

Surprise leads to noisier perceptual decisions

Marta I Garrido

Wellcome Trust Centre for Neuroimaging, University College London, London WC1N 3BG, England;
e-mail: m.garrido@fil.ion.ucl.ac.uk

Raymond J Dolan

Wellcome Trust Centre for Neuroimaging, University College London, London WC1N 3BG, England;
e-mail: r.dolan@ucl.ac.uk

Maneesh Sahani

Gatsby Computational Neuroscience Unit, University College London, London WC1N 3AR, England;
e-mail: maneesh@gatsby.ucl.ac.uk

Received 8 October 2010, in revised form 8 March 2011; published online 15 May 2011

Abstract. Surprising events in the environment can impair task performance. This might be due to complete distraction, leading to lapses during which performance is reduced to guessing. Alternatively, unpredictability might cause a graded withdrawal of perceptual resources from the task at hand and thereby reduce sensitivity. Here we attempt to distinguish between these two mechanisms. Listeners performed a novel auditory pitch–duration discrimination, where stimulus loudness changed occasionally and incidentally to the task. Responses were slower and less accurate in the surprising condition, where loudness changed unpredictably, than in the predictable condition, where the loudness was held constant. By explicitly modelling both lapses and changes in sensitivity, we found that unpredictable changes diminished sensitivity but did not increase the rate of lapses. These findings suggest that background environmental uncertainty can disrupt goal-directed behaviour. This graded processing strategy might be adaptive in potentially threatening contexts, and reflect a flexible system for automatic allocation of perceptual resources.

Keywords: prediction, change detection, uncertainty, attention, sensitivity, oddballs.

1 Introduction

The ability to detect changes or novelty in the environment is crucial for adaptive behaviour. Combined behavioural, neurophysiological, and neuroimaging research suggests that surprising stimuli evoke distinct patterns of neural activity, and are often associated with behavioural interference (Dalton and Lavie 2004; Schroeger and Wolff 1998; Watkins et al 2007). Surprising stimuli can be difficult to ignore. In visual search tasks, nontargets with unique features will slow down responses to targets, even if the dimensions of distractor uniqueness are irrelevant to the task (Theeuwes 1992). A similar phenomenon has been shown in hearing, where the presence of unexpected acoustic features retards and impairs task execution (Schroeger and Wolff 1998), even if listeners are explicitly instructed to ignore distractors (Dalton and Lavie 2004).

Subsequent studies have elaborated the impact of auditory distraction on behavioural performance (for an example see Parmentier and Andres 2010). Whilst many studies have documented perceptual errors associated with unexpected stimuli, or changes in stimulus features (Horvath and Winkler 2010; Schroeger and Wolff 1998; Sussman et al 2003), these studies have not attempted to disentangle the potentially separate contributions made to this effect by changes in sensitivity and in lapse rate. These two possible mechanisms of interference have different behavioural consequences. On some occasions the distracting stimulus, or feature, may capture the resources needed for sensory or perceptual processing completely, thus leading to a complete lapse in the processing of other stimuli. At such times a listener would be reduced to making guesses about the task-relevant stimulus irrespective of its true value, leading to an increase in stimulus-independent behavioural noise. Alternatively, an unexpected change may cause only a partial withdrawal of perceptual processing that

leads to reduced sensitivity at the task. In this case, listeners may still respond accurately to easy tasks, and therefore this effect is apparent as stimulus-dependent noise. Importantly, a failure to take into account the effects of guessing when modelling the data can lead to biased sensitivity estimates (Treutwein and Strasburger 1999; Wichmann and Hill 2001).

In this study we ask whether the presence of task-irrelevant surprising events can significantly shape perceptual decisions, in terms of not only their speed but also their sensitivity and lapse rates. Using signal detection theory, we specifically address whether surprising events that lead to poorer performance in a perceptual discrimination task are caused by increases in either stimulus-dependant or stimulus-independent sensory noise, or both. By separately modelling these two phenomena, we are able to tease them apart. Moreover, our design did not require listeners to rely on an internal representation or memory reference of common events to complete the task, as they do in oddball detection tasks (Dalton and Lavie 2004, 2007; Schroeger and Wolff 1998; Watkins et al 2007). Therefore, our design allowed us to isolate the effects of distraction on sensory or perceptual processing, rather than on memory or retrieval processes.

2 Methods

2.1 Participants

Twenty healthy volunteers (five male, fifteen female; 18–34 years; mean age = 23.5 years; standard deviation = 4.5 years) participated in the experiment. All reported normal hearing and two of them had substantial musical training. Each participant provided signed informed consent before the study, which proceeded under local ethical committee guidelines. Psychometric data of four subjects were excluded from the population statistics due to poor sampling, leaving sixteen subjects (four male, twelve female; 18–31 years; mean age = 23.4 years; standard deviation = 4 years). All participants were compensated for their time at the end of the experiment.

2.2 Perceptual discrimination task

Participants sat comfortably in front of a computer screen in a dimly illuminated sound-proofed room while performing a simultaneous two-alternative forced-choice (AFC) perceptual discrimination task. A sequence of pairs of broadband noisy sounds was played at a sampling frequency of 44 100 Hz and presented diotically. Of each pair, one sound was shaped to elicit the percept of a high pitch (400 Hz) and the other a low pitch (128 Hz). Each sound token was an instance of iterated rippled noise (IRN), generated by iteratively combining a noise stimulus with the same signal delayed by $1/f$ s, to achieve a pitch percept of f Hz (Patterson et al 2000; Yost 1996). The number of iterations, which controls pitch strength or saliency, was kept constant (to thirty-two iterations) for both high and low pitched sounds. IRN was chosen instead of pure tones so that both sounds spanned the same frequencies, allowing equal intensity sounds to have equal loudness, and minimising differential stimulus-specific adaptation effects (Ulanovsky et al 2003). Onsets were simultaneous, and listeners were asked to report by means of button press on a computer keyboard which of the sounds (the high or the low pitch) lasted the longer, or equivalently whether a composite sound ended with a high or low pitch tail. The pairs of sounds were played every 2 s.

Stimuli were blocked by the predictability of their (task-irrelevant) loudness. In a surprise or oddball block, intensity was mostly constant at 70 dB SPL (with probability ~ 0.8), but occasional pairs were either louder (80 dB SPL; probability ~ 0.1) or softer (60 dB SPL; probability ~ 0.1) than this median standard loudness. In each predictable or control block, all pairs had the same intensity, selected from the three different levels used in the surprising condition. Thus, there were six different types of trials: surprising-block trials with median (S_M), louder (S_L), and softer (S_S) intensities; and predictable or control-block trials for median

(P_M), louder (P_L), and softer (P_S) blocks (see figure 1 for a schematic of the experimental design). All participants performed a practice session lasting 3–6 minutes (90–180 trials) in which they became familiar with the stimuli and the task. In the practice session the duration of these sounds varied between 100 ms and 1100 ms. On each trial the difference in duration (the ‘tail’) between high and low pitch sounds could be one of eleven pseudorandomised and equally spaced values between –1000 ms and 1000 ms. The negative values correspond to trials where the longer sound had the higher pitch. Pairs with longer high pitch sounds and those with longer low pitch sounds were equally probable.

The experiment proper had six sessions in total: five oddball blocks of about thirteen minutes each (the surprise condition; 1995 trials), and one control block of about ten minutes (the predictable condition; 495 trials: three subblocks with 165 trials at each of the three loudness levels). In these sessions, the durations of the pairs of sounds varied in the range of 100–1276 ms and were adjusted in order to obtain a good sampling of each participant’s psychometric function. This design is not dissimilar to the oddball detection paradigm described by Schroeger and Wolff (1998), in which a change occurs at a particular stimulus dimension (in this case loudness), incidental to the task-relevant dimension (relative duration).

The stimuli and task were written in MATLAB, using the Cogent 2000 toolbox (<http://www.vislab.ucl.ac.uk/cogent.php>).

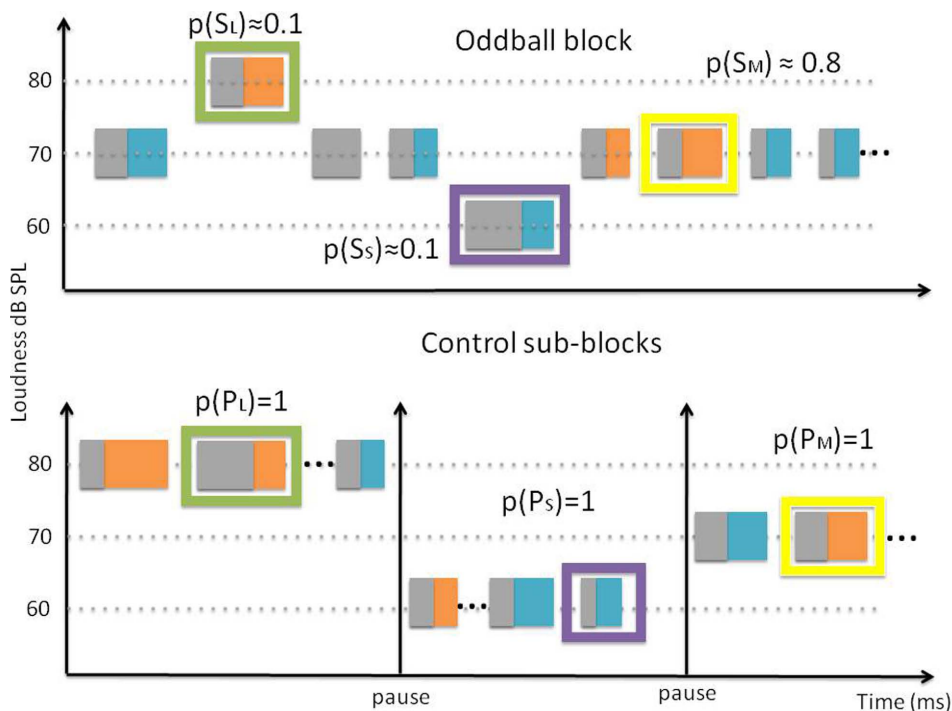


Figure 1. Experimental design and task schematic. Participants performed a two-alternative forced-choice perceptual discrimination task where they had to report which of two, high or low pitch, sounds lasted the longer, in every trial. In oddball blocks (surprise condition) both sounds occasionally changed their loudness. In control blocks (predictable condition) the loudness was kept constant. Blue and orange rectangles correspond to high and low pitch sounds played alone, and grey to the superposition of the two.

2.3 Sampling procedure

The choice of sound durations in the main experiment was based on the sensitivities, or slopes (ρ), estimated in the practice session, to ensure good sampling of each subject-specific

psychometric function. The goal was to ensure a number of points in the central high-slope region of the function sufficient to estimate that slope accurately, while also sampling the flanking regions of asymptotic performance where lapses are more evident. Initially, we used equally spaced observations (differences in sound durations, or tail lengths) that fell within a distance of $1.57/\rho$ on either side of μ , the point of subjective equality (Levitt 1971). This procedure was adopted for nine subjects. However, for four of these subjects this sampling procedure was not optimal, and fewer than three points fell within the central region. Very likely, this reflected continued learning during the main sessions, which improved sensitivity by as much as a factor of 2. We excluded these four subjects from our analysis of choice data. As a result of this observation, we placed the samples for the next subject at $1/2\rho$ [0, 0.25, 0.5, 1, 2, 4] (with ρ derived from the practice sessions, as before) and then further refined the sample placement to $1/2\rho$ [0, 0.125, 0.25, 0.5, 1, 3] for the remaining ten subjects. In this way we maximised the number of observations on the high-slope section of a subject's specific psychometric function, while ensuring at least one data point in both extremes. We used reaction time (RT) data from all twenty subjects, and all choice data from the sixteen subjects for whom we had at least three measurements in the central region. Each of the eleven points of the psychometric curve was sampled fifteen times on average.

2.4 Psychophysics modeling and fitting scheme

Our model assumes that the listener's estimate of the duration tail, τ , is normally distributed around the true duration tail, t , with variance σ^2 , such that: $p(\tau/t) = N(\tau; t, \sigma^2)$ (Macmillan and Creelman 2005). This corresponds to a general psychophysical treatment of sensory noise. We found that each listener's psychometric curve was well fitted by a cumulative normal psychometric function for binomially distributed response counts which also models lapses (Whiteley and Sahani 2008; Wichmann and Hill 2001). This allowed us to separately model two different sources of noise: (1) stimulus-centred sensory noise expressed as random additive perturbations to the listener's estimate of the duration tail, and (2) a stimulus-independent noise (independent of t) that results from motor errors or complete lapses in attention. While complete lapses lead to guessing, and reduce task performance to chance level, irrespective of the stimulus, stimulus-centred noise reflects a more graded form of distraction that partially withdraws perceptual resources from the task at hand, and thereby reduces perceptual sensitivity.

We used eleven different true offsets, t_i (we use the subscript i to indicate offset in the equations below), in six different conditions (S_M, S_L, S_S, P_M, P_L, P_S) (indexed by j), with N_{ij} trials in each condition. The number of trials n_{ij} in which observers answer 'low pitch' for duration tail t_i in the j th condition is assumed to be drawn from a binomial distribution,

$$P(n_{ij}) = \binom{N_{ij}}{n_{ij}} p_{ij}^{n_{ij}} (1 - p_{ij})^{N_{ij} - n_{ij}}, \quad (1)$$

where p_{ij} , the modelled probability of answering low pitch, depends on t_i according to a modified cumulative Gaussian,

$$p_{ij} = (1 - \lambda_j) \left\{ \frac{1 + \operatorname{erf}[\sqrt{\pi}\rho_j(t_i - \mu_j)]}{2} \right\} + \frac{1}{2}\lambda_j \quad (2)$$

and $\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-y^2} dy$ is the standard error function.

The contribution of stimulus-dependent sensory noise is accounted for by the first term in equation (2) and is assumed to have a Gaussian shape centred around μ_j with width σ_j . Therefore, the estimate of the slope (ρ_j) of the cumulative Gaussian provides an estimate of

the variance through the relation $\hat{\sigma}_j = 1/\sqrt{2\pi}\hat{p}_j$. Stimulus-independent noise is modelled by the second term in equation (2), where λ_j corresponds to the probability of a lapse, which, for a binary decision, leads to a guessing probability of 0.5. We fit this model separately to the data observed under each condition and estimated parameters for the centre or bias μ_j , slope ρ_j , and lapse rate λ_j , for each curve. The fitting algorithm was based on a conjugate gradient maximum-likelihood method, minimize.m (Rasmussen 2001). Initial conditions for the bias parameter were based on a grid search, ranging from -50 ms to 200 ms in steps of 1 ms. Slopes were initialised on the slopes obtained in the practice session, and the lapse rate was initialised at 0.01 .

3 Results

3.1 RT is slower in surprising than in predictable contexts

RT data from twenty subjects were extracted for each of the six conditions ($S_M, S_L, S_S, P_M, P_L, P_S$) and analysed using an ANOVA with two factors: context and loudness. The context could be either surprising (S = surprising or oddball block) or predictable (P = predictable or control block), while the loudness factor had three levels: median, louder, and softer. We found that responses to louder sounds were significantly slower ($p = 0.004$), indicating a degree of perceptual interaction. Critically for our purposes, however, we found a significant main effect of context revealing slower responses to events in the surprising context at the same loudness level ($p = 0.023$) (see figure 2).

For analysis of both RT and choice data (described below), we excluded the first trial of the surprising blocks and subblocks, as well as all standard trials that immediately followed an oddball event. This was implemented because a local change, such as a standard sound after a deviant, could be perceived as a local deviant. Data from the five oddball sessions were then concatenated and treated as a single oddball or surprising block.

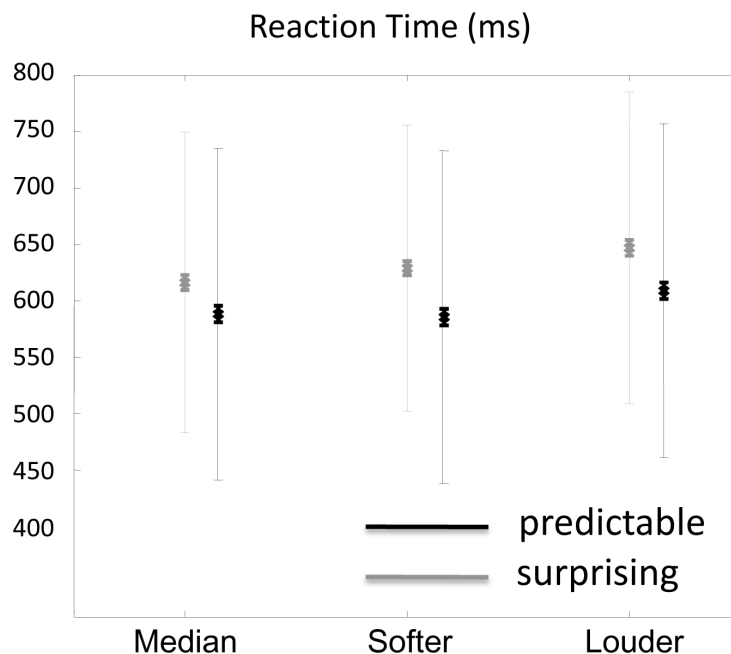


Figure 2. Surprised-induced noise slows down sensory decisions. Reaction time data for surprising and predictable conditions in median, softer, and louder sound levels. The thick segments indicate standard errors and the lines correspond to standard deviation.

3.2 Perceptual uncertainty in surprising and predictable contexts

We used signal detection theory to compare how listeners' sensitivity to duration differences in sounds of equal loudness differed between a surprising and a predictable condition, or context. As described above, we fitted separate psychometric curves for the data from each of the 3×2 conditions, and obtained specific bias μ_j , slope ρ_j , and lapse rate λ_j parameters for each condition indexed by j . Figure 3 shows example data from one listener in the six different conditions. Crosses correspond to observed data points and smooth lines show psychometric functions fit to data in the surprise (grey) and predictable (black) contexts.

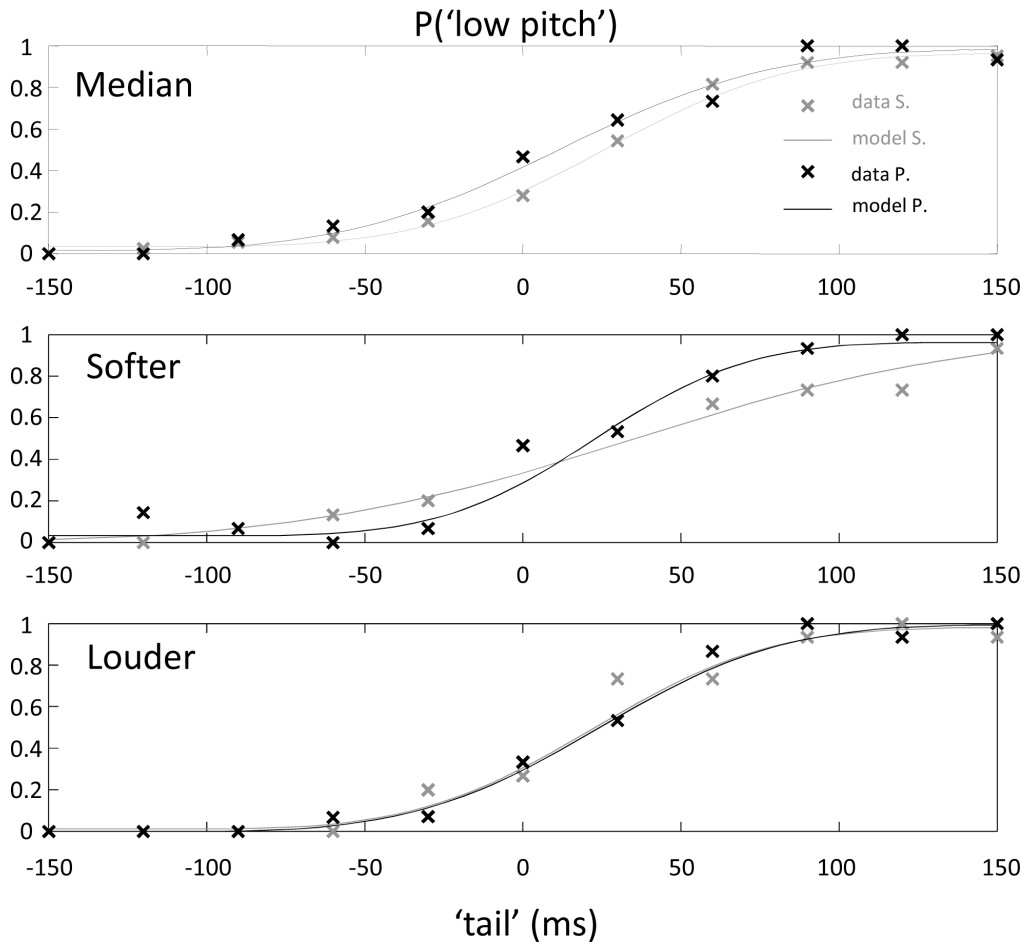


Figure 3. Example psychophysical data from one listener. Probability of responses for the surprising (S), and predictable (P), conditions in the three matched perceived loudness levels (standard = 70 dB SPL, soft = 60 dB SPL, and loud = 80 dB SPL). Crosses show data points and smooth lines show psychometric functions fitted to the data.

We performed a Kolmogorov-Smirnov test to assess whether these parameters were normally distributed and found that the slopes were not. The slopes were then normalised by means of a logarithmic transformation. Figure 4 shows, for all subjects, the slope (a measure of the sensitivity) and the lapse rate estimated in the surprising condition plotted against the predictable condition for the three levels of loudness (standard, soft, and loud). This figure shows that for most of the subjects the sensitivity is higher in the predictable compared with the surprising context. In a similar vein to our analysis of the RT data, we performed separate ANOVA tests for each of the psychometric parameters estimated: bias, log slope, and lapse rate. We found that, while surprise did not introduce a bias effect, nor an increase in lapse rate, it did diminish sensitivity ($p = 0.028$), irrespective of sound level. We also found a

trend-level enhancement in sensitivity for increasing loudness ($p = 0.064$) and an interaction between sound level and surprise for lapse rate ($p = 0.081$) that approached, but failed to reach, significance. The interaction trend (visible in figure 4) was for a greater effect for surprising events (S_L and S_S) than for standard events in the surprising context (S_M).

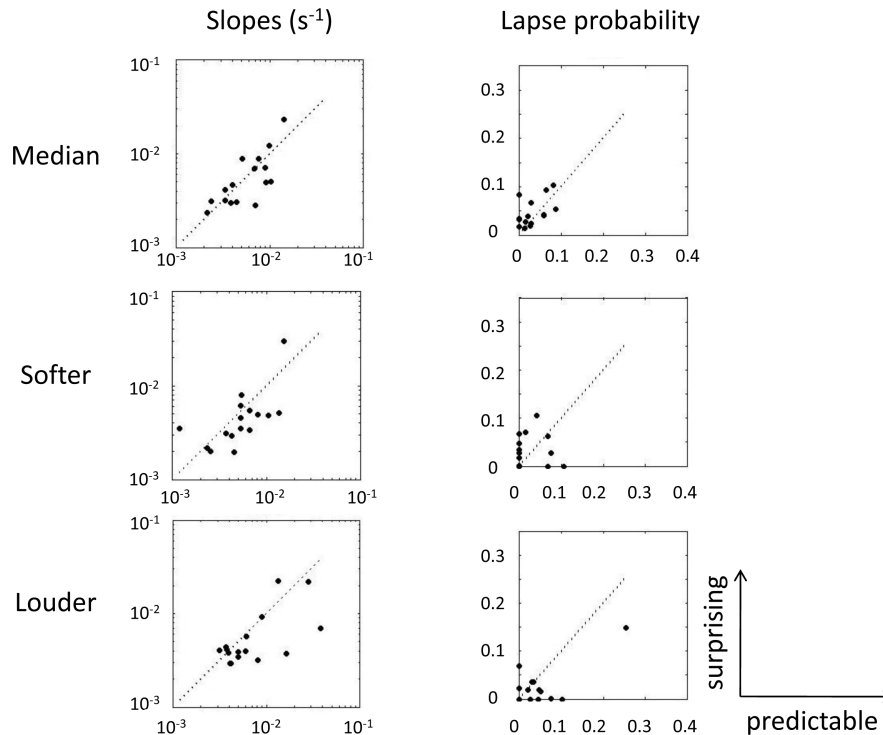


Figure 4. Surprised-induced noise in sensory decisions. Scatter plot of the slopes (or inverse variance) estimated for individual listeners. The dashed line represents equal slopes for the surprising and predictable conditions; that is, no effect of surprise. For most listeners, variance is higher (slopes are lower) in the surprising than in the predictable context, for median, softer, and louder stimuli. There is no significant evidence for an effect on the lapse rate.

4 Discussion

Our study shows that surprising contexts decrease sensitivity in auditory perceptual decisions. Unpredictable stimulus changes increased internal noise and this led to poorer sensory judgments, even though these changes were incidental to the actual judgment task.

In this paper we explicitly modelled the separate contributions of stimulus-dependent noise, (σ), and stimulus-independent noise, or lapse rate (λ), to auditory perceptual decisions. The first component reflects a graded form of distraction that partially withdraws perceptual resources from the task at hand, and thereby reduces perceptual sensitivity. Lapse rate, on the other hand, reflects complete distraction that leads to guessing, reducing task performance to chance level on some trials. Crucially, we show that the significant fall in performance in a surprising context is caused by decreases in sensitivity. We did not find an effect on the lapse rate.

This finding is consistent with conclusions drawn previously, although the earlier study did not, in fact, model the specific contributions of perceptual sensitivity and lapse rate separately. It should be noted that in our paradigm (as in those cited above) all surprising events also correspond to a change in stimulus parameters. Therefore, we cannot say whether

the effect we find is specifically due to surprise or to a change, and have used surprise and unpredictable change without distinction throughout the report.

Our findings reveal that surprising contexts lead to stimulus-dependant and not stimulus-independent noisier perceptual decisions. Importantly, this result holds when the judgment may be made using sensory information available within a single trial, rather than relying on a comparison between current inputs and an internal reference based on previous trials. This suggests that the effect is not due to interactions between surprise and memory. This is the case even when the surprising change occurs in a task-irrelevant dimension of the very same auditory object, which suggests that listeners cannot selectively attend to one dimension of an object, while simultaneously ignoring another dimension of the same object (Duncan 1984). If they could, then all perceptual resources would be used to attend to the task-relevant dimension and there should be no decreases in performance. The fall in performance, caused by surprising changes in a task-irrelevant object dimension, provides evidence that listeners were not able to ignore the irrelevant surprise.

We show slower responses to oddballs, replicating the findings of Schroeger and Wolff (1998). This is in keeping with the idea that unpredictable events are monitored in parallel with the current attentional goal (Vuilleumier 2005). In this context, unpredictable events are distractors that partially drag attention, or perceptual resources, toward themselves (Itti and Baldi 2009), and away from the task at hand. We interpret this, together with the fact that louder events slowed listeners' responses, in terms of a behavioural salience account for surprising and loud stimuli, which fits nicely with the theory that hearing evolved as an 'early warning system' (Scharf 1998).

Alternatively, the results can be interpreted in the context of conjectures on the costs and benefits of invalid and valid spatial cues in visual attention tasks. The disadvantage of surprising, and the advantage of predictable, contexts on the accuracy of decisions about acoustic information is analogous to the effects of expectations based on invalid and valid cues, respectively. In fact, the experimental design causes the listener to build an internal expectation for a specific standard loudness level. A violation of this expectation, which can be seen as analogous to an invalid cue trial, might increase internal perceptual noise and thus lead to slower and poorer performance. In the Posner (1980) paradigm these expectations are built based on task-relevant cues that have a direct impact on task performance. These data, however, show the same fall in performance even if expectations are formed based on cues incidental to the task, and where the expectations themselves are implicitly built and task irrelevant.

In summary, our results suggest that surprising events, or environmental uncertainty, can cause partial disengagement from a current task and disrupt goal-directed behaviour. This is likely to be ecologically advantageous in surprising and potentially threatening contexts, and might be part of a flexible system for automatically allocating attentional resources.

Acknowledgements. We are grateful to Sukhbinder Kumar for sharing the code that generated iterated rippled noise and to Steve Fleming for useful discussions. We thank the volunteers for participating in this study. This work has been funded by the Wellcome Trust and the Gatsby Charitable Foundation.

References

- Dalton P, Lavie N, 2004 "Auditory attentional capture: effects of singleton distractor sounds" *Journal of Experimental Psychology: Human Perception and Performance* **30** 180–193 ◀
- Dalton P, Lavie N, 2007 "Overriding auditory attentional capture" *Perception & Psychophysics* **69** 162–171 doi:10.3758/BF03193739 ◀
- Duncan J, 1984 "Selective attention and the organization of visual information" *Journal of Experimental Psychology: General* **113** 501–517 doi:10.1037/0096-3445.113.4.501 ◀

- Horvath J, Winkler I, 2010 "Distraction in a continuous-stimulation detection task" *Biological Psychology* **83** 229–238 doi:10.1016/j.biopsycho.2010.01.004 ◀
- Itti L, Baldi P, 2009 "Bayesian surprise attracts human attention" *Vision Research* **49** 1295–1306 doi:10.1016/j.visres.2008.09.007 ◀
- Levitt H, 1971 "Transformed up-down methods in psychoacoustics" *The Journal of the Acoustical Society of America* **49** 467–477 doi:10.1121/1.1912375 ◀
- Macmillan N A, Creelman C D, 2005 *Detection Theory: A User's Guide* (Mahwah, NJ: Lawrence Erlbaum Associates) ◀
- Parmentier F B, Andres P, 2010 "The involuntary capture of attention by sound: novelty and postnovelty distraction in young and older adults" *Experimental Psychology* **57** 68–76 ◀
- Patterson R D, Yost W A, Handel S, Datta A J, 2000 "The perceptual tone/noise ratio of merged iterated rippled noises" *The Journal of the Acoustical Society of America* **107** 1578–1588 doi:10.1121/1.428442 ◀
- Posner M I, 1980 "Orienting of attention" *The Quarterly Journal of Experimental Psychology* **32** 3–25 doi:10.1080/0033558008248231 ◀
- Rasmussen C E, 2001 "minimize.m", software package, http://www.mit.edu/~rsalakhu/code_DBM/91minimize.m ◀
- Scharf B, 1998 "Auditory attention: the psychoacoustical approach" in *Attention* Ed. H Pashler pp 75–117 (Hove, Sussex: Psychology Press) ◀
- Schroeger E, Wolff C, 1998 "Behavioral and electrophysiological effects of task-irrelevant sound change: a new distraction paradigm" *Cognitive Brain Research* **7** 71–87 doi:10.1016/S0926-6410(98)00013-5 ◀
- Sussman E, Winkler I, Schroeger E, 2003 "Top-down control over involuntary attention switching in the auditory modality" *Psychonomic Bulletin & Review* **10** 630–637 doi:10.3758/BF03196525 ◀
- Theeuwes J, 1992 "Perceptual selectivity for color and form" *Perception & Psychophysics* **51** 599–606 doi:10.3758/BF03211656 ◀
- Treutwein B, Strasburger H, 1999 "Fitting the psychometric function" *Perception & Psychophysics* **61** 87–106 doi:10.3758/BF03211951 ◀
- Ulanovsky N, Las L, Nelken I, 2003 "Processing of low-probability sounds by cortical neurons" *Nature Neuroscience* **6** 391–398 doi:10.1038/nn1032 ◀
- Vuilleumier P, 2005 "How brains beware: neural mechanisms of emotional attention" *Trends in Cognitive Sciences* **9** 585–594 doi:10.1016/j.tics.2005.10.011 ◀
- Watkins S, Dalton P, Lavie N, Rees G, 2007 "Brain mechanisms mediating auditory attentional capture in humans" *Cerebral Cortex* **17** 1694–1700 doi:10.1093/cercor/bhl080 ◀
- Whiteley L, Sahani M, 2008 "Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes" *Journal of Vision* **8** 2–2 doi:10.1167/8.3.2 ◀
- Wichmann F A, Hill N J, 2001 "The psychometric function: I. Fitting, sampling, and goodness of fit" *Perception & Psychophysics* **63** 1293–1313 doi:10.3758/BF03194544 ◀
- Yost W A, 1996 "Pitch of iterated rippled noise" *The Journal of the Acoustical Society of America* **100** 511–518 doi:10.1121/1.415873 ◀