are evaluated based on this weighing (Yon and Frith, 2021; Lawson et al., 2021; Heilbron and Chait, 2018; Kanai et al., 2015; Feldman and Friston, 2010; Weissbart et al., 2020; Friston, 2017).<sup>2</sup>

Belief updating has predominantly been studied in research on probabilistic decision making. Using relatively slow tasks that depend on evidence accumulation, multiple studies have shown that participant responses reflect estimated precision (e.g., Nassar et al., 2012; Lawson et al., 2021; Krishnamurthy et al., 2017). For example, in a seminal experiment (Nassar et al., 2012) participants performed a task that required them to predict the next number in a series. The numbers were drawn from a Gaussian distribution with a mean that changed abruptly at random intervals. Results indicated that performance and baseline pupil diameter - hypothesized to index arousal and autonomic state - were modulated by belief uncertainty, as modelled with a Bayesian model of evidence accumulation. More recently, an EEG study (Weissbart et al., 2020) showed that the cortical activity tracks word surprisal in continuous natural speech, and that this tracking is modulated by precision. Here we focus on understanding evidence accumulation in a faster-paced domain, asking what statistics of rapidly unfolding sensory signals are automatically tracked by the brain during listening.

The acoustic environment is rich in statistical regularities that unfold on many different time scales. These statistics characterize auditory textures (water, fire, a roaring engine, the hum of a crowd), reflect physical constraints of sources (locomotion sounds, vocalization), and register the reoccurrence of events across time (the repeating call of a bird in the forest, or motifs in music). The auditory system continuously analyses and tracks these regularities as they appear and disappear from our surroundings - even when this information is not immediately relevant to behaviour (Sohoglu and Chait, 2016; Barczak et al., 2018; Asokan et al., 2021; Demarchi et al., 2019). Experimentally manipulating the statistical structure of auditory sequences therefore offers valuable insights into how the brain accumulates and processes distributional information, leading to adaptive changes in expectations. In particular, it offers a means to probe potential heuristics used by the system, to ascertain which statistical information is tracked, and to disentangle more ‘automatic’ processes from those related to carrying out behavioural goals.

<sup>2</sup> It is somewhat ironic that the term “precision” has varying interpretations across fields. In our case, we adhere to its original statistical definition as “the inverse variance of a distribution,” specifically referring to the width of the internal distribution(s) estimated by the model and human listeners. We think using “precision” in this way is important due to the importance of this concept in sensory models (as laid out by Yon and Frith (Yon and Frith, 2021)).

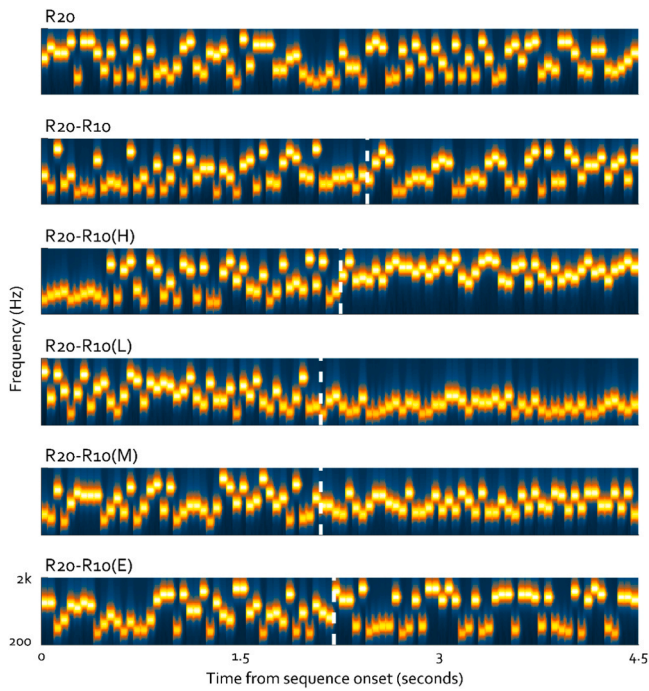
In this vein, observations from paradigms based on introducing deviants in structured sound sequences have shown that deviant-evoked brain responses reflect precision-weighted prediction errors (Heilbron and Chait, 2018; Garrido et al., 2011; Hsu and Hämäläinen, 2021; Sedley et al.; Hsu et al., 2019; SanMiguel et al., 2021). Recent dynamic causal modelling (DCM) work has suggested that encoding of precision and prediction errors may be dissociated (Lecaiguard et al., 2022).

Nevertheless, deviant-evoked responses provide only an indirect and momentary measure of listeners’ sensitivity to ongoing statistical structure. To circumvent this issue, Barascud et al. (2016) used the dynamics of brain responses to statistically structured auditory streams to investigate the process by which listeners automatically acquire an internal model of regularities in the environment. Here, naïve listeners were exposed to sequences of short tone-pips containing transitions from randomly ordered to regularly repeating patterns. The emergence of regularity was associated with a prototypical brain response pattern: a gradual increase followed by a plateau in sustained EEG power, underpinned by auditory cortical, frontal, and hippocampal sources. Increases in sustained EEG power toggled to increases in sequence predictability have been observed for a variety of patterned stimuli (Sohoglu and Chait, 2016; Southwell et al., 2017; Southwell and Chait, 2018; Herrmann et al., 2021; Herrmann and Johnsrude, 2018). The greater EEG amplitude for regular over random stimuli cannot easily be interpreted as a response to physical attributes of the signal. Adaptation, for example, would be expected to result in the opposite pattern (Megela and Teyler, 1979; Pérez-González and Malmierca, 2014). Rather, Barascud et al. (2016) showed that the sustained response appears to vary consistently with the predictability of the ongoing stimulus sequence and hypothesized that it might reflect precision estimated from the unfolding trial.

However, this previous work suffers from a major limitation: these studies only used transitions from random to fully deterministic sequences, i.e., precisely repeating patterns. This makes it difficult to disambiguate the predictions made under a Bayesian framework - emphasizing the process of evidence accumulation under uncertainty, and distinguishing “change probability” tracking and “precision tracking” - from those from memory-based accounts, which emphasize mechanisms engaged by the repeating tone sequences (Skerritt-Davis and Elhilali, 2018; Harrison et al., 2021; Andriillon et al., 2015).

Informed by an initial behavioural study that established listeners’ sensitivity to changes in rapid, stochastic sound sequences, we combined EEG and computational modelling to investigate brain tracking of these sequences. Passively listening EEG participants were exposed to structured tone-pip arrangements (Fig. 1) that contained transitions between a range of stochastic patterns, manifested as abrupt changes in the mean and/or variance of the unfolding sequence. Importantly, the sequences did not contain ‘transients’ - (new tones that had not been encountered earlier in the trial); changes in the underlying distribution could only be inferred by tracking the unfolding statistical properties of the sequence. This thus allows us to ask directed questions concerning what statistical information about the sequence structure is reflected in EEG response dynamics.

Predictions were guided by a Bayesian predictive inference model (Skerritt-Davis and Elhilali, 2021b; Fig. 2) previously used to study a variety of statistical tracking phenomena (Nassar et al., 2010; Soltani and Izquierdo, 2019). To explain the response of human observers in volatile environments, the model monitors sequences containing hidden changes in the underlying statistical structure. As it “listens” to the unfolding auditory sequence, it derives a predictive distribution for the next observation using sufficient statistics collected over a local context of previous observations, the extent of which adapts to the volatility of the observed sequence. The model monitors two key inferred statistics to guide predictions about upcoming input: (1) *change probability*, the likelihood that a change in sequence statistics has occurred; and (2) *precision* (inferred reliability), the inverse variance of predictive distributions derived from the relevant past inputs. As highlighted above,



**Fig. 1.** Example spectrograms of stimuli in Experiment 1. The stimuli were sequences of concatenated tone-pips (50 ms) with frequencies drawn *with replacement* from a pool of 20 fixed values, spaced from 222 Hz to 2000 Hz in steps of  $\sim 2$  semitones. Six different conditions were used. The control condition, R20, was generated by randomly sampling from the full pool with replacement; The remaining conditions contained a change partway through the trial, manifested as a reduction in alphabet size from 20 to 10 tones. In R20-R10 the retained tones were sampled equally from the full pool. In R20-R10(H) the retained tones were sampled from the two highest frequency bands (the 10 highest frequencies; 707 Hz – 2000 Hz). In R20-R10(L) the retained tones were sampled from the two lowest frequency bands (10 lowest frequencies; 222 Hz – 630 Hz). In R20-R10(M) the retained tones were sampled from the two central bands (10 middle frequencies; 397 Hz – 1122 Hz). In R20-R10(E) the retained tones were sampled from the highest (1259–2000 Hz) and lowest (222–354 Hz) bands (5 highest and 5 lowest frequencies). For illustrative purposes, the plotted sequences are of a fixed length. In the actual experiment, transitions (marked as dashed white lines here) occurred between 2 and 2.5 seconds post sequence onset, with the overall sequence length varying from 4 to 4.5 seconds.

sensitivity to both change probability and precision is hypothesized to play a fundamental role in perception. Accurate estimation of change probability is critical for survival in natural environments that often change abruptly and unpredictably. Tracking of precision is hypothesized to be critical to observers' ability to operate under different degrees of uncertainty (Yon and Frith, 2021; Behrens et al., 2007; Payzan-LeNestour and Bossaerts, 2011; Gold and Stocker, 2017).

By relating model outputs to observed brain responses, we asked (a) whether EEG responses exhibit sensitivity to the evolving structural patterns within auditory sequences; and (b) whether the dynamics of these transition-related responses align with change point estimation, or alternatively tracking of precision (as hypothesized in Barascud et al., 2016).

## 2. Results

### 2.1. Listeners are sensitive to the statistics of stochastic tone patterns

We developed a new stimulus paradigm, based on the tone patterns used in Barascud et al. (2016), to study sequences that contain transitions between stochastic patterns that are differentially statistically structured (i.e., sequences that contain abrupt changes to the underlying generative parameters). Our basic stimulus is a random pattern of 50 ms

tone-pips whose frequencies are uniformly sampled (with replacement) from a discrete pool of 20 values (R20; an example is shown in Fig. 1; see Methods). We created transitions from the full pool of 20 values to another random sequence, but where the tones are drawn from a reduced set ("alphabet size") of 10 frequencies (R20-R10). To address the core question - whether listeners are sensitive to changes in distributional statistics - we manipulated R20-R10 such that the 10 retained frequencies are sampled from different distributions (Fig. 1).

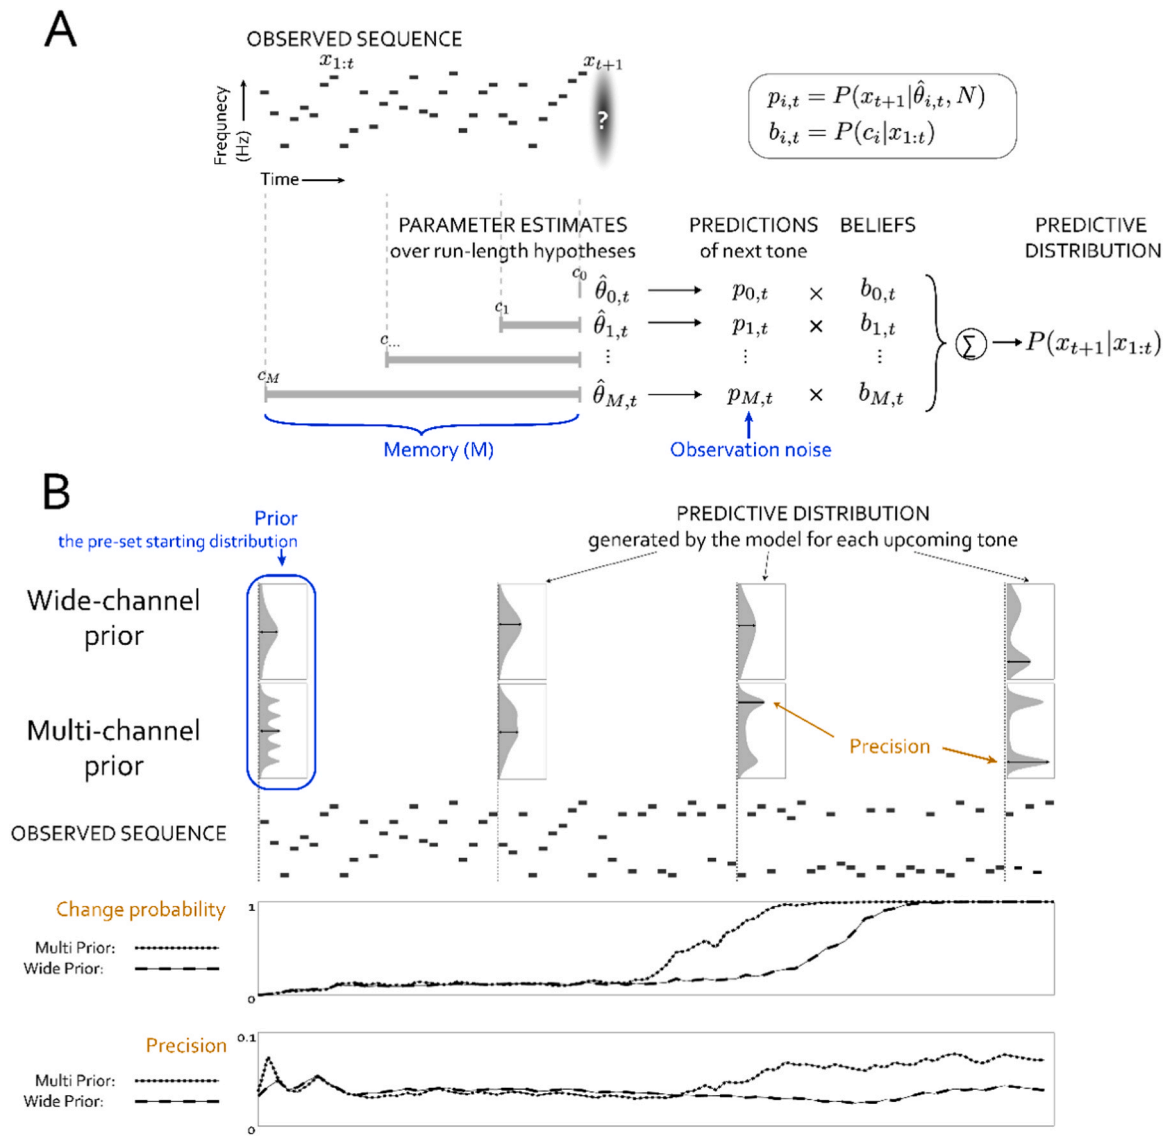
The full pool of 20 frequencies was divided into four logarithmically equally sized frequency bands, each containing five frequency values. All transitions featured a reduction to 10 frequencies. These were selected from (a) the full pool of frequencies (R20-R10); (b) the highest (R20-R10(H)) spectral band; (c) lowest band (R20-R10(L)); (d) middle band (R20-R10(M)); and (e) edge bands (R20-R10(E); see the figure legend of Fig. 1 and Methods for details).

**R20-R10(H)** and **R20-R10(L)** cause a change in both the mean and variance of the spectral distribution, and thus should evoke a transition to a single-peaked spectral predictive distribution. **R20-R10(M)** drives only a change in variance, but again a single-peaked spectral distribution. **R20-R10(E)** will drive only a change in variance, but a more complex, double-peaked spectral distribution. Finally, **R20-R10** is only associated with a reduction in alphabet size, but neither a change in mean nor variance.

As an initial step to characterizing listeners' sensitivity to such changes in spectral statistics over time, we ran a behavioural study where participants were instructed to press a button immediately upon hearing a transition within a sound sequence. To allow participants to develop the most appropriate listening strategy, each stimulus type was presented in a separate block. In each block, 20 trials of the transition condition were randomly intermixed with 20 trials of its non-change control (R20). The order of the five blocks was randomised across participants.

In principle, there are at least two broad strategies which listeners could employ in the process of transition detection: (1) tracking the zeroth order statistics of the unfolding sequence (i.e., the probability of each frequency); (2) monitoring the distributional statistics of tone frequencies (i.e., the probability distribution across the spectrum), either by tracking a *wide-band* (cross spectral) distribution covering the entire spectral range, or by tracking several concurrent (*'multiband'*) spectral distributions reflecting activity in distinct spectral channels. In (1) the transition will manifest as an increase in incidence in those channels (from probability of 1/20 to 1/10) that are retained after the transition, and in a decrease in incidence in the rest of the channels (from a probability of 1/20 to 0). If listeners do employ *Strategy 1*, we expect no difference in performance across conditions because the same number of tones (frequency channels) are retained in each. In contrast, *Strategy 2* would predict graded performance across conditions as a function of the similarity between pre/post transition means and variances. Here, R20-R10(H) and R20-R10(L) both contain changes in the mean and variance and hence should be detected most readily, with somewhat poorer detection in R20-R10(M) (change in variance only). R20-R10(E) transitions should be yet more difficult to detect because they are associated with the emergence of a double peaked post-transition distribution. Finally, R20-R10 should be most difficult because there is no change in variance or mean.

Fig. 3A summarizes the behavioural performance (mean  $d'$  for R20-R10(H) R20-R10(H):  $2.56 \pm 0.27$ , R20-R10(L):  $2.72 \pm 0.36$ , R20-R10(M):  $1.61 \pm 0.25$ , R20-R10(E):  $0.96 \pm 0.17$ , R20-R10:  $0.09 \pm 0.12$ ). A repeated measures ANOVA on  $d'$  data showed a main effect of condition ( $F(4, 36) = 32.99, p < 0.001$ ). Post-hoc tests revealed that the  $d'$  for R20-R10 was significantly lower than all other conditions ( $p \leq 0.001$ ). The second lowest condition was R20-R10(E) followed by R20-R10(M) ( $p = 0.022$ ). The highest  $d'$  values were associated with R20-R10(H) and R20-R10(L), which did not statistically differ from each other ( $p = 0.60$ ) but were significantly higher than for the other transition types ( $p \leq 0.003$ ). One-sample t-tests (2-tailed) for comparisons against



**Fig. 2.** Schematic of perceptual model and model outputs. (A) At time  $t$ , the model contains multiple parameter estimates collected over run-lengths into the past up to the memory constraint ( $M$ ). Each estimate yields a prediction for the next tone, with increased uncertainty due to observation noise. Upon observing  $x_{t+1}$ , the model updates the run-length beliefs using the predictive probability for each hypothesis. Note that the prediction for length  $M$  is used to update all beliefs with length greater than or equal to  $M$ , thus limiting the number of past observations used in the update. (B) Outputs from the model for an example observed sequence (here the unfolding stimulus is R20-R10(E)). The predictive distribution (Gaussian distribution shown on the top), generated by the model for each upcoming tone, combines predictions across run-length hypotheses weighted by their beliefs, thus “integrating out” run-length. At the beginning of each trial, a starting distribution (“Prior”), is initiated; Two prior types were modelled - a single Gaussian across the entire spectral range (wide-channel prior) or a Gaussian with multiple components (multi-channel prior). Change probability is the probability that at least one changepoint has occurred, as inferred using the run-length beliefs. To adapt the model to behavioural responses, the model detects a change if the change probability exceeds threshold. Precision is quantified as the maximum value of the normalised predictive distribution (marked as a black line in each predictive distribution). In this example, the change in stimulus statistics is associated with a dynamic increase in change probability and precision, which occurs earlier in the multi-channel prior model. The model parameters (memory, observation noise and priors) are in blue. The outputs (change probability and precision) are plotted in the bottom for each prior distribution (wide- and multi-channel prior).

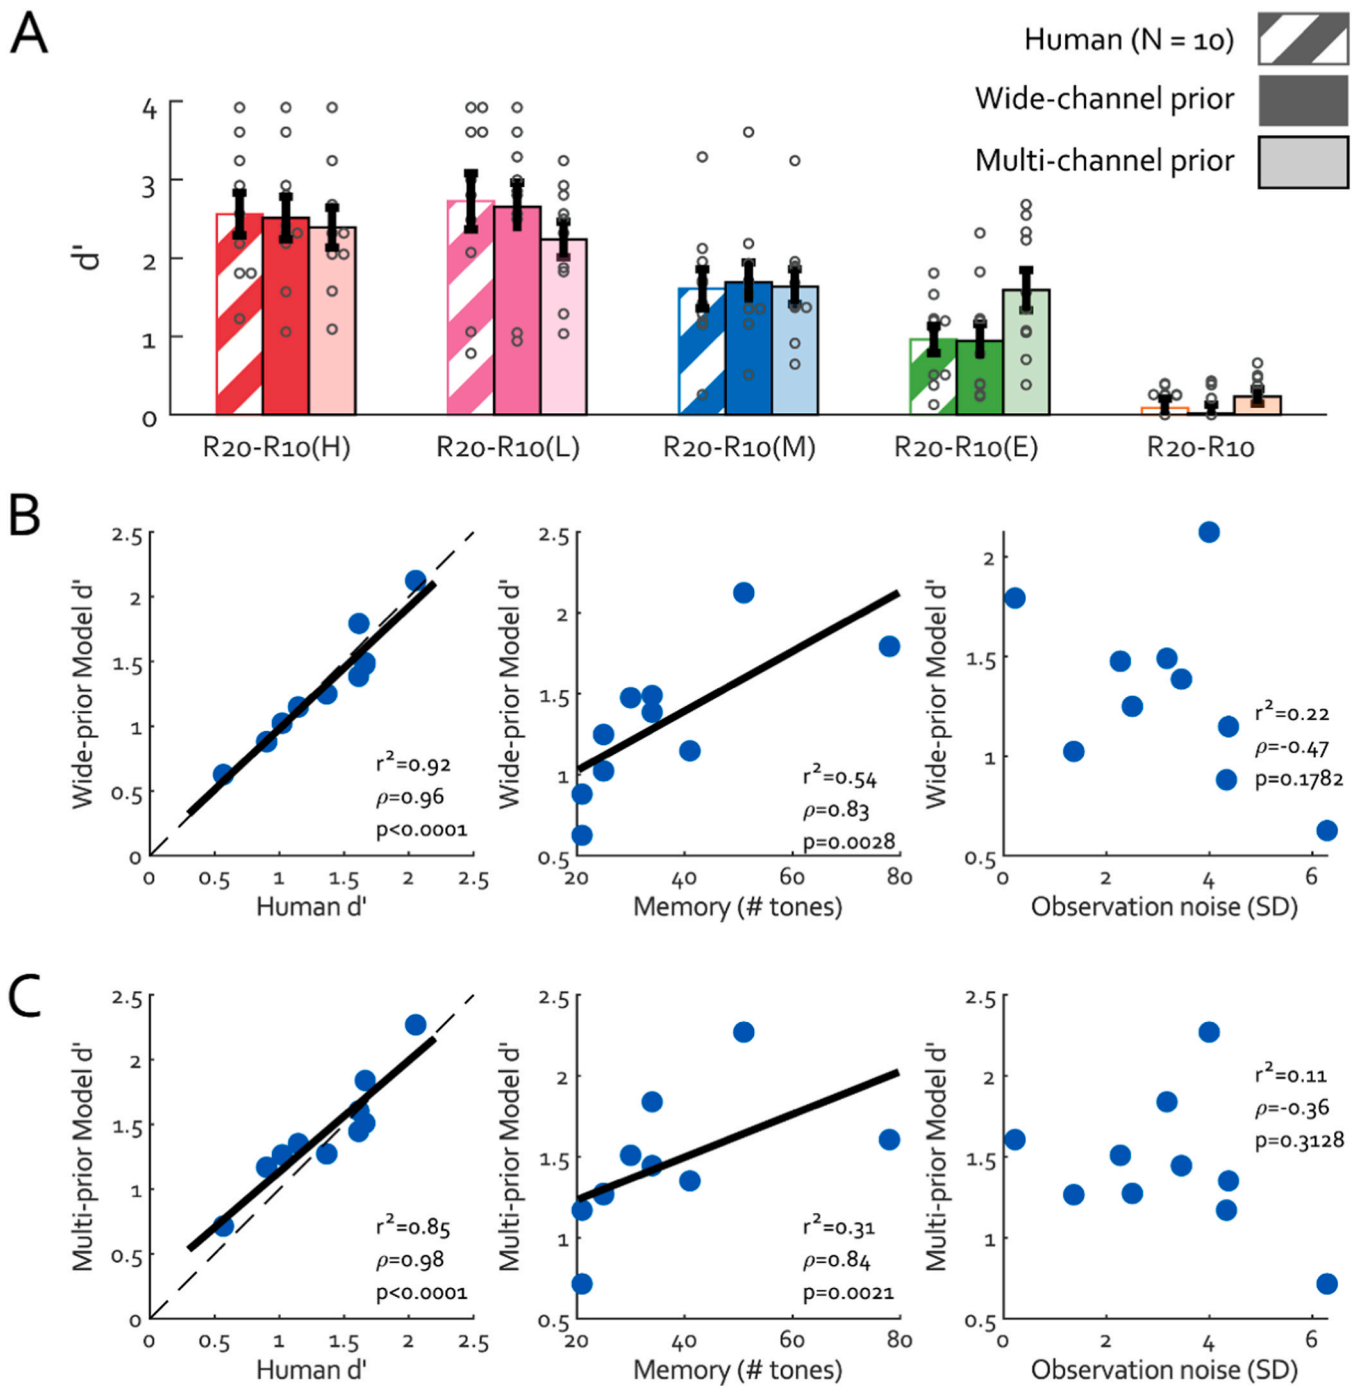
0 were applied to the  $d'$  data. Results established that performance on the H, L, M and E conditions was significantly above floor ( $p < 0.001$ ), whilst that for R20-R10 did not differ from floor ( $t(9) = 0.76, p = 0.47$ ).

This results pattern demonstrates that behavioural performance is consistent with Strategy 2, where listeners track the distributional statistics of tone frequencies. Detection of R20-R10(H) and R20-R10(L) was most accurate, followed by R20-R10(M), 20-R10(E), with R20-R10 being least accurate.

## 2.2. A Bayesian prediction model, D-REX, accurately fits behavioural data

We used a Bayesian sequential prediction model (Dynamic Regularity Extraction D-REX) (Skerritt-Davis and Elhilali, 2021b) to fit the behavioural data (Fig. 2).

As it is exposed to the unfolding tone-pip sequence, D-REX sequentially generates a predictive distribution for the frequency of the next tone given the previous observations (Fig. 2A). This distribution takes the form of a mixture of Gaussians, reflecting the model’s uncertainty about the underlying sound statistics. At the beginning of each trial, the model is initialised with a prior distribution representing the listener’s



**Fig. 3.** Behavioural and modelling results for Experiment 1, with results from humans (striped bars) and the D-REX models (solid bars, dark for wide-channel prior and light for multi-channel prior models). Coloured bars are group mean values, with grey circles indicating individual means. All error bars are 1 SEM. (B) & (C) Overall fit between human and model performance. Correlations between human and modelling results indicate that the D-REX model with wide-channel prior (B) and multi-channel prior (C) provides a good fit for human performance on an individual level. (Left) Human  $d'$  plotted against the model  $d'$ . Blue dots indicate individual participants. The black line indicates the linear fit of the individual data. Spearman's correlation values are shown at the bottom right corner of the graph. Overall  $d'$  was computed by collapsing across all five transition conditions. We then examined the relationship between model performance ( $d'$ ) and two parameters derived from the model fit to each participant: memory (M) and observation noise. (Middle) Like the relationship with the actual human  $d'$ , memory size is positively correlated with the model  $d'$ , while observation noise is not.

initial assumptions about the source distribution. We considered two types of priors (Fig. 2B): a *wide-channel prior* (single Gaussian) and a *multi-channel prior* (multimodal Gaussian with five channels). The former is modelled using a *single* Gaussian centred on the log-frequency range of the stimuli, and the latter's prior distribution uses 5 channels spread evenly in log-frequency (between  $-16$ – $16$  semitones relative to the centre frequency of the stimuli). These priors correspond to two

schemes for the listening strategy #2 as discussed above: (a) listeners employ a *single wide-band* (cross spectral) distribution covering the entire spectral range; or (b) listeners employ a *within-channel* listening strategy, approximating a uniform distribution over frequency where several concurrent, partially overlapping distributions are being tracked simultaneously. To capture human listeners' limited perceptual capacity, the model incorporates two perceptually plausible constraints: a

finite memory, constraining the maximum number of previous observations (i.e., number of tone-pips in the sequence) used to generate the predictive distribution; and observation noise, representing individual differences in pitch perception as well as variations in perceived loudness.

We applied the model to the stimuli in Experiment 1 to fit human listeners' behavioural performance. For each participant, the model was run on the entire experimental session in the same way they were delivered to the human participants. The comparison between human and model performance for each transition condition is shown in Fig. 3A.

A repeated measures ANOVA on model  $d'$  data showed a main effect of condition for both the wide-prior model ( $F(2.82, 25.36) = 33.48$ ,  $p < 0.001$ ,  $\eta^2 = 0.79$ ) and the multi-prior model ( $F(2.29, 20.62) = 23.46$ ,  $p < 0.001$ ,  $\eta^2 = 0.72$ ). *Post-hoc* tests for the wide-prior model revealed that the  $d'$  for R20-R10 was significantly lower than all other conditions ( $p = 0.016$  when compared with R20-R10(E) and  $p < 0.001$  compared with R20-R10(H)/(L)/(M)). Overall, the pattern was similar to that observed in humans: the second lowest  $d'$  was R20-R10(E), followed by R20-R10(M) (mean difference =  $0.75 \pm 0.27$ ,  $p = 0.047$ ). The highest  $d'$  values were associated with R20-R10(H) and R20-R10(L) which did not differ statistically ( $p = 1.00$ ) but were significantly higher than the other transition types (wide-prior:  $p \leq 0.043$ ). One-sample  $t$ -tests (2-tailed) for comparisons against 0 were applied to the  $d'$  data. Again, the wide-prior model was very similar to human performers in that it could not accurately detect the change in R20-R10 (not different from floor,  $t(9) = 0.15$ ,  $p = 0.89$ ). *Post-hoc* tests for the multi-prior model revealed a slightly different pattern: unlike the single-prior model, a one sample  $t$ -test against 0 showed above chance performance for R20-R10 ( $t(9) = 2.52$ ,  $p = 0.033$ ). Furthermore, the multi-prior model showed no difference between R20-R10(M) and R20-R10(E) conditions, outperforming human listeners in the latter.

Overall  $d'$  was computed by collapsing across all five transition conditions. As shown in Fig. 3B,C the model provides a good fit to (mean) human performance (wide-prior:  $r^2 = 0.92$ ,  $p < 0.0001$ ; multi-prior:  $r^2 = 0.85$ ,  $p < 0.0001$ ). We further examined the relationship between human performance and two additional model perceptual parameters: memory and observation noise. The human performance ( $d'$ ) is positively correlated with memory (wide-prior:  $r^2 = 0.54$ ,  $p = 0.028$ ; multi-prior:  $r^2 = 0.31$ ,  $p = 0.021$ ), suggesting that listeners modelled to have longer integration horizons were empirically better able to detect sequence transitions. No correlation with observation noise was observed (Figure 3C right). Overall, this model, and in particular the 'wide-channel' prior, sensitized to unimodal distributional statistics, successfully captured human performance across a range of transition types. It also demonstrated that most listeners based their transition detection on a context length of 20–40 tones (1–2 seconds).

### 2.3. The human brain automatically tracks the statistics of stochastic acoustic patterns

The initial data above are based on behavioural responses to perceived changes in sequence characteristics. To understand the more general case, we turned to EEG to ask whether the brains of *naïve listeners* are sensitive to changes in stochastic acoustic patterns - and if so, what sequence statistics are being monitored. In Experiment 2 A, participants were exposed to the R20-R10(H), R20-R10(M), R20-R10(E) transition conditions, along with R20 as a no-change control (50 % of trials). (Unlike the behavioural task, all conditions were presented in random order, trial-by-trial). The R20-R10(L) condition was omitted to reduce experiment length, due to its similarity to R20-R10(H). In Experiment 2B, listeners listened to the R20-R10 (full spectrum) condition along with R20 as a control. In all experiments that follow, participants were naïve to the stimuli and watched a silent movie while listening.

Based on previous results showing increases in the sustained EEG

power with increases in predictability (Barascud et al., 2016; Southwell et al., 2017; Southwell and Chait, 2018; Herrmann et al., 2021; Herrmann and Johnsrude, 2018), we expected that listeners' sensitivity to statistical transitions would be reflected in an EEG divergence between the responses to R20-R10(H/M/E) versus no-change R20.

As shown in Fig. 4, all conditions evoked a similar EEG response: an N1-P2 onset complex followed by a persisting sustained response, and an offset response (see the topographies in Fig. 4A). In the four change conditions, clear evoked responses are observed post-transition (Fig. 4, A-D). The sustained response gradually diverged from R20 after the transition. Topographies (Fig. 4F) show that this amplitude increase in sustained response is attributed to a steady rise of negative activity in central channels, consistent with effects observed previously (Barascud et al., 2016; Southwell et al., 2017). As discussed also elsewhere (Barascud et al., 2016; Southwell et al., 2017), the divergence between conditions is not outwardly consistent with low level adaptation effects. An adaptation-based explanation would in fact predict the opposite pattern: R10(H/M/E) sequences are associated with a *higher* probability of occurrence per tone, and hence should cause *greater* adaptation than R20.

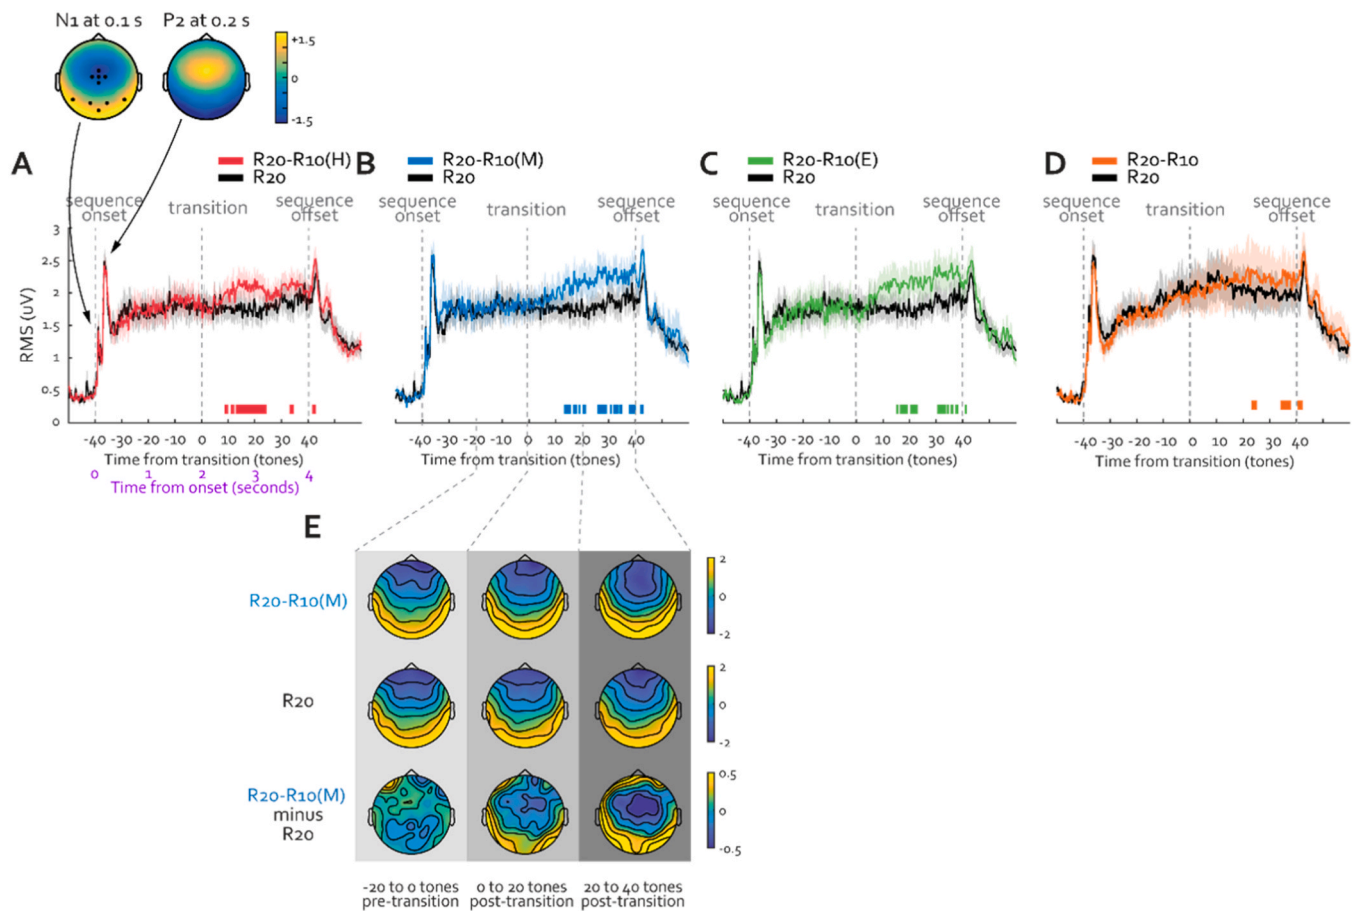
The sustained response amplitude did not differ across conditions (Repeated measures ANOVA on mean power post-transition between H, M and E;  $F(1.60, 39.98) = 0.17$ ,  $p = 0.80$ ,  $\eta^2 = 0.007$ ). Considering the behavioural detection performance as an indicator of the salience of the pattern change, the absence of variation in the sustained response amplitude among these three conditions implies that the sustained response is likely not a direct reflection of the behavioural salience of the transition itself. Instead, as will be shown in Experiment 3, below, it appears to be linked to the informational content inherent in the post-transition sequence.

### 2.4. The EEG sustained response in passively listening, individuals reflects sequence statistics

A potential explanation for the observed effects in Experiment 2 is that the amplitude of the sustained response is not associated with stimulus statistics *per se*, but instead reflects a heightened attentional state triggered by the *change* in the unfolding sequence. (Note that transitions were always task irrelevant, and participant attention was directed toward a silent movie). To examine this possibility, Experiment 3 selected one of the conditions - R20-R10(M), our "anchor condition" - and introduced a *reverse* transition R10(M)-R20, where the spectral distribution changes from reduced to full. (The stimulus set also included two no-transition conditions: R10M and R20).

If the gradual increase in the sustained EEG response merely reflects *detection of a statistical change*, then responses to R10M-R20 and R20-R10M should be similar because both contain statistical changes partway through the sequence. However, if the increased EEG power is *specific to the statistical properties of the spectral content before and after the transition*, R10M-R20 should evoke a different response compared to R20-R10M.

The latter predictions held (Fig. 5). The response to R20-R10(M) shows a slow rise in the sustained response (Fig. 5A), replicating the pattern observed in Experiment 2. The *opposite* pattern is observed in the response to the R10(M)-R20 transition, relative to its no-change control R10(M) (Fig. 5B). Shortly after the R10(M)-R20 transition, there is an MMN-like response, corresponding to the presentation of a novel frequency (the first tone of the R20 portion of the sequence, see the high-pass filtered response at the bottom of Fig. 5B). This observation confirms that listeners are sensitive to the introduction of novel frequencies in the R10(M)-R20 condition. Following this response, a sharp drop in the sustained response is observed, later settling at a level below that of R10(M) for the remainder of the trial. This suggests that the intrinsic statistics of each stochastic tone distribution, rather than the transition between distributions, may drive the sustained EEG response. This interpretation is further supported by the observation that the R10(M)



**Fig. 4.** EEG results of Experiments 2 A and 2B. (A–D) Plotted are the group mean of EEG responses from naïve participants passively listening to sequences while watching a silent film. Root mean square (RMS) was computed for each participant over 10 auditory sensitive electrodes, selected based on the N1 response evoked by the sequence onset (left topography in A). Shaded areas are  $\pm 1$  SEM (standard error of the mean). Colour-coded horizontal lines at graph bottom indicate time intervals where bootstrap statistics show significant differences between each transition condition and R20 (black trace). All trials had a fixed total length of 4 seconds (80 tones), with a potential ( $p = 0.5$ ) transition in tone pool size starting at 2 seconds. The sequence onset, transition and offset are all marked with grey vertical dashed lines. Data in (A) to (C) are from Experiment 2 A ( $N = 26$ ) and data in (D) is from Experiment 2B (a different group of participants;  $N = 26$ ). (E) Topographies for R20-R10(M) and the control R20 are plotted over three key time periods: (1) 1 second before transition (–20th to 0th tone pre-transition) (2) 1 second after transition (1st to 20th tone post-transition) and (3) the last second before offset (21st to 40th tone).

no-change condition elicits a greater sustained response amplitude than the R20 no-change condition (Fig. 5C).

Overall, this pattern is consistent with the hypothesis that the sustained EEG responses reflect a brain state, as opposed to a transition-related response, and moreover that sustained EEG responses track the statistical structure of the sequence. The specific observation of a higher sustained response evoked by R10M than R20 is consistent with the hypothesis from Barascud et al. (2016) that the sustained response may track sequence precision: the predictive distribution for R10(M) is expected to be narrower (hence more “precise”) than that for R20.

### 2.5. Relating D-REX outputs to EEG responses

To determine whether the EEG brain response pattern in naïve listeners can be explained as arising from a process that accumulates distributional statistics, we asked which D-REX model outputs might correspond to the brain responses we observed. As discussed in the introduction, D-REX tracks two key inferred statistics: (1) *Change probability* (the probability that at least one changepoint has occurred); and (2) *Precision* (the reliability of the prediction). In D-REX, precision is quantified as the maximum value of the normalised predictive distribution (labelled in Fig. 2B).

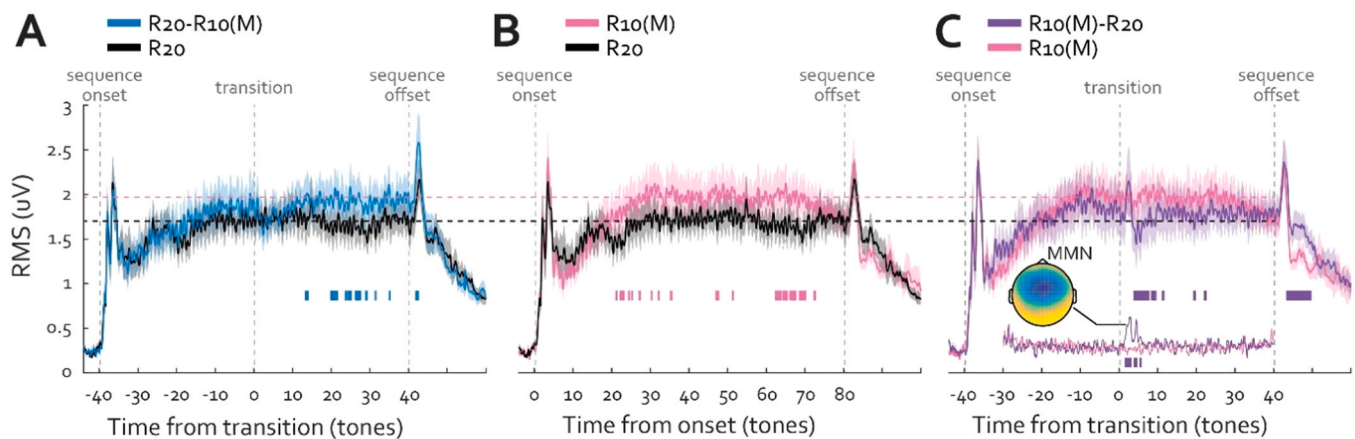
To simulate the model outputs for the EEG stimuli, we applied the model to the stimuli in Experiment 2. As no behavioural response was

acquired in Experiment 2 (passive listening), ideal observer model parameters were used (infinite memory and zero noise). The model was then run on the entire experimental session in the same way it was delivered to each of the 26 participants (i.e., 600 trials, including 100 trials for each transition condition and 300 trials of R20).

Fig. 6 shows the model outputs for the “wide-channel” and “multi-channel” priors. Change probability (Fig. 6, A & C) grows post-transition, but at rates that differ markedly across conditions. Transition from the full-frequency tone pool (R20) to higher frequency tones only (R10(H)) evoked the steepest increase in change probability, followed by transitions to the middle frequency tone pool (R10(M)), and the high and low ‘edge’ tone pool (R10(E)). Precision also grows after the transition, reflecting attunement of the model to the underlying statistics of the sequence (Fig. 6, B & D).

Notably, the wide-channel prior model outputs did not differentiate R20-R10 from R20, either in change probability (Fig. 6A) or precision (Fig. 6B). The EEG response (Fig. 6E) is clearly not consistent with Wide-Channel change probability (Fig. 6A) or precision outputs (Fig. 6B). By contrast, the “multi-channel” prior model, in particular the “precision” output (Fig. 6D) provides a good qualitative match to the EEG responses, with responses to R20-R10(H), R20-R10(M) and R20-R10(E) reaching a similar level after the transition, and with activity in R20-R10 rising at a slower rate.

However, the stimulus set in Experiment 2 is not optimal for



**Fig. 5.** EEG results in Experiment 3 ( $N = 22$ ). EEG responses to R20-R10(M), R10(M) and R10(M)-R20 against their respective controls plotted against time. The shaded area shows 1 SEM (standard error of the mean). Colour-coded horizontal lines at graph bottom indicate time intervals where bootstrap statistics show significant differences between each change condition and its no-change control. (A) The response to R20-R10(M) and its control R20 replicates the pattern observed in Experiment 2 A (Fig. 4B). (B) Compared with R20, the brain response to the no-change R10(M) shows a generally higher sustained response, emerging from 21.5 tones after sequence onset. The dashed pink horizontal line across plots A, B, and C indicates the sustained response amplitude of R10(M), and the dashed black line indicates that of R20. (C) The brain response to R10(M)-R20 shows a sharp drop in the sustained response after the transition and then settles at a lower-level post-transition compared with R10(M). The MMN response to the first novel tone in the post-transition sequence can be clearly observed in the 2Hz-high-pass-filtered trace, which is included in the bottom of the plot. The corresponding topography which is consistent with that commonly observed for MMN is also provided.

adjudicating between multi-channel model predictions because the pattern of results for change probability and precision is quite similar for the three transition stimuli (R20-R10(H), R20-R10(M), R20-R10(E)). In the final experiment we address this issue by employing a different stimulus set for which the model should predict divergent dynamics.

### 2.6. Automatic brain responses track the precision of unfolding sound sequences

We introduced a new type of transition, going from the full set of 20 to five tones equally distributed across the spectrum (R20-R5), comparing it against our “anchor stimulus” - R20-R10(M) (Fig. 7A). R20-R5 was chosen because it has very different underlying stochasticity (see below), which in turn drives different model predictions. In a pilot experiment (see Methods Experiment 4 A) we confirmed that R20-5 evokes similar behavioural performance to R20-R10(M).

Model outputs (change probability and precision against time) for the wide- and multi-channel prior models are plotted in Fig. 7, B-E. The wide-channel prior model shows substantially higher change probability (Fig. 7B) and precision (Fig. 7C) for the R20-R10(M) condition compared to R20-R5. By contrast the multi-channel prior model shows very similar change probability for both transition conditions (Fig. 7D), but higher and faster-rising precision for R20-R5 compared to R20-R10(M) (Fig. 7E). This is consistent with the fact that the spectral likelihood distribution for that condition is narrower than for R20-R10(M).

Which model output best predicts passive EEG responses to the same stimuli? In the R20-R10(M) condition, we again find a slow rise in the sustained response (Fig. 7F, blue curve), consistent with that seen in Experiments 2A and 3. Critically, the R20-R5 sequence evokes a markedly different response (Fig. 7F, grey curve), one uniquely mirroring the precision pattern predicted by the multi-channel prior model (Fig. 7E). Significantly greater EEG amplitude for R20-R5 versus R20-R10(M) emerged from 422 ms after onset and was sustained for the remainder of the trial.

A follow-up, channel-by-channel analysis was conducted to assess the consistency of this effect across scalp loci. We first identified the electrodes ( $N = 40$  out of a total of 64) that reliably responded to the transitions (a significant difference between R20-R10(M) and R20 and R20-R5 and R20 in the interval 1–2 seconds post-transition; see *Methods - EEG channel-by-channel analysis*). For each of the transition-sensitive electrodes, we then computed the amplitude difference between R20-

R5 and R20-R10(M) at 1–2 seconds post transition. These values are illustrated in Fig. 8. This analysis confirmed that no electrode exhibited  $R10M > R5$ , providing further evidence against the “wide-channel” model. The vast majority of the transition-sensitive electrodes exhibited a larger response to R5 relative to R10(M), indicative of precision tracking, and in-line with the RMS analysis above.

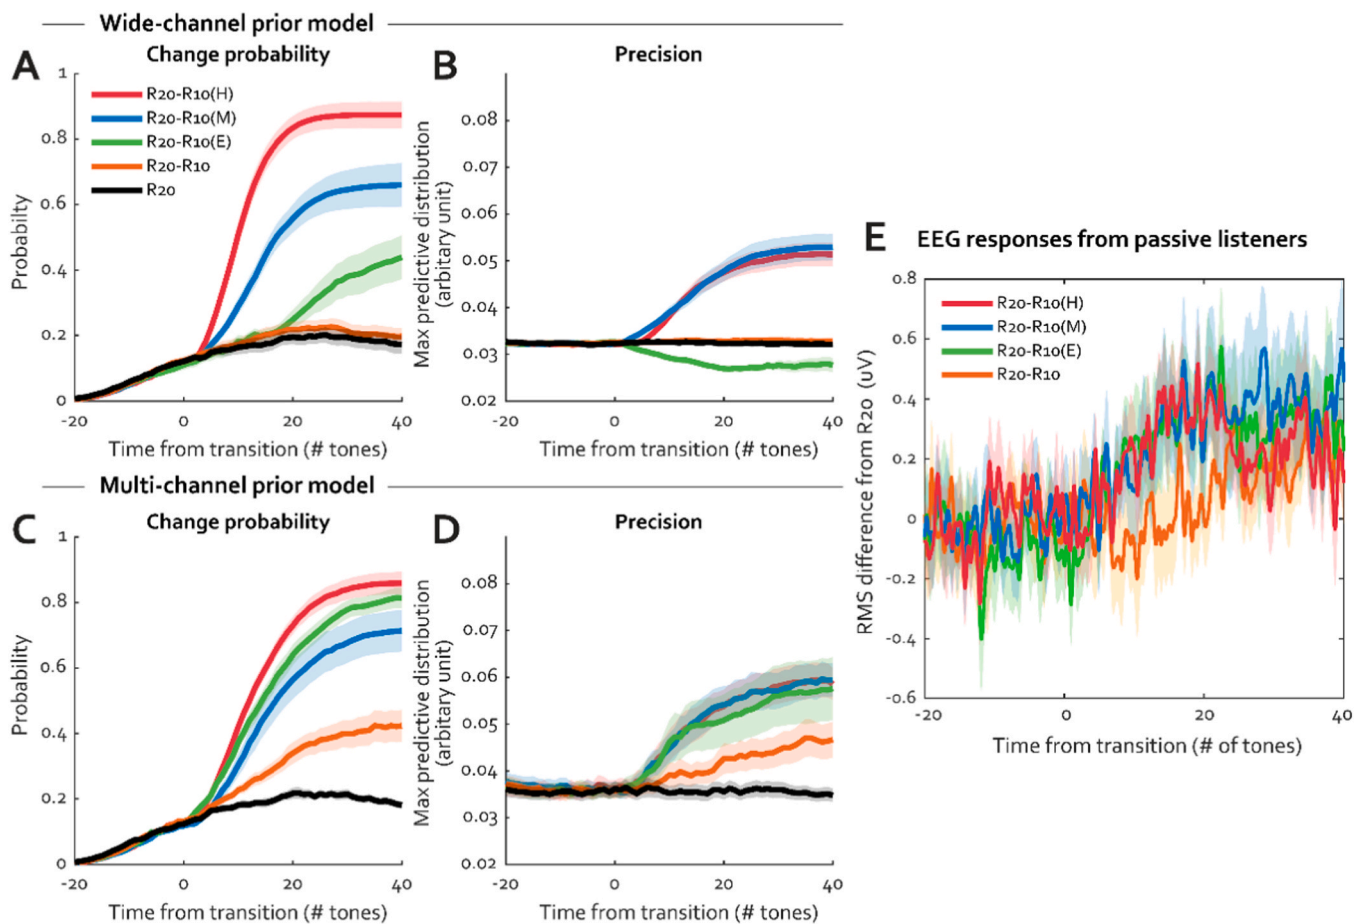
The close similarity between the D-REX Multi-Channel precision output and these EEG data suggests that the passively listening brain utilises a within-channel listening strategy, and that the brain responses driving the EEG sustained response reflect the precision of the internal model of the sound sequence constructed during listening.

An alternative explanation for the observed effects might involve spectral resolution or tracking ability. However, previous results e.g. from Barascud et al. (2016) challenge this interpretation. They showed that sustained responses associated with sequences consisting of the same frequency components but arrange such that they are either predictable (high precision) or less predictable (low precision) are associated with different sustained responses. Conversely, here R10(E) differs in gross frequency spacing from R10(H), but they have similar EEG responses. Thus, we suggest that spectral resolution or attentional spectral tracking capacity is unlikely to account for these effects.

### 3. Discussion

Listeners automatically track the statistics of unfolding sounds even when these are not relevant to behaviour. Transitions from R20 (uniform sampling from the full frequency pool) to various constrained sampling regimes drove an increase of the sustained EEG response. This was observed to a range of distributional statistics, and even in situations (R20-R10) where behavioural detection of these transitions in Experiment 1 was at floor. Similar EEG responses were previously reported for deterministic patterns (Sohoglu and Chait, 2016; Barascud et al., 2016; Southwell et al., 2017; Southwell and Chait, 2018; Herrmann et al., 2021; Herrmann and Johnsrude, 2018). That they are also seen for stochastic transitions supports the hypothesis that EEG sustained response modulation reflects a general process of belief updating.

The EEG data were compared to a Bayesian predictive inference model fit to the acoustic scenes experienced by listeners. This comparison showed that EEG responses were most closely associated with coding precision, a representation of the uncertainty associated with the current predictive distribution.



**Fig. 6.** Time-course of modelling outputs and EEG responses from human listeners. Change probability (A) and precision (B) are plotted as a function of time (in number of tones) from the transition for the wide-channel prior model. Outputs of the multi-channel prior model are shown in C and D. For each transition condition, the change probability and precision (see Fig. 2B for explanation) were averaged across 100 trials presented to the model, mirroring the evoked response EEG analysis (also based on averaging brain activity across trials). Generally, change probability and precision exhibit a rise, followed by a plateau shortly after the stimulus transition. The specific dynamics reflect the statistics of the stimulus. The unique ‘undershoot’ response observed in the wide-channel prior precision time series for the R20-R10(E) condition (green line in panel B) is a function of both the prior distribution and the underlying distribution of the stimuli. Here, the wide-channel model prediction widens yet further in response to the bimodal stimuli, leading to a reduction in precision from the initial predictive distribution. (E) EEG responses to the four transition conditions (R20-R10(H), R20-R10(M), R20-R10(E) and R20-R10) are overlaid together (previously shown in Fig. 4 A-C). Because the responses are taken from two different experiments (Experiment 1 A and 1B; see Fig. 4) each change condition (R20-R10(H), R20-R10(M), R20-R10(E), R20-R10) is normalised by subtracting the relevant R20 baseline activity. This is done to highlight the specific effects of the transitions on the EEG response, independent of the baseline activity. The shaded area in A-D is 1 standard deviation of bootstrap of 100 trials presented to the model. The shaded area in E indicates 1 SEM (across participants). All outputs are plotted against time (in #tones) from the transition.

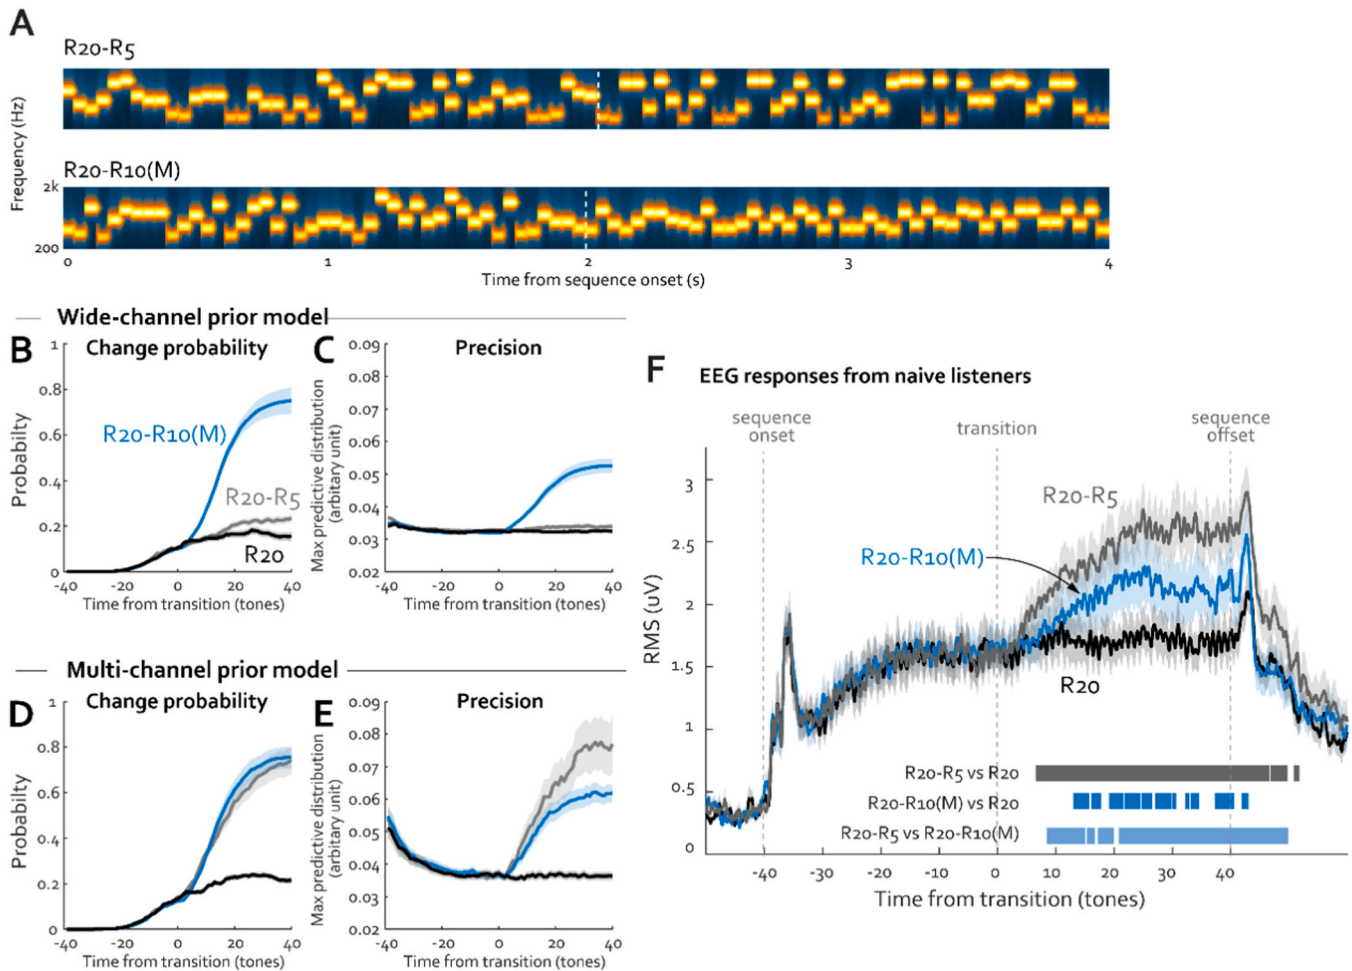
### 3.1. Listeners are sensitive to stochastic regularities

Accumulating work demonstrates that the human brain is tuned to the statistical properties of the stochastic ongoing acoustic environment (Skerritt-Davis and Elhilali, 2021a, 2018; Garrido et al., 2013; Winkler et al., 1990; Overath et al., 2007; Furl et al., 2011; Rubin et al., 2016). By modelling brain responses to two-tone sequences of different base probabilities, Rubin et al. (2016) showed that trial-wise neural responses in auditory cortex are well explained by the probability of occurrence of each tone frequency, calculated from the recent history of the sequence. The models that best fit neural responses tracked a relatively long stimulus history (~10 tones) but encoded a coarse representation of stimulus properties, based on a small set of 0-th order statistics. In line with this conclusion, Garrido et al. (2013) (also see Daikhin and Ahissar, 2012) demonstrated that MEG responses to probe (outlier) tones are sensitive to the statistical context (mean and variance of frequency) of randomly generated tone-pip sequences. Larger responses occurred to the same probe tone when presented in a context with low, as opposed to high, frequency variance. However, through its emphasis on the effect of various statistical contexts on the processing of *deviant* tones, this work

has provided largely indirect insight into the brain and perceptual mechanisms through which listeners track the statistics of stochastic auditory environments.

More recently, several investigations have focused on the factors that affect change detection in stochastic sound sequences. Boubenec et al. (2017) used a broadband ‘tone cloud’ stimulus composed of tone-pips randomly drawn from a spectral probability distribution. They demonstrated that listeners’ change detection accuracy depended on the spectral properties of the changed distribution. Change- and decision-evoked EEG responses were identified in centro-parietal electrodes, consistent with other evidence integration tasks (O’Connell et al., 2012; Kelly and O’Connell, 2013; Yao et al., 2020). Using fractal sequences, Skerritt-Davis and Elhilali (2021a, 2018) found that listeners’ change detection and surprisal-associated brain responses were consistent with an observer model that tracks the higher-order statistics of temporal covariance structures in the sequences.

Whereas these previous studies used active monitoring of changes-points in auditory sequences, the present study demonstrates automatic statistical tracking during passive listening. The brain responses implicated in belief updating originated from central scalp locations,



**Fig. 7.** Comparing the brain response and modelling responses to R20-R5 and R20-R10(M) in Experiment 4. (A) Example spectrograms of the two transition conditions used in Experiment. R20-R5 contains a reduction in alphabet size partway through the sequence (from 20 to 5, sampled equally from the full pool). (B) Like Fig. 6, A-D, time-courses of model outputs — change probability and precision— are plotted from the transition. The outputs of both wide-channel prior model (B and C) and the multi-channel prior model (D and E) are shown. The outputs were averaged across 100 trials presented to the model. The shaded area in B-E is 1 standard deviation of bootstrap. (F) EEG responses from naïve passive listening listeners in Experiment 4B. Plotted are the group mean of RMS over 10 auditory sensitive electrodes. Shaded areas are 1 SEM (standard error of the mean). Colour-coded horizontal lines at graph bottom indicate time intervals where bootstrap statistics show significant differences between each transition condition and R20 (black trace). The bottom horizontal line in light blue indicates intervals where a significant difference between two transition conditions was observed.

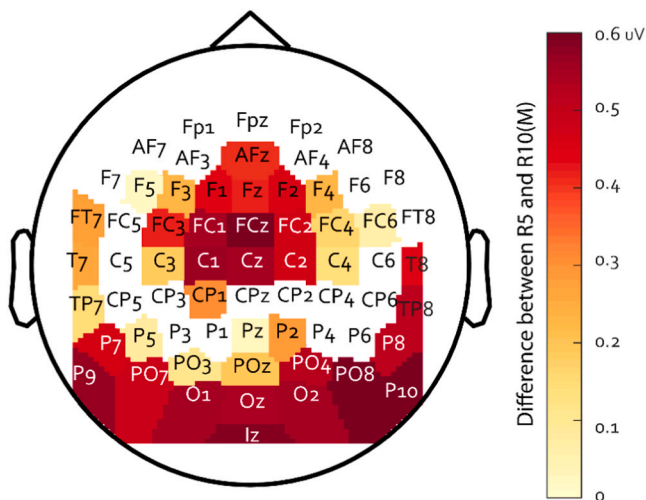
likely reflecting auditory cortical activity. The electrodes exhibiting the sustained response modulation were also those that responded most strongly to auditory onsets, suggesting that auditory cortex is involved in tracking these statistics. This contrasts with Boubenec et al., 2017 where evidence accumulation-activity was concentrated in centro-parietal electrodes. The difference between our findings and those in (Boubenec et al., 2017) could potentially be attributed to task-related factors. In our study, the observed responses reflect automatic tracking, in contrast to Boubenec et al., 2017 where participants actively monitored for sound changes, engaging explicit “evidence accumulation” mechanisms.

### 3.2. Bayesian tracking of rapidly unfolding sounds

We used a Bayesian model of statistical tracking (“dynamic regularity extraction”, D-REX (Skerritt-Davis and Elhilali, 2021b) to constrain our interpretation of EEG dynamics. A detailed account of the model is presented elsewhere (Skerritt-Davis and Elhilali, 2021b), see also (Skerritt-Davis and Elhilali, 2021a, 2018). D-REX is similar to frameworks widely used to model various aspects of human behaviour under uncertainty (Behrens et al., 2007; Daw et al., 2006;

Krishnamurthy et al., 2017; Nassar et al., 2010, 2012; Wilson et al., 2013; Wilson and Niv, 2012). Specifically, the model is based on key assumptions in predictive coding: the brain (a) builds beliefs by collecting statistical estimates over time; and (b) simultaneously maintains multiple hypotheses regarding the stimulation context it has just experienced. The latter assumption incorporates a focus on detecting changes in the underlying generative distribution, often referred to as “change point detection.” Each sequentially encountered stimulus is used to update these hypotheses in an “observe-predict-update” loop (Knill and Pouget, 2004; Tenenbaum et al., 2006; Daunizeau et al., 2010).

Prior work in perceptual decision making — in both visual and auditory fields — used a modelling approach similar to D-REX to show that observers track change point probability. In other words, observers estimate sensory uncertainty from immediate and past signals, and use it to optimize their prediction of future variance (Nassar et al., 2010, 2012; Krishnamurthy et al., 2017; Norton et al., 2019; Beierholm et al., 2020). In the context of auditory processing, D-REX has been used to model a variety of perceptual phenomena related to understanding the extent to which auditory information is represented in memory (Skerritt-Davis and Elhilali, 2021a, 2018, 2021b). In tasks that required listeners to monitor for changes in rapid sound sequences, the model successfully fit



**Fig. 8.** Channel-by-channel analysis of responses to R20-R5 and R20-R10(M) in Experiment 4. Plotted is the mean difference during the 1–2 second post-transition period between these conditions in each electrode. A higher positive value signifies a higher amplitude response to R5 compared to R10(M).

performance. The model results suggested that listeners track marginal mean variance, and covariance (temporal dependency) between successive elements. The free model parameters: *memory* (maximum context window) and *observation noise* accounted for individual variability in sensitivity. Similar effects were seen in our behavioural experiment (Experiment 1). We observed high correlations between model performance and individual subject behaviour, in particular for the memory parameter which suggested that listeners use a context of 20–80 tones to inform their decision about change occurrence.

### 3.3. EEG responses in naïve listeners reflect within-channel precision tracking

To explain the dynamics of EEG responses, we focused on two key instantaneous outputs of the model: (1) its beliefs about *the probability that change point has occurred*, i.e., that the generative distribution underlying the stimulus sequence has changed; and (2) *the precision of its predictive representation*, reflecting the accuracy (or conversely "expected uncertainty"; (O'Reilly, 2013)) with which future inputs can be predicted. We also evaluated two potential listening strategies by initiating the model with two types of priors: either a single Gaussian (*Wide Single-channel prior*), reflecting the hypothesis that listeners initially track a single wide-band spectral distribution; or a multimodal Gaussian (*Multi-channel prior*), where listeners employ a within-frequency-channel listening strategy, where several concurrent spectral distributions are initially tracked.

Relating the EEG responses to outputs of the D-REX model demonstrated that the sustained response pattern was consistent with the "multi-channel prior" model. This is in line with many existing demonstrations (Chi et al., 2005; Santoro et al., 2014; Bitterman et al., 2008) that the auditory system's "default mode" is to process sound via multiple separate spectral channels. Interestingly, it appears that behavioural performance (Experiment 1, Fig. 3) was better modelled with the "wide-channel prior" model, hinting at a potential distinction between active statistical tracking and passive listening, an initial finding which should be followed up in future work. One difference worth noting is that the EEG experiment used randomized transitions across trials. In contrast, the behavioural experiment used blocked trials which may have guided listener expectations to some degree.

Earlier observations in the context of deterministic sequences showed that the sustained response is modulated by sequence precision: Regular patterns consistently evoke larger sustained EEG responses than

do random patterns (Barascud et al., 2016; Southwell et al., 2017; Southwell and Chait, 2018; Herrmann et al., 2021; Herrmann and Johnsrude, 2018). In addition, more predictable random patterns (i.e., smaller alphabet size) evoke higher sustained EEG responses than less predictable random patterns (Barascud et al., 2016). Another study (Sohoglu and Chait, 2016) also reported similar effects in the context of artificial "scenes" that were populated by temporally regular vs. random streams. Extending these observations to stochastic patterns, we show that the sustained response recorded in naïve passive-listening listeners reflects a brain state that varies systematically with sequence precision. This is consistent with precision weighting being an optimal evidence integration strategy (Yon and Frith, 2021), and establishes the sustained M/EEG response as a potential neural signature for precision tracking.

The ability to accurately estimate precision allows an organism to operate optimally under different levels of environmental volatility by reweighing signals according to their inferred reliability (Yon and Frith, 2021; Diederer et al., 2016). This increases the salience of deviants in a high precision context, indicative of a genuine change in the environment. Generally, high precision signifies a stable and predictable environment, which organisms tend to favour (Yon and Frith, 2021). From a survival standpoint, it likely triggers a sense of relative safety, activating the parasympathetic nervous system and facilitating a "rest-and-digest" mode (Milne et al., 2021). This state may be reflected by the large shift in sustained response we observe.

That EEG responses do *not* reflect change probability tracking might be somewhat surprising given the key importance of sensitivity to environmental change to perception and survival. However, this finding is consistent with observations from Boubenec et al (Boubenec et al., 2017), who also did not find any evidence for change-point detection-related activity in auditory cortex. One possibility is that change probability is being tracked in auditory cortex but in a way that is not reflected in the EEG activity we analysed.

Furthermore, it is important to note, that our focus here was on a specific Bayesian model, that was selected due to its widespread use in the literature (Nassar et al., 2010; Wilson et al., 2013), and for the EEG data, the comparison remains qualitative. While the results suggest that EEG responses are closely related to Bayesian precision, to draw stronger inferences about Bayesian versus non-Bayesian processes, it is essential to demonstrate that the findings generalize across a range of parameter values and models, including approximate or heuristic models. It is possible—and perhaps likely (Nassar et al., 2010; Krishnamurthy et al., 2017; Spaak, 2024)—that the brain does not perform full Bayesian updating but instead computes simpler metrics, such as indices of entropy, temporal variability, or similar measures.

Moreover, we deliberately employed a very simple stimulus to maintain full control over the auditory statistics presented to listeners and to enable straightforward tracking and quantification. While this approach ensures robust interpretations of the data, naturalistic stimuli typically exhibit richer and more diverse statistical structures than those examined here. Understanding how these findings might extend to more complex, ecologically valid conditions is an important challenge for future research.

### 3.4. The neural correlates of the EEG sustained response

The neural-circuit underpinnings of the sustained-response modulation effects remain elusive. According to one proposal, precision is encoded by gain in superficial pyramidal cells in sensory cortex (Feldman and Friston, 2010; Aukstulewicz et al., 2017), potentially via neuromodulation. In particular, tonic Acetylcholine (ACh) has been shown to be modulated by environmental uncertainty (Dalley et al., 2001; Yu and Dayan, 2005; Bland and Schaefer, 2012) (also see (Pérez-González et al., 2024)).

An alternative, though not mutually exclusive, account links precision tracking to inhibitory mechanisms. Growing evidence from animal electrophysiology and modelling (albeit using very simple sounds)

demonstrates that inhibitory plasticity plays a critical role in shaping brain responses to repeating stimuli. Specifically, increased inhibition onto excitatory neurons tuned to familiar (repeating/predictable) stimuli might constitute a cortical mechanism for precision tracking (Yarden et al., 2022; Natan et al., 2015, 2017; Richter and Gjorgjieva, 2022; Schulz et al., 2021). Under this hypothesis the modulations in the sustained response may be reflecting increased inhibitory activity. Indirect evidence from recent dynamic causal modelling (DCM) work has also suggested that encoding of precision is attributable to inhibitory mechanisms (Lecaigard et al., 2022).

A specific role for inhibition in controlling responses to predictable sensory stimuli is also consistent with behavioural results showing that rather than attracting attention (which would be expected from an excitatory effect), high-precision patterns are easier to ignore (Southwell et al., 2017), and are associated with reduced arousal (Milne et al., 2021). M/EEG (or BOLD) cannot straightforwardly resolve inhibitory from excitatory activity. Progress in testing this account therefore depends on targeted cellular-level or pharmacological investigations.

## 4. Materials and methods

### 4.1. Ethics statement

Experimental procedures were approved by the research ethics committee of University College London, and written informed consent was obtained from each participant. Participants were paid for their participation.

### 4.2. Stimuli

The stimuli were constructed by selecting tone-pips (50 ms long, gated with 5 ms raised-cosine ramps) from a pool of 20 frequencies ('alphabet'; 222–2000 Hz, ~12.2 % steps or ~2 semitones, as in Barascud et al., 2016). The frequency steps provide sufficient spacing to allow differentiation, while the short unit ensures that listeners cannot explicitly track the sequence (Warren and Ackroff, 1976; Warren, 2008; Warren and Obusek, 1972).

The basic stimuli were random sequences constructed by sampling frequencies uniformly at random from the full pool ("alphabet" = 20). This condition is referred to as "R20" (Fig. 1, top). Transition stimuli were created by altering the sampling regime partway through the sequence.

The full pool of 20 frequencies was divided into four logarithmically equally sized frequency bands, each containing five frequency values. Different conditions were created by reducing the *alphabet* to 10 or 5 frequencies, and manipulating how these are selected. Spectrograms of example stimuli can be found in Fig. 1. Sound examples can be found online [<https://osf.io/4dkn3/>].

In **R20-R10**, the fully random (R20) sequence transitioned into a sequence that contained 10 frequencies (R10) uniformly sampled from across the full pool (2 from each band).

- (a) **R20-R10(H)** contained a transition into a sequence comprising 10 frequencies sampled from the two highest bands (707 Hz to 2000 Hz).
- (b) **R20-R10(L)** contained a transition into a sequence comprising 10 frequencies sampled from the two lowest bands (222 Hz to 630 Hz).
- (c) **R20-R10(M)** contained a transition into a sequence comprising 10 frequencies sampled from the two middle bands (397 Hz to 1122 Hz).
- (d) **R20-R10(E)** contained a transition into a sequence comprising 10 frequencies sampled from the highest (1259 Hz to 2000 Hz) and lowest (222 Hz to 354 Hz) bands.

In Experiment 4 we also used

- (e) **R20-R5** which contained a transition into a sequence comprising 5 frequencies uniformly sampled from across the full pool (1 from each band; 1 selected at random from the remaining values). (Fig. 7A)

All stimuli were generated anew for each trial in each subject. In the behavioural experiments (Experiment 1 and Experiment 4 A), the transition time was jittered between 40 and 50 tones after sequence onset. In all EEG experiments (Experiments 2 A, 2B, 3 and 4), the sequence length for each trial was fixed to 80 tones and the change position was fixed to the 41st tone.

A summary of the different conditions across all experiments is shown in Table 1 below.

### 4.3. Experiment 1

#### 4.3.1. Participants

Ten participants (10 females; aged 19–24, average 21.6) took part in the experiment.

#### 4.3.2. Stimuli

The stimulus set included R20 (control) and five transition conditions: R20-R10(H), R20-R10(L), R20-R10(M), R20-R10(E), and R20-R10. Each condition was presented in a separate block (20 transition conditions plus 20 R20 controls; block order was randomized across participants). In a pilot study, we used a mixed design presenting all conditions in a randomised order. The results were similar to those obtained in this blocked design. We chose the blocked design in the main experiment for two reasons: (1) It allows the report of false alarms for each change condition, which is important for model fitting. (2) Given the small number of trials, the blocked design allowed participants to quickly optimise their listening strategy and refine their performance for each condition in the most efficient way.

Transition times were jittered by assigning a random length to the pre-transition segment; the transition time was therefore unpredictable. Trials were 80–90 tone-pips long, lasting 4–4.5 seconds.

To make this obvious change unpredictable, the transition time and sequence length were also jittered in exactly the same way as for the experimental transition conditions. The inter-trial interval (ITI) was randomised between 1 and 1.5 seconds.

The stimuli were delivered to the participants' ears by Sennheiser HD558 headphones (Sennheiser, Germany) via a UA-33 sound card (Roland Corporation) at a comfortable listening level, self-adjusted by each participant. Stimulus presentation and response recording were controlled with the Psychtoolbox package (Psychophysics Toolbox Version 3) (Brainard, 1997) in MATLAB (The MathWorks, Inc.).

Participants were tested in a darkened, acoustically shielded room

**Table 1**

Different stimulus conditions used across all six experiments in the present study.

Condition	1 (Behaviour)	2 A (EEG)	2B (EEG)	3 (EEG)	4 A (Behaviour)	4 (EEG)
R20	x	x	x	x	x	x
R20-R10	x		x			
R20-R10 (H)	x	x				
R20-R10 (L)	x					
R20-R10 (M)	x	x		x	x	x
R20-R10 (E)	x	x				
R10(M)- R20				x		
R10(M)				x		
R20-R5					x	x

(IAC triple-walled sound-attenuating booth). Each experimental session lasted about 1 hour and began with a short practice session, followed by the main experiment. Participants were instructed to fixate their gaze on a white cross at the centre of the computer screen while listening to the stimuli and to respond by pressing a keyboard button as soon as they detected the transitions. Feedback (in the form of a tick or a cross) was provided after each trial. Each block lasted about 6 minutes and participants were allowed short breaks between blocks.

#### 4.4. Experiment 2 - EEG

##### 4.4.1. Participants

**Experiment 2 A:** Data from 26 naïve participants (18 females; aged 21–38; average 24.1; 2 left-handed) are reported. No participant was excluded in Experiment 2 A.

**Experiment 2B:** Data from 26 naïve participants (17 females; aged 20–32, average 23.8, all right-handed) are presented. One additional participant was rejected due to excessive missing data (one block of data failed to record).

All participants reported normal hearing with no history of neurological disorders. None of the participants was a professional musician or known to possess absolute pitch.

##### 4.4.2. Stimuli

**Experiment 2 A:** The stimulus set included R20-R10(H), R20-R10(M), R20-R10(E) and their non-change control R20. Overall, 600 stimuli (100 each of R20-R10(H), R20-R10(M), R20-R10(E), and 300 of R20) were presented to the listeners in a random order.

**Experiment 2B:** The stimulus set consisted of two stimulus conditions: R20-R10 and its no-change control R20. In total 200 stimuli (100 each of R20-R10 and R20) were presented in random order.

All stimuli were 4 seconds long (80 tones) with the transition occurring at 2 seconds post onset (at the 41st tone-pip). The inter-trial interval (ITI) varied between 1.5 and 2 seconds. Stimuli were generated offline in MATLAB and presented binaurally with the Psychophysics Toolbox (Brainard, 1997; Kleiner et al., 2007) using EarTone in-ear earphones at a comfortable listening level of ~60 dB SPL self-adjusted by each participant.

##### 4.4.3. Procedure

After the placement of the electrodes, participants were seated in front of a computer monitor and instructed to watch a silent subtitled documentary throughout data collection. Attention to the silent film was not explicitly quantified. In a previous study (Barascud et al., 2016), we employed a simple 1-back visual memory task to keep naïve listeners distracted. However, when piloting this study, we found that a silent film was more effective in preventing fatigue and keeping participants awake during the session. The silent film was chosen by participants from a selection of calm, nature-related documentaries, promoting engagement.

The monitor was adjusted to present the subtitles at eye level. The experiment was organized into 10-minute blocks (6 blocks in Experiment 2 A; 2 blocks in Experiment 2B). Short breaks were provided between blocks, but participants were asked to remain still in the breaks.

EEG was recorded through a 64 channel Biosemi system (Biosemi Active Two AD-box ADC-17, Biosemi, Netherlands) with a 64-electrode cap, referenced to the CMS-DRL ground, which functions as a feedback loop driving the average potential across the montage as close as possible to amplifier zero. Acquisition was continuous with a sampling rate of 2048 Hz.

#### 4.5. Experiment 3 - EEG

##### 4.5.1. Participants

All participants from Experiment 2B also took part in this experiment; however, data from four participants were excluded due to

excessive missing data (one block of data failed to record). Data from 22 participants are presented (13 females; aged 21–32, average 24, all right-handed). All participants, first completed Experiment 3, followed by Experiment 2B.

##### 4.5.2. Stimuli

This experiment contained four stimulus conditions: R20-R10(M) and its control R20, and R10(M)-R20 and its control R10(M). Each were presented 100 times in random order. The R10(M)-R20 stimuli were constrained such that the first tone after the transition was novel (i. e., not contained in the preceding R10M sequence). This ensured that statistically, the transition occurred at the same time in all stimuli. The experiment was organized into four 10-minute blocks. Short breaks were provided between blocks, but participants were asked to remain still in the breaks.

#### 4.6. Experiment 4 A – Behavioural Pilot

##### 4.6.1. Participants

Five new participants (5 females; aged 21–23, average 22) participated in the pilot.

##### 4.6.2. Stimuli and procedure

The experiment contained 80 trials: 20 trials of R20-R5, 20 trials of R20-R10(M) and 40 trials of R20. All were randomly mixed. Otherwise, the procedure was identical to Experiment 1.

#### 4.7. Experiment 4 – EEG

##### 4.7.1. Participants

28 new participants (18 females, aged between 22 and 34, average 24.1, 2 left-handed) participated this experiment. Four participants were excluded, three due to excessive missing data (one block of data failed to record) and one due to faulty earphones. Data from 24 participants (15 females, aged 22–34, average 24.2; 1 left-handed) are presented.

##### 4.7.2. Stimuli

This experiment contained three conditions: R20-R5, R20-R10(M) and their common control R20. Transition times were as in Experiment 2, above. R5 sequences were created by randomly retaining 5 frequencies from the full pool – 1 from each band and an additional value randomly selected from across all bands. Overall, 600 stimuli (150 R20-R5, 150 R20-R10(M) and 300 of R20) were presented to the listeners in a random order. The experiment was organized into six 10-minute blocks. Short breaks were provided between blocks, but participants were asked to remain still in the breaks.

#### 4.8. Data analysis

##### 4.8.1. EEG data processing

The EEG analysis was performed using the Fieldtrip toolbox (Oostenveld et al., 2010) in MATLAB (2015b, The MathWorks, Inc., Natick, Massachusetts, United States).

Data were first smoothed by convolution with a square window of size 1/50 Hz to suppress 50 Hz and harmonics, and then low-passed at 30 Hz (two-pass, fifth-order Butterworth filter). Next, the data were downsampled to 256 Hz, and 5500 ms epochs (from 500 ms pre-onset to 1000 ms post-offset) were extracted for each trial. In Experiments 2 A, 3, and 4, the number of R20 trials was larger than that of the other conditions. To equate SNR between conditions, the number of trials contributing to the estimation of R20 activity was reduced to match the trial number of the other transition conditions by random selection. Data were then detrended using the NoiseTools toolbox (de Cheveigné and Arzounian, 2018) and average-referenced across the full set of channels.

De-noising source separation (DSS) analysis was applied to maximize reproducibility across trials (de Cheveigné and Parra, 2014; de

Cheveigné and Simon, 2008). In Experiments 2 and 4 DSS was applied on collapsed data (across conditions). In Experiment 3 DSS was run separately for each condition. The rationale was that the transition condition R10(M)-R20 might evoke a brain response (e.g., MMN) which might not share the components with other conditions.

For each participant, the first five DSS components (i.e., the five most reproducible components) were selected and projected back into sensor space. Trials which deviated from the mean by more than twice the standard deviation were automatically flagged as outliers (normally less than 10 %) and discarded from further analyses. Subsequently, the data were averaged across trials.

To identify additional (fast) activity, which may potentially be masked by the slow changes, the same analysis as described above was also performed on 2 Hz high-pass filtered data (two-pass, sixth-order Butterworth filter).

The root-mean-square (RMS) of the measured current at each time point was calculated over a subset of 10 channels (5 central, 5 occipital: FCz, Fz, Cz, FC1, FC2, P9, P10, Iz, O1 and O2) for each participant and condition and used as a measure of brain activation over time. These channels were selected because they were the most activated during the N1 onset response (around 100 ms after stimulus onset; see Fig. 4 A) in the grand average data (computed by collapsing across conditions) and thus considered to best capture auditory activity.

#### 4.8.2. EEG timeseries statistical analysis

For illustration purposes, group-RMS (RMS of individual RMSs) is shown, but statistical analysis is always conducted across participants. The difference between the (squared) RMS waveforms of each transition condition and R20 was calculated for each participant, and subjected to bootstrap re-sampling (5000 interactions; balanced) (Efron and Tibshirani, 1994). The difference was deemed significant if the proportion of bootstrap iterations that fell above or below zero was more than 95 % (i.e.,  $p < 0.05$ ) for 10 or more consecutive samples (i.e.,  $> 39$  ms). The significant intervals determined in this way are indicated as horizontal lines, below the EEG trace. The time at which the first significant sample is observed is considered the earliest time point at which the brain has detected that the statistics of the stimulus have changed.

The threshold for the number of consecutive samples was determined by comparing activity within the initial 2000 ms across conditions. Stimuli are identical during this period (R20); therefore, any significant differences are fully attributable to noise. We estimated the largest number of consecutive samples to show a significant effect in that interval by rerunning the analysis on the  $-500$  to  $2000$  ms interval 1000 times and taking the longest obtained “spurious” significant interval as the threshold for the statistical analysis.

#### 4.8.3. EEG channel-by-channel analysis (Experiment 4)

This analysis sought to assess the consistency of the EEG effect  $R5 > R10(M)$  across various EEG electrodes in Experiment 4. Initially, we identified the electrodes sensitive to transitions, where both R20-R5 and R20-R10(M) exhibited significant deviations from R20 within the 1–2 second post-transition timeframe. This selection process yielded 40 out of the 64 available electrodes.

For each of the ‘transition-sensitive’ electrodes, we conducted baseline correction on the absolute time series of R20-R5 and R20-R10 (M) using a 0.5-second pre-transition interval, followed by subtraction. The mean difference during the 1–2 second post-transition period was computed. A higher positive value signifies a more pronounced response to R5 compared to R10(M). These values are depicted in Fig. 8 with colour coding.

In Fig. 8, the majority of the transition-sensitive electrodes display large positive values. Notably, none of the electrodes exhibited negative values, indicating that none of the channels demonstrated a response pattern where  $R5 < R10(M)$ .

#### 4.8.4. Behavioural analysis

The dependent measure was  $d'$  (Macmillan, 1991), computed as  $d' = z(\text{Hit rate}) - z(\text{False Alarm rate})$ . For the purposes of computing the  $d'$  score, transition trials in which the participant responded after the nominal transition time were considered as hits. Responses to no change trials, or before the onset of the transition were labelled as false alarms.

Statistical analysis was performed using SPSS. In all analyses of variance (ANOVA), the Greenhouse-Geisser correction was applied when the assumption of sphericity was violated (Greenhouse and Geisser, 1959). The  $\alpha$  level was set a priori to 0.05.

#### 4.9. Computational modelling: D-REX model

An auditory perceptual model - Dynamic Regularity Extraction (D-REX) - was used to capture and predict the unfolding statistics within sequential experimental stimuli (Skerritt-Davis and Elhilali, 2018). The model is an extension of a general Bayesian Online Changepoint Detection model (Adams and MacKay, 2007) which uses Bayesian inference to perform sequential predictions given previous observations in the presence of unknown changepoints.

The model operates on a basic assumption that the generative model underlying all sensory observations is a Gaussian distribution estimated by parameters  $\theta$  that changes at unknown times (i.e., changepoints), where any observations after a change are independent of observations before the change. The Bayesian model attempts to predict the next observation despite the unknown changepoint by estimating predictive distributions over many context windows of previous observations and inferring the time since the last change (see Fig. 2).

Estimation of model parameters is performed online. At time  $t$ , context-specific predictive distributions  $p_{i,t} = P(x_{t+1}|c_i) = P(x_{t+1}|\hat{\theta}_{i,t})$  are estimated by sufficient statistics  $\hat{\theta}_{i,t}$  based on each context ( $c_i$ ) of previous observations since  $t = c_i$ :  $\hat{\theta}_{i,t} = \{\hat{\mu}_{i,t}, \hat{\sigma}_{i,t}\}$ , where  $\hat{\mu}_{i,t}$ ,  $\hat{\sigma}_{i,t}$  are the sample mean and sample variance, respectively, estimated from observations  $x_{c_i:t} = \{x_{c_i}, x_{c_i+1}, \dots, x_t\}$  (Murphy, 2007). Under the original Bayesian changepoint detection framework, all previous observations are used to infer the belief that the last changepoint occurred at  $c_i$ :  $b_{i,t} = P(c_i|x_{1:t})$ . The predictive distribution of the next observation then marginalizes out the unknown changepoint:  $P(x_{t+1}|x_{1:t}) = \sum_i p_{i,t} b_{i,t}$ .

Two perceptual parameters are also used in this model: (1) memory ( $m$ ), which limits the context  $c_i$  used to estimate sufficient statistics for generating predictions; and (2) observation noise ( $n$ ), which represents limitations on perceptual fidelity by adding independent Gaussian noise to the variance of the predictive distribution. By fitting these parameters to each participant’s data, we have accounted for individual differences in auditory perception, including those related to pitch perception and loudness preferences. Hence, the predictive distribution at time  $t$  estimated from the context of previous observations since  $c_i$  is:

$$\hat{\theta}_t^{(c_i)} = \begin{cases} \left\{ \hat{\mu}_{i,t}, \tilde{\sigma}_{i,t} \right\}, & c_i > t - m \\ \left\{ \hat{\mu}_{m,t}, \tilde{\sigma}_{m,t} \right\}, & c_i \leq m \end{cases}$$

where  $\tilde{\sigma}_{i,t} = \hat{\sigma}_{i,t} + n^2$  is the sample covariance with added observation noise  $n$ . The predictive distribution is then estimated by at most the last  $m$  observations. An extended version of the D-REX model was used wherein the sufficient statistics underlying the predictive probability track changes in a multivariate Gaussian distribution in line with previous work (Skerritt-Davis and Elhilali, 2021b). The model performs online estimation of the parameters of a multivariate Gaussian distribution from sequential observations (Engel and Heinen, 2010).

To derive behavioural responses from the model that are akin to the human listeners’ change detection task, we define the probability that any change has occurred in the observed sequence before time  $t$  as:  $P(c_i > 1|x_{1:t}) = \sum_{i:c_i > 1} b_{i,t}$ . In other words, the probability that the last

change  $c_i$  occurred after the beginning of the sequence  $t = 1$ . A threshold ( $\tau$ ) applied to this change probability over time is used to define the change response from the model (see Fig. 2b).

Two versions of the model were compared: 1) A wide-channel model which uses a single Gaussian distribution as the starting prior; and 2) a multi-channel model which uses a multimodal Gaussian prior. In the latter case, the prior distribution uses 5 channels spread evenly in log-frequency (between  $-16$  and  $16$  semitones relative to the centre frequency of the stimuli). The choice of prior distributions was based on a maximum entropy prior for the stimuli employed, where the wide-channel prior was based on the statistics of the R20 stimulus and the multi-channel model contained components spanning the entire frequency range. The choice of five channels in the multi-channel prior was motivated by the intention to mitigate any potential bias in the results towards the frequency bands employed in stimulus generation.

The model is implemented in MATLAB and its code is available online <https://github.com/JHU-LCAP/DREX-model>. The model implementation is overall similar to that described in detail in a previous study (Skerritt-Davis and Elhilali, 2021b) with the following modifications: First, all observations are assumed to be independent (no temporal covariance). Second, the hazard rate, which sets the prior probability of change occurrence, is set to 0.01, essentially an “ideal” value approximately matching the true occurrence of changes in the stimulus statistics. The hazard rate, while likely a learned parameter by the listener reflecting the expected volatility of the stimulus, does not have a qualitative impact on the model responses when it is set within a reasonable range of the true value for the stimuli employed, and therefore was removed from any parameter fitting in order to reduce dimensionality in the search space. For modelling of behavioural experiments, D-REX model parameters ( $m, n, \tau$ ) were individually matched to each subject by performing a grid search and comparing model responses to subjects’ responses to the same set of stimuli. The model detection rate (i.e., percentage of trials wherein a change was detected) within each stimulus condition (both false positive and true positive detections) was collected for each iteration in the search procedure. We selected the parameters resulting in the least mean squared error in average detection rate across conditions between model and listener behaviour. For EEG experiments, a single ideal observer model was used with infinite memory and zero noise.

### CRedit authorship contribution statement

**Frederic Dick:** Writing – review & editing, Writing – original draft, Supervision, Resources, Methodology, Investigation, Conceptualization. **Mounya Elhilali:** Writing – review & editing, Supervision, Software, Methodology, Investigation, Formal analysis. **Benjamin Skerritt-Davis:** Writing – review & editing, Visualization, Validation, Software, Methodology, Formal analysis. **Sijia Zhao:** Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Maria Chait:** Writing – review & editing, Writing – original draft, Supervision, Resources, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

### Declaration of Competing Interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Acknowledgements

This work was supported by a BBSRC project grant (BB/P003745/1) to MC and Office of Naval Research (ONR) grant (N00014-23-1-2050) to ME. We are grateful to Daniel Yon and Clare Press for comments and discussion.

## Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at [doi:10.1016/j.pneurobio.2024.102696](https://doi.org/10.1016/j.pneurobio.2024.102696).

### Data availability

Stimuli examples, anonymised data as well as the analysis code to reproduce the figures is available at [<https://doi.org/10.17605/OSF.IO/4DKN3>].

### References

- R.P. Adams, D.J.C. MacKay, Bayesian Online Changepoint Detection. *arXiv:0710.3742 [stat]* (2007).
- Andrillon, T., Kouider, S., Agus, T., Pressnitzer, D., 2015. Perceptual learning of acoustic noise generates memory-evoked potentials. *Curr. Biol.* 25, 2823–2829.
- Asokan, M.M., Williamson, R.S., Hancock, K.E., Polley, D.B., 2021. Inverted central auditory hierarchies for encoding local intervals and global temporal patterns. *Curr. Biol.* 31, 1762–1770.e4.
- Auksztulewicz, R., Barascud, N., Cooray, G., Nobre, A.C., Chait, M., Friston, K., 2017. The cumulative effects of predictability on synaptic gain in the auditory processing stream, 0291–17 *J. Neurosci.* 0291–17.
- Barascud, N., Pearce, M.T., Griffiths, T.D., Friston, K.J., Chait, M., 2016. Brain responses in humans reveal ideal observer-like sensitivity to complex acoustic patterns. *Proc. Natl. Acad. Sci.* 113, E616–E625.
- Barczak, A., O’Connell, M.N., McGinnis, T., Ross, D., Mowery, T., Falchier, A., Lakatos, P., 2018. Top-down, contextual entrainment of neuronal oscillations in the auditory thalamocortical circuit. *PNAS* 115, E7605–E7614.
- Behrens, T.E.J., Woolrich, M.W., Walton, M.E., Rushworth, M.F.S., 2007. Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221.
- Beierholm, U., Rohe, T., Ferrari, A., Stegle, O., Noppeney, U., 2020. Using the past to estimate sensory uncertainty. *eLife* 9, e54172.
- Bitterman, Y., Mukamel, R., Malach, R., Fried, I., Nelken, I., 2008. Ultra-fine frequency tuning revealed in single neurons of human auditory cortex. *Nature* 451, 197–201.
- Bland, A.R., Schaefer, A., 2012. Different varieties of uncertainty in human decision-making. *Front. Neurosci.* 6.
- Boubenec, Y., Lawlor, J., Górska, U., Shamma, S., Englitz, B., 2017. Detecting changes in dynamic and complex acoustic environments. *eLife* 6, e24910.
- Brainard, D.H., 1997. The psychophysics toolbox. *Spat. Vis.* 10, 433–436.
- de Cheveigné, A., Arzoumanian, D., 2018. Robust detrending, rereferencing, outlier detection, and inpainting for multichannel data. *NeuroImage* 172, 903–912.
- de Cheveigné, A., Parra, L.C., 2014. Joint decorrelation, a versatile tool for multichannel data analysis. *NeuroImage* 98, 487–505.
- de Cheveigné, A., Simon, J.Z., 2008. Denoising based on spatial filtering. *J. Neurosci. Methods* 171, 331–339.
- Chi, T., Ru, P., Shamma, S.A., 2005. Multiresolution spectrotemporal analysis of complex sounds. *J. Acoust. Soc. Am.* 118, 887–906.
- Daikhin, L., Ahissar, M., 2012. Responses to deviants are modulated by subthreshold variability of the standard. *Psychophysiology* 49, 31–42.
- Dalley, J.W., McGaughy, J., O’Connell, M.T., Cardinal, R.N., Levita, L., Robbins, T.W., 2001. Distinct changes in cortical acetylcholine and noradrenaline efflux during contingent and noncontingent performance of a visual attentional task. *J. Neurosci.* 21, 4908–4914.
- Daunizeau, J., den Ouden, H.E.M., Pessiglione, M., Kiebel, S.J., Stephan, K.E., Friston, K.J., 2010. Observing the observer (I): meta-bayesian models of learning and decision-making. *PLOS ONE* 5, e15554.
- Daw, N.D., O’Doherty, J.P., Dayan, P., Seymour, B., Dolan, R.J., 2006. Cortical substrates for exploratory decisions in humans. *Nature* 441, 876–879.
- Demarchi, G., Sanchez, G., Weisz, N., 2019. Automatic and feature-specific prediction-related neural activity in the human auditory system. *Nat. Commun.* 10.
- Diederen, K.M.J., Spencer, T., Vestergaard, M.D., Fletcher, P.C., Schultz, W., 2016. Adaptive prediction error coding in the human midbrain and striatum facilitates behavioral adaptation and learning efficiency. *Neuron* 90, 1127–1138.
- Efron, B., Tibshirani, R.J., 1994. *An Introduction to the Bootstrap*. Chapman and Hall/CRC, New York.
- Engel, P.M., Heinen, M.R., 2010. Incremental Learning of Multivariate Gaussian Mixture Models. In: da Rocha Costa, A.C., Vicari, R.M., Tonidandel, F. (Eds.), *Advances in Artificial Intelligence – SBIA 2010*. Springer, Berlin, Heidelberg, pp. 82–91.
- Feldman, H., Friston, K.J., 2010. Attention, uncertainty, and free-energy. *Front. Hum. Neurosci.* 4.
- Friston, K.J., 2017. Precision psychiatry. *Biol. Psychiatry: Cogn. Neurosci. Neuroimaging* 2, 640–643.
- Furl, N., Kumar, S., Alter, K., Durrant, S., Shawe-Taylor, J., Griffiths, T.D., 2011. Neural prediction of higher-order auditory sequence statistics. *NeuroImage* 54, 2267–2277.
- Garrido, M.I., Dolan, R.J., Sahani, M., 2011. Surprise leads to noisier perceptual decisions. *i-Perception* 2, 112–120.
- Garrido, M.I., Sahani, M., Dolan, R.J., 2013. Outlier responses reflect sensitivity to statistical structure in the human brain. *PLoS Comput. Biol.* 9, e1002999.
- Glaze, C.M., Kable, J.W., Gold, J.I., 2015. Normative evidence accumulation in unpredictable environments. *eLife* 4, e08825.

- Gold, J.I., Stocker, A.A., 2017. Visual decision-making in an uncertain and dynamic world. *Annu. Rev. Vis. Sci.* 3, 227–250.
- Greenhouse, S.W., Geisser, S., 1959. On methods in the analysis of profile data. *Psychometrika* 24, 95–112.
- Harrison, A.W., Mannion, D.J., Jack, B.N., Griffiths, O., Hughes, G., Whitford, T.J., 2021. Sensory attenuation is modulated by the contrasting effects of predictability and control. *NeuroImage* 237, 118103.
- Heilbron, M., Chait, M., 2018. Great expectations: is there evidence for predictive coding in auditory cortex? *Neuroscience* 389, 54–73.
- Herrmann, B., Johnsrude, I.S., 2018. Neural signatures of the processing of temporal patterns in sound. *J. Neurosci.* 38, 5466–5477.
- Herrmann, B., Araz, K., Johnsrude, I.S., 2021. Sustained neural activity correlates with rapid perceptual learning of auditory patterns. *NeuroImage* 238, 118238.
- Hsu, Y.-F., Hämäläinen, J.A., 2021. Both contextual regularity and selective attention affect the reduction of precision-weighted prediction errors but in distinct manners. *Psychophysiology* 58, e13753.
- Hsu, Y.-F., Waszak, F., Hämäläinen, J.A., 2019. Prior precision modulates the minimization of auditory prediction error. *Front. Hum. Neurosci.* 13.
- Kanai, R., Komura, Y., Shipp, S., Friston, K., 2015. Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philos. Trans. R. Soc. B* 370, 20140169.
- Kelly, S.P., O'Connell, R.G., 2013. Internal and external influences on the rate of sensory evidence accumulation in the human brain. *J. Neurosci.* 33, 19434–19441.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., Broussard, C., 2007. What's new in psychtoolbox-3. *Perception* 36, 1–16.
- Knill, D.C., Pouget, A., 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719.
- Krishnamurthy, K., Nassar, M.R., Sarode, S., Gold, J.I., 2017. Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nat. Hum. Behav.* 1, 0107.
- Lawson, R.P., Bisby, J., Nord, C.L., Burgess, N., Rees, G., 2021. The computational, pharmacological, and physiological determinants of sensory learning under uncertainty. *Curr. Biol.* 31, 163–172.e4.
- Lecaignard, F., Bertrand, O., Caclin, A., Mattout, J., 2022. Neurocomputational underpinnings of expected surprise. *J. Neurosci.* 42, 474–486.
- Macmillan, N.A., 1991. *Detection Theory: A User's Guide*. CUP Archive.
- Megela, A.L., Teyler, T.J., 1979. Habituation and the human evoked potential. *J. Comp. Physiol. Psychol.* 93, 1154–1170.
- Milne, A., Zhao, S., Tampakaki, C., Bury, G., Chait, M., 2021. Sustained pupil responses are modulated by predictability of auditory sequences. *J. Neurosci.* 41, 6116–6127.
- Murphy, K.P., 2007. *Conjugate Bayesian analysis of the Gaussian distribution*. 1, 29 (2007).
- Nassar, M.R., Wilson, R.C., Heasly, B., Gold, J.I., 2010. An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* 30, 12366–12378.
- Nassar, M.R., Rumsey, K.M., Wilson, R.C., Parikh, K., Heasly, B., Gold, J.I., 2012. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* 15, 1040–1046.
- Natan, R.G., Briguglio, J.J., Mwilambwe-Tshilobo, L., Jones, S.I., Aizenberg, M., Goldberg, E.M., Geffen, M.N., 2015. Complementary control of sensory adaptation by two types of cortical interneurons. *eLife Sci.* 4, e09868.
- Natan, R.G., Rao, W., Geffen, M.N., 2017. Cortical interneurons differentially shape frequency tuning following adaptation. *Cell Rep.* 21, 878–890.
- Norton, E.H., Acerbi, L., Ma, W.J., Landy, M.S., 2019. Human online adaptation to changes in prior probability. *PLoS Comput. Biol.* 15, e1006681.
- O'Connell, R.G., Dockree, P.M., Kelly, S.P., 2012. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat. Neurosci.* 15, 1729–1735.
- O'Reilly, J.X., 2013. Making predictions in a changing world—inference, uncertainty, and learning. *Front. Neurosci.* 7.
- Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.-M., 2010. FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* 2011, 156869.
- Overath, T., Cusack, R., Kumar, S., von Kriegstein, K., Warren, J.D., Grube, M., Carlyon, R.P., Griffiths, T.D., 2007. An information theoretic characterisation of auditory encoding. *PLoS Biol.* 5, e288.
- Payzan-LeNestour, E., Bossaerts, P., 2011. Risk, unexpected uncertainty, and estimation uncertainty: bayesian learning in unstable settings. *PLoS Comput. Biol.* 7, e1001048.
- Pérez-González, D., Malmierca, M.S., 2014. Adaptation in the auditory system: an overview. *Front Integr. Neurosci.* 8.
- Pérez-González, D., Lao-Rodríguez, A.B., Aedo-Sánchez, C., Malmierca, M.S., 2024. Acetylcholine modulates the precision of prediction error in the auditory cortex. *eLife* 12, RP91475.
- Richter, L.M.A., Gjorgjieva, J., 2022. A circuit mechanism for independent modulation of excitatory and inhibitory firing rates after sensory deprivation. *Proc. Natl. Acad. Sci.* 119, e2116895119.
- Rubin, J., Ulanovsky, N., Nelken, I., Tishby, N., 2016. The representation of prediction error in auditory cortex. *PLoS Comput. Biol.* 12, e1005058.
- SanMiguel, I., Costa-Faidella, J., Lugo, Z.R., Vilella, E., Escera, C., 2021. Standard tone stability as a manipulation of precision in the oddball paradigm: modulation of prediction error responses to fixed-probability deviants. *Front. Hum. Neurosci.* 15.
- Santoro, R., Moerel, M., Martino, F.D., Goebel, R., Ugurbil, K., Yacoub, E., Formisano, E., 2014. Encoding of natural sounds at multiple spectral and temporal resolutions in the human auditory cortex. *PLoS Comput. Biol.* 10, e1003412.
- Schulz, A., Miehl, C., Berry II, M.J., Gjorgjieva, J., 2021. The generation of cortical novelty responses through inhibitory plasticity. *eLife* 10, e65309.
- W. Sedley, P.E. Gander, S. Kumar, C.K. Kovach, H. Oya, H. Kawasaki, M.A. Howard, T.D. Griffiths, *Neural signatures of perceptual inference*. *eLife* 5, e11476.
- Skerritt-Davis, B., Elhilali, M., 2018. Detecting change in stochastic sound sequences. *PLoS Comput. Biol.* 14, e1006162.
- Skerritt-Davis, B., Elhilali, M., 2021a. Neural encoding of auditory statistics. *J. Neurosci.* 41, 6726–6739.
- Skerritt-Davis, B., Elhilali, M., 2021b. Computational framework for investigating predictive processing in auditory perception. *J. Neurosci. Methods* 360, 109177.
- Sohoglu, E., Chait, M., 2016. Detecting and representing predictable structure during auditory scene analysis. *eLife* 5, e19113.
- Soltani, A., Izquierdo, A., 2019. Adaptive learning under expected and unexpected uncertainty. *Nat. Rev. Neurosci.* 20, 635–644.
- Southwell, R., Chait, M., 2018. Enhanced deviant responses in patterned relative to random sound sequences. *Cortex* 109, 92–103.
- Southwell, R., Baumann, A., Gal, C., Barascud, N., Friston, K., Chait, M., 2017. Is predictability salient? A study of attentional capture by auditory patterns. *Philos. Trans. R. Soc. B* 372, 20160105.
- E. Spaak, *The “Bayesian” brain, with a bit less Bayes*. arXiv arXiv:2111.09063 [Preprint] (2024). (<https://doi.org/10.48550/arXiv.2111.09063>).
- Tenenbaum, J.B., Griffiths, T.L., Kemp, C., 2006. Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* 10, 309–318.
- Warren, B.M., Ackroff, J.M., 1976. Two types of auditory sequence perception. *Percept. Psychophys.* 20, 387–394.
- Warren, R.M., 2008. *Auditory Perception: An Analysis and Synthesis*. Cambridge University Press, Cambridge. (<http://ebooks.cambridge.org/ref/id/CBO9780511754777>).
- Warren, R.M., Obusek, C.J., 1972. Identification of temporal order within auditory sequences. *Percept. Psychophys.* 12, 86–90.
- Weissbart, H., Kandyilaki, K.D., Reichenbach, T., 2020. Cortical tracking of surprisal during continuous speech comprehension. *J. Cogn. Neurosci.* 32, 155–166.
- Wilson, R., Niv, Y., 2012. Inferring relevance in a changing world. *Front. Hum. Neurosci.* 5.
- Wilson, R.C., Nassar, M.R., Gold, J.I., 2013. A mixture of delta-rules approximation to bayesian inference in change-point problems. *PLoS Comput. Biol.* 9, e1003150.
- Winkler, I., Paavilainen, P., Alho, K., Reinikainen, K., Sams, M., Naatanen, R., 1990. The effect of small variation of the frequent auditory stimulus on the event-related brain potential to the infrequent stimulus. *Psychophysiology* 27, 228–235.
- Yao, J.D., Gimoto, J., Constantinople, C.M., Sanes, D.H., 2020. Parietal cortex is required for the integration of acoustic evidence. *Curr. Biol.* 30, 3293–3303.e4.
- Yarden, T.S., Mizrahi, A., Nelken, I., 2022. Context-dependent inhibitory control of stimulus-specific adaptation. *J. Neurosci.* 42, 4629–4651.
- Yon, D., Frith, C.D., 2021. Precision and the Bayesian brain. *Curr. Biol.* 31, R1026–R1032.
- Yu, A.J., Dayan, P., 2005. Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692.