Sensitivity to Pain Expectations: A Bayesian Model of Individual Differences

Hoskin, R.[1,3], Berzuini, C. C. [2,3], Acosta-Kane, D.[4], El-Deredy, W. [1,5], Guo, H.[2], & Talmi, D.[1,4]

1. Division of Neuroscience and Experimental Psychology, University of Manchester
2. Centre for Biostatistics, University of Manchester.
3. Joint first authors
4. Princeton University
5. School of Biomedical Engineering, University of Valparaiso, Chile

Word count (Excl. Abstract, Supplementary Material, Appendix & References): 8,944

Correspondence: Deborah Talmi, Division of Neuroscience and Experimental Psychology, Faculty of Biology, Medicine and Health, University of Manchester, Oxford Road, Manchester, United Kingdom, M13 9PL. Deborah.Talmi@manchester.ac.uk +44 (0) 161 275 1968

Abstract

The thoughts and feelings people have about pain (referred to as 'Pain expectations') are known to alter the perception of pain. However we know very little about the cognitive processes that underpin pain expectations, or what drives the differing effect that pain expectations have between individuals. This paper details the testing of a model of pain perception which formalises the response to pain in terms of a Bayesian prior-to-posterior updating process. Using data acquired from a short and deception-free predictive cue task, it was found that this Bayesian model predicted ratings of pain better than other, simpler models. At the group level, the results confirmed two core predictions of Hierarchical Predictive Processing; that expectation alters perception and that increased uncertainty in the expectation reduces its impact on perception. The addition to the model of parameters relating to trait differences in pain expectation, improved its fit, suggesting that such traits (which have been relatively neglected in experimental work on pain perception) play a significant role in perception beyond those expectations triggered by the pain cue. When model parameters were allowed to vary by participant, the model's fit improved further. This final model produced a characterisation of each individual's sensitivity to pain expectations. These findings are relevant for the understanding of the cognitive basis of pain expectations, and potentially useful for guiding patient stratification and clinical experimentation.

Key words: Pain; Expectation; Uncertainty; Bayes; Perception; Placebo effect

1. Introduction

The experience of pain, like all other sensory experience, is a result not only of the objective reality, such as the degree of tissue damage, but also of the sufferer's beliefs about pain. Conscious and unconscious thoughts and beliefs that people have about imminent pain are referred to as 'pain expectations' (Schrooten, Vlaeyen, & Morley, 2012). The effect of pain expectations on pain experience is evident in both laboratory and clinical experiments (e.g. Atlas & Wager, 2012; Bingel et al., 2011; Colloca, & Benedetti, 2006; Peerdeman et al., 2016; Tracey, 2010). In addition, neural correlates of the effect of expectation on pain perception have been established (Brown., et al 2008a; Ploghaus et al., 1999; Seymour et al., 2004; Tracey, 2010; Wager et al., 2004; Watson et al., 2009). Despite these advances there remains a significant gap in our knowledge regarding the specific cognitive processes that underpin pain expectations.

Hierarchical Predictive Processing (HPP) provides the dominant theoretical framework for understanding the effects of expectation on perception (Clark, 2013, Friston, 2003), including pain perception (Buchel, Geuter, Sprenger, & Eippert, 2014; Tabor, Thacker, Moseley, & Körding, 2017; Van den Bergh, Witthöft, Petersen, & Brown, 2017). HPP stipulates that perception is biased towards the expected level of pain, and that this bias will be stronger when the expectation is more certain, since expectation uncertainty causes the suppression of top-down, prior-driven signals, leading to greater importance being placed on the bottom-up sensory input. A number of studies confirm the specific predictions of HPP for pain perception, beyond the established biasing effect that pain expectations exert on pain experience. The effect of pain expectations is enhanced when expectations are more precise (Brown, Seymour, Boyle, El-Deredy, & Jones, 2008b; Colloca, Petrovic, Wager, Ingvar, & Benedetti, 2010); and formal Bayesian models have been shown to provide a good fit to pain reports in placebo analgesia studies (Anchisi, Zanon, Plaghki, Kirsch, & Claggett, 2015; Jung, Lee, Wallraven, & Chae, 2017) and to neural data from the anterior insula (Geuter, Boll, Eippert, & Büchel, 2017).

Although HPP guides much of the research on the effect of expectations on perception, some findings appear to contradict its prediction concerning the impact of the precision of expectations on pain experience. Watkinson et al., (2013) delivered two different distributions of pain stimuli, a unimodal and a

bimodal distribution. The two distributions had the same mean pain level, but the bimodal distribution had a larger variance. In direct contradiction with HPP, and in support of the alternative, Range Frequency Theory (Parducci., 1963), the effect of expectations was stronger in the bimodal condition, despite the greater uncertainty generated by the larger stimulation variance. Likewise, Yoshida et al., (2013) altered participant's expectations relating to upcoming pain stimulations by presenting fictitious pain ratings of the stimulus prior to their delivery. They found that increasing pain uncertainty contributed independently to increased pain experience, in contradiction with HPP. In addition, it is not clear from Yoshida et al. data that uncertainty significantly modulated the bias induced by the mean of the fictitious pain ratings, as predicted by HPP. These findings suggest that further work is needed before HPP can be accepted as a viable framework for understanding pain expectations.

Over and above the conflicting findings around the impact of expectation uncertainty, we also know very little about how specific processes that underlie pain expectations are integrated to inform the pain percept. To unpack the construct of pain expectations to its underlying information processing mechanisms requires a statistical model of the pain perception process to be formulated which accommodates parameters that describe different facets of pain expectations. Furthermore, in order for such a model to be useful in understanding an individual's response to both pain and placebo analgesia, it needs to be capable of predicting the effect of pain expectations not just at a group level, but also at an individual level. For example such a model would need to be able to identify individuals that experience high levels of pain because they have a trait-like bias to expect high levels of pain (i.e. always expecting high pain, independent of context), and to distinguish such individual from those who experience high pain because they are highly pessimistic in assessing the information they are given about the placebo treatment, or because they are over-confident in their mildly-pessimistic expectations, so that they rely less on their sensory data. That level of understanding is necessary in order to predict an individual's response, providing a basis for a clinical tool that can support clinical decision making.

This paper details the construction and testing of a model of impact of pain expectations on pain perception using experimental data garnered from a novel predictive cue task. The purpose of this model

construction was threefold. Firstly, we wished to assess the effect of expectation uncertainty on pain perception, in light of the conflicting past results mentioned above. Our second objective was to identify, using the experimental data, a number of putative cognitive processes that give rise to the umbrella term 'pain expectations'. Finally, leading on from the second objective, we wished to assess whether a model could be constructed that would enable individuals to be distinguished based on the aforementioned cognitive processes.

We collected two sets of experimental data via two separate experiments using independent samples. A series of increasingly complex statistical models were constructed and their performance was compared using the data from the first experiment. The simplest model (Model 1) where the pain participants experienced was influenced only by the actual, delivered pain was effectively used as a baseline to compare to three other models, which successively included additional facets of pain expectations (Table 1). In the 'Mean-only Bayesian model' (Model 2), pain experience was additionally influenced by the mean of the expected pain, computed based on the information provided by the cue. The 'Mean-and-variance Bayesian model' (Model 3), inspired by HPP, took into consideration the variance of expected pain, a function of the discrepancy between the possible pain levels indicated by the cue. Like Model 2, Model 3 also allowed pain experience to be affected by the mean of the expected pain, but here the impact of mean expected pain was modulated by its uncertainty. Finally the 'Full Bayesian model' (Model 4) additionally included the effect of cue-independent pre-experimental pain expectations in addition to the cue-dependent expectations used in models 2 and 3. In this formulation, cue-dependent pain expectations relate to those triggered on each trial by the cue, while cue-independent expectations capture more stable, trait-like differences in the propensity to believe that pain will be greater or weaker, independent of the effect of specific local cues. Finally, to satisfy the objective of creating a model that can characterise pain expectations at the individual level, we compared the winning group-level model to an individual-level variant that included individual-level random effects (Model 5). We considered the winning model to be able to usefully characterise pain expectations at an individual level if the 'individual-level' variant of the model was significantly better at predicting pain perception than its equivalent group level variant. The

second experimental dataset was used to validate the findings from the first dataset, thus providing evidence that the winning model could predict outcomes in a dataset other than the one from which it was constructed (cf. Maloney & Zhang 2010).

We hypothesised that the aforementioned model comparison will provide support for the HPP framework. More specifically we predict that Model 3 will provide a better fit to the data than models 1 and 2 (thus showing that expectation uncertainty significantly affects pain perception) and that the value of model parameter describing the effect of expectation uncertainty will support the predict from HPP that increased uncertainty reduces the influence of pain expectations. Based on the vast clinical literature on individual differences in pain catastrophizing (Sullivan, Bishop, & Pivik, 1995) we hypothesised that Model 4 will provide a better fit for the data than Model 3, suggesting that both cue-dependent and cue-independent expectations have separate influences on pain perception. Finally, we hypothesised that when individual-level random effects are added to the winning model this will significantly improve the fit of the model to the data even when that model was penalised for the increased complexity, suggesting that it is possible to characterise sensitivity to pain expectations at an individual level.

2. Method

2.1 General experimental design

Two experiments were conducted on independent samples. Experiment 2 was conducted for validation, with task delivery modifications introduced to enhance the translational impact of the approach. During each experiment participants performed a ‘pain rating’ task (Figure 1.). In each trial of the task, participants’ expectations regarding the upcoming stimulation were manipulated. This manipulation was achieved by offering the participants a choice between two ‘cues’ expressing different probability distributions for the intensity of the upcoming stimulation in terms of the possible intensities and corresponding probabilities. After selecting one of the cues the participant then experienced a pain stimulation level generated in accordance with the probabilities represented by the selected cue. Finally,

the participant was required to rate the intensity of the stimulation they received. This choice mechanism was used (in contrast to just presenting participants with one cue) to ensure that the participants both attended to, and understood, the pain probability cues they were presented with. The choice mechanism also acted to give the participants a sense of control over the upcoming pain, so as to minimise the possible influence of learned helplessness on their pain experience (Bhat et al., 2010). The trials were designed in such a way that one cue (herein referred to as the 'target') was always preferable to the other (herein referred to as the 'lure') in terms of expected pain intensity. Trials where the participants chose the lure (or failed to make a choice at all) were discarded from the data analysis as in such instances we could not be sure that the participants had understood the cues. The targets were designed in accordance with a factorial design involving the systematic manipulation of the relevant aspects of pain expectation.

The task utilised here is novel in the context of pain research. However, it has been used extensively to examine how people value financial outcomes (Talmi, Atkinson, & El-Deredy, 2013) as well as taste and pain outcomes (Hird et al., 2017). It differs from placebo paradigms because the manipulated quantity is the individual's expectation about the impending pain, rather than about the efficacy of treatment. One particular feature of our task (in contrast with placebo and other cue-based paradigms) is the avoidance of deception. In placebo paradigms participants are told that an ineffective intervention (e.g. an inert cream) is in fact a proven treatment (e.g. an active analgesic). In current cue-based paradigms (e.g. Atlas et al., 2010), cues intentionally mislead participants, for example informing them that an impending pain stimulus will be high intensity, when it may in fact be of moderate intensity. We avoided deception firstly because the integrity of data resulting from deception paradigms is reliant on the participant remaining naïve to the deception throughout the experiment, and secondly because deception is problematic as far as translation into clinical settings is concerned. Files needed to run the experimental paradigm and model detailed in this paper are available to download from BioRxiv.

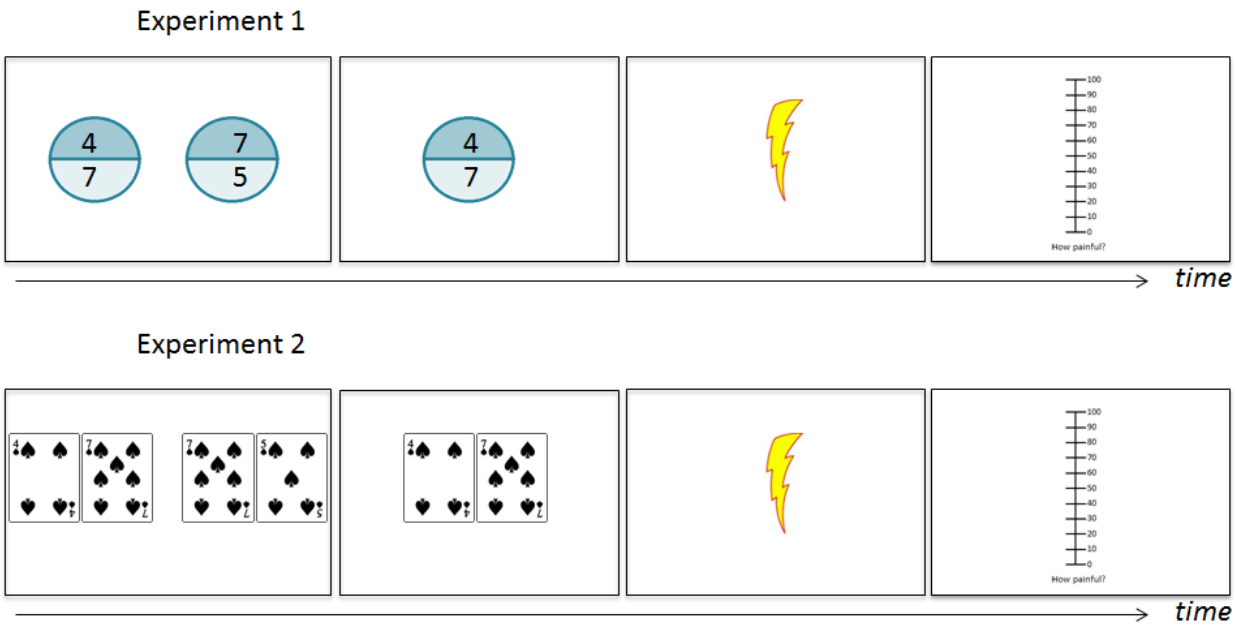

Figure 1. Schematic of one experimental trial in experiments 1 (top) and 2 (bottom). *Each trial began with a cue selection phase (left panel) where the participant had to choose between two cues. Participants who understand these cues should choose the left, target cue to minimise the chances of receiving strong pain. Once a cue was selected it appeared on its own for 2s (second panel). That cue was assumed to influence participants' pain expectations. Here the selected cue gives the participant 50% chance of receiving pain level '4' and 50% of receiving pain level '7'. With this cue selected the expected pain level for a participant who is not particularly anxious or overly optimistic is 5.5. Participants then received a painful electric stimulation corresponding to one of the levels depicted on the selected cue, e.g. a '4' (third panel). Finally the participant rates their pain experience (fourth panel). In Experiment 1 the size of the portion on which the number was displayed corresponded to the probability of receiving pain of that level. In Experiment 2 all probabilities were fixed to 50%.*

2.3 Experiment 1: Method

2.3.1 Experimental design

Figure 1 shows an example of a trial from Experiment 1. In this example the target cue presented a 50% chance of receiving level 4 pain and 50% chance of receiving level 7 pain. This target therefore provided an

expected pain intensity of 5.5, and thus it is superior to the lure cue, which presents a 50% chance of receiving level 5 pain and 50% chance of receiving level 7 pain (expected intensity = 6). The cues were designed to provide a factorial design involving the manipulation of the following variables across the sequence of trials experienced by each individual: the expected value of the pain (the mean: 4.5, 5.5), the number of different pain levels that could potentially be delivered (2, 4), the prediction error (the difference between expected level of pain and intensity of the delivered stimulation: 0.5,1.5) and its sign (positive or negative). Using the trial depicted in Figure 1 as an example, if the target cue was selected and the level 7 pain subsequently delivered, then the prediction error would be +1.5 (7 - 5.5). The positive sign of the prediction error indicates that the pain level that was experienced was higher than the average pain level participants would expect on the basis of the cue they selected. In summary, the experimental corresponded to a 2 (mean expected pain: 4.5, 5.5) x 2 (number of cue options: 2, 4) x 2 (direction of the prediction error: positive, negative) x 2 (the size of the prediction error: 0.5, 1.5) design.

2.3.2 Participants

Sixteen undergraduates from the Manchester University School of Psychological Sciences (14 female, Mean age 19.6, σ = 1.36) participated in the study for course credit. Participants were excluded if they had a history of psychiatric or neurological disorders. All participants achieved criterion of 85% accurate performance (defined as choosing the target cue). The study received ethical approval from the North West 6 Research Ethics committee (Greater Manchester South).

2.3.3 Materials

Acute experimental pain was delivered to the participants using electric stimulation, shown to produce effects of expectations that are equivalent to those produced with the more typical laser stimulation (Hird, Jones, Talmi, & El-Deredy, 2018). Compared to laser pain, which is used more often in the laboratory, the electric stimulation method is safe to use repeatedly and easier to implement in clinical settings. The electrical stimulations were delivered to the back of the right hand via a ring electrode built in-house

(Medical Physics, Salford Royal Hospital) attached to a Digitimer DS5 Isolated Bipolar Constant Current Stimulator (www.digitimer.com). To ensure adequate conductance between the electrode and the skin, the back of each participant's hand was prepared with Nuprep Skin Preparation Gel and Ten20 Conductive Paste prior to the electrode being attached. The experimental paradigm itself was delivered via a laptop using Cogent2000 on a Matlab platform (www.Mathworks.com). The inputs to the DS5 machine were sent from Matlab via the Spike2 software (www.ced.co.uk) and a 1401plus data acquisition interface (www.ced.co.uk).

Cues consisted of charts which were 240 pixels in diameter and used colours that were equiluminant (thus making each chart equiluminant). These charts used combinations of the 5 pain levels (3,4,5,6,7) and 6 different proportions (25%, 33%, 50%, 67%, 75%, 100%) to manipulate aspects of the pain expected (Supplementary Table2.)

2.3.4 Procedure

Participants were given an information sheet prior to the study informing them of the justification for the study and of the use of electrical stimulation. On arriving at the laboratory participants were first asked to sign a consent form, and then to confirm that they had read the information sheet. The electrode was then attached to the back of the participant's right hand. Once the electrode was attached the participants undertook a pain calibration procedure. This procedure was necessary firstly to ensure that the participant could tolerate the stimulations, and secondly to ensure that the stimulations were psychologically equivalent across participants. Once the pain calibration procedure was complete participants were given instructions relating to the pain rating task. Participants were then given 6 practice trials of that task before undertaking the task proper.

2.3.4.1 Pain calibration.

During this procedure participants received a series of stimulations, starting from 0.2V, and incrementing by 0.2V at each step. Participants rated each stimulation on a scale from 0 – 10 where a score of 0 reflected not being able to feel the stimulation, 3 reflected a stimulation level that was on the threshold of being painful, 7 related to a stimulation that was deemed 'painful but still tolerable' and 10 related to 'unbearable pain'. The scaling procedure was terminated once the participant reported the level of pain as being equivalent to '7' on the scale. This calibration procedure was performed twice to allow for initial habituation/sensitisation to the stimulation. The voltage levels rated as '3', '4', '5', '6', and '7' on the second run of the calibration procedure were used as the different pain levels to be delivered during the pain rating task. The average voltage used for each pain level are shown in Supplementary Table 1.

2.3.4.2 Pain rating task.

Each trial began with a 250ms fixation cross, before two cues, presented as two pie charts, appeared side-by-side on the screen (+/- 160 pixels from the centre: Figure 1). Participants were required to choose between these cues using the mouse. Numbers on the 'slices' of each pie chart indicated the level of stimulation that the slice related to, while the size of the slice depicted the probability of the stimulation level being delivered if the chart was selected. Participants had up to 8 seconds to make their selection. The target cue was always clearly preferable to the lure in terms of expected pain value (see General Methods). If the participant failed to make any selection within the 8s time limit then the lure chart was automatically selected. The selected cue was then presented alone in the centre of the screen for 2 seconds, at which point the cue disappeared and one of the pain levels depicted on the cue was delivered to the back of the hand according to the probabilities depicted on the cue. Finally a visual-analogue scale (Hjermstad et al., 2011) appeared 500ms after the offset of the stimulation. Participants were required to rate their experience of the stimulation on this scale using the mouse, with the trial only proceeding once a rating had been given. Participants were told that even though the stimulation they would receive would always be one of those predicted by the selected cue, they should rate the pain they actually felt. After the rating

was provided an inter-trial interval occurred of a length selected randomly from 1000, 1500, 2000, 2500 & 3000ms (average 2000ms).

2.3.4.3 Experimental task.

Participants undertook 5 blocks of trials, with 32 trials in each block, giving a total of 160 trials. Each block consisted of a randomised presentation of 24 types of trials, depicted in Supplementary Table 2. In 16 of these the target cue corresponded to one of the 16 experimental conditions according to the 2x2x2x2 experimental design described above. Four additional trial types were required to complete the design (see Supplementary Table 2). Four further trials types, involving target cues which depicted the delivery of pain levels 3, 4, 5 and 6 with 100% probability, were included for the purpose of a manipulation check. A self-timed break was given every 20 trials.

2.4 Experiment 2: Method

The design of Experiment 2 was similar to that of Experiment 1. The following changes were made with the potential clinical application of the paradigm in mind. Firstly sets of playing cards were used as cues instead of pie charts (Figure 2). This allowed the results from Experiment 1 to be (potentially) replicated using cues that may be more readily understood by a clinical audience. Secondly only trial types involving 2 different intensities with equal (50%) probabilities were included in the factorial design. This meant that only 3 aspects of pain expectation were manipulated within participants; the expected pain level (4.5, 5, 5.5), the size of the prediction error (0.5, 1.5, 2), and the direction of the prediction error (positive or negative). This change allowed a reduction to the number of trials appearing in the experiment, and therefore a reduction in the length of the experiment. This was thought beneficial due to the limited time available in the clinic for diagnostic procedures.

2.4.1 Participants

Thirty-two undergraduates from the Manchester University Psychology Department (26 female, Mean age 19.3, σ = 1) participated in the study for course credit. Participants were excluded if they had a history of psychiatric or neurological disorders. The study received ethical approval from the North West 6 Research Ethics committee (Greater Manchester South). Data from three participants was excluded because they did not achieve criterion of 85% accurate performance (defined as choosing the target cue during the cue selection phase) leaving a sample of 29 (24 female, mean age 19.3, σ = 1).

2.4.2 Materials and procedure

The materials were the same as in Experiment 1, but instead of charts, the cues consisted of spade suit playing cards. Each cue took the form of a pair of playing cards (100x130 pixels in size), with each card corresponding to one slice of the charts used in Experiment 1. During the selection phase the two cards from each pair of cards were presented 225 and 100 pixels from the centre of the screen, with one pair being presented to the left of the screen and the other to the right. After selection, the two cards comprising the selected cue were presented +/- 75 pixels from the centre of the screen.

Participants undertook 5 blocks of trials. Each block consisted of a randomised presentation of 16 types of trials depicted in Supplementary Table 3. In 8 of these the target cue corresponded to one of the 8 experimental conditions corresponding to a 2 (the expected pain level: 4.5, 5.5) x 2 (the size of the prediction error: 0.5, 1.5) x 2 (the direction of the prediction error: positive or negative) experimental design. In 4 trial types the target cue corresponded to one of the 4 experimental conditions corresponding to a 2 (the size of the prediction error: 0.5, 1.5) x 2 (the direction of the prediction error: positive or negative) experimental design. We employed these two factorial designs together in order to compare their efficacy, and decide which design was best for future experiments. Four additional trial types involving target cues which depicted the delivery of pain levels 3, 4, 5 and 6 with 100% probability were again included for the purpose of a manipulation check. In total, therefore, each participant undertook 80 trials,

with self-timed breaks given every 20 trials. The average voltage used for each pain level during the experiment are shown in Supplementary Table 1.

2.5 Statistical analysis

2.5.1 Modelling approach

HPP suggests that our perceptual systems are organised in a hierarchical manner. Expectations (priors) held at higher levels of the perceptual hierarchy, are compared to incoming pain at lower levels. Errors in prediction are then projected back to the higher level, causing the resulting pain experience (posterior) to be updated, which in turn alters the basis on which subsequent priors are calculated (Buchel et al., 2014). Our model is influenced by this framework, formalising response to pain in terms of a Bayesian prior-to-posterior updating process.

For participant *i* during trial *j* we represent the pain stimulation that was delivered as X_{ij} and the pain they reported as R_{ij}, both expressed on a 0-100 scale. The ratings were linearly transformed ($R_{ij} \rightarrow R'_{ij} = a_i + b_i \times R_{ij}$) using the pain responses on the 100% probability trials to standardize the responses across participants so as to avoid issues of differential scale use. Thus $a_i$and b_i were selected to minimize the squared error between participants' ratings and the delivered intensity. Before the pain stimulation was delivered, participant *i* at the *j*th trial also saw a cue. Let the cue consist of information Z_{ij}, involving the possible magnitudes ($v_{1ij} \ldots v_{m_{ij}ij}$) of the incoming pain stimulus, also expressed on a 0-100 scale, and their corresponding probabilities of occurrence, ($w_{1ij} \ldots w_{m_{ij}ij}$). This information provides a specification of a probability distribution Z_{ij} with mean q_{ij} and standard deviation sd_{ij} respectively given by

$$q_{ij} = \sum_{k=1}^{m_{ij}} w_{kij} v_{kij} \qquad sd_{ij} = \sqrt{(\sum_{k=1}^{m_{ij}} w_{kij}(v_{kij} - q_{ij})^2)}$$

On the *j*th trial, an individual *i* exposed to cue Z_{ij} represents the predicted intensity of the upcoming pain as a probability distribution P($R_{ij} \mid Z_{ij}$). Using the terminology of HPP, this is the 'prior distribution'.

We formulate $P(R_{ij} \mid Z_{ij})$ as a product of two normal distributions, one incorporating cue-independent information and the other incorporating the cue information (see appendix for derivation):

1) $P(R_{ij} \mid Z_{ij}) \propto Normal(q_{ij}, \rho + \eta \times sd_{ij}) \times Normal(\mu, \upsilon)$

where *Normal(x, y)* stands for the normal distribution with mean *a* and standard deviation *b,* occasionally truncated to make sure that random variables satisfy their range constraints, whenever appropriate.

The cue-independent $Normal(\mu, \upsilon)$ represents a stable, trait-like bias to expect pain of a certain intensity independently of the cue. Higher values of parameter μ (the mean of the cue-independent prior) reflect the expectation of higher levels of pain, independent of any cue information. Higher values of parameter υ (the standard deviation of the cue-independent prior) reflect greater uncertainty (and therefore lesser influence) of this cue-independent prior. Participants with high μ and low υ can be thought of as pessimistic about pain. The cue-dependent $Normal(q_{ij}, \rho + \eta \times sd_{ij})$ incorporates two parameters which modulate the effect of the cue information Z_{ij} on the likelihood of the rating R_{ij}. The parameter ρ controls the variance of the likelihood function generated by the cue. In effect it represents the precision of this prior. η modulates this precision according to the cue's variance. The influence of the cue increases with lower values of ρ and η. Together, these parameters determine the extent to which the cue information modulates the prior distribution, and thus the subsequent experience of pain.

The stimulation causes the prior distribution to be updated into a posterior. The posterior distribution P($R_{ij} \mid X_{ij}, Z_{ij}$) governs the pain intensity that the individual experiences and rates. Following from Bayes' theorem, the posterior is given, up to a proportionality constant, by the product of the likelihood of the stimulation and the prior:

$$P(R_{ij} \mid X_{ij}, Z_{ij}) \propto P(X_{ij} \mid R_{ij}, Z_{ij}) \times P(R_{ij} \mid Z_{ij})$$

In the appendix we show that evaluating the likelihood, $P(X_{ij} \mid R_{ij}, Z_{ij})$, results in the addition of a $Normal(X_{ij}, \beta)$ term in the posterior:

2) $P(R_{ij}|X_{ij},Z_{ij}) \propto Normal(X_{ij},\beta) \times P(R_{ij}|Z_{ij})$

In this term, each predicted value of R_{ij} is compared with the external sensory information, associating each such value with a measure of 'surprise', which our model takes to increase with an increasing distance between R_{ij} and the actual intensity of the stimulation, X_{ij}. The value of β determines the extent that the perceptual mechanism will tolerate a given disparity between the predicted experience and the actual sensation, with smaller values of β implying greater sharpness of the likelihood function around X_{ij} and hence a reduced tolerance. Plugging Equation 1 into Equation 2 we get the full model:

3) $P(R_{ij}|X_{ij},Z_{ij}) \propto Normal(X_{ij},\beta) \times Normal(q_{ij},\rho + \eta \times sd_{ij}) \times Normal(\mu,v)$

Here the prior is combined with the surprise in accord with the Bayesian formalization of the integration of multiple sources of evidence.

2.5.2 Model specification

Data from each experiment were analysed separately with the following models:

1. Baseline model: This model assumes that pain experience depends solely on the pain stimulation. No effect is attributed to pain expectations, whether arising from the cue, or from cue-independent effects of the prior.

$P(R_{ij}|X_{ij},Z_{ij}) = P(R_{ij}|X_{ij}) = Normal(X_{ij},\beta)$

2. Mean-only Bayesian model: Pain experience depends not only on the pain stimulation, but also on the mean of the cue, q_{ij}, modulated by the parameter ρ. This model does not incorporate the effects of uncertainty.

$P(R_{ij}|X_{ij},Z_{ij}) \propto Normal(X_{ij},\beta) \times Normal(q_{ij},\rho)$

3. Mean and Variance Bayesian model: Pain experience depends not only on the pain stimulation, but also on the mean and variance of the cue. These influence the uncertainty of the prior distribution of R_{ij} such that cues with higher sd_{ij}(modulated by η) influence the posterior less than precise cues.

$$P(R_{ij}|\, X_{ij}, Z_{ij}) \propto Normal(\, X_{ij}, \beta) \times Normal(q_{ij}, \rho + \eta \times sd_{ij})$$

4. *Full Bayesian model:* In this model pain experience is determined jointly by the pain stimulation and two tiers of expectations: stable cue-independent effects akin to optimism/pessimism trait and effects informed by the cue. Cue-independent effects are represented by a probability distribution with mean μ and variance υ, which is integrated with the information provided by the cue to form the final prior distribution.

$$P(R_{ij}|\, X_{ij}, Z_{ij}) \propto Normal(\, X_{ij}, \beta) \times Normal(q_{ij}, \rho + \eta \times sd_{ij}) \times Normal(\, \mu, \upsilon)$$

5. *Variable full Bayesian model:* In this model pain experience is also determined jointly by the pain stimulation and two tiers of expectations, however in this model the parameters are now allowed to vary by participant.

$$P(R_{ij}|\, X_{ij}, Z_{ij}) \propto Normal(\, X_{ij}, \beta_i) \times Normal(q_{ij}, \rho_i + \eta_i \times sd_{ij}) \times Normal(\, \mu_i, \nu_i)$$

2.5.3 Inference

Data from each experiment were analysed separately. For each model and each experiment, we calculated the Deviance Information Criterion (DIC, Spiegelhalter et al., 2002), which penalises models for increasing complexity. Using data from Experiment 1, we first compared fixed-parameter models 1-4 where all participants shared the same values for the parameters. The baseline model (Model 1) serves as the null hypothesis, as it assumes that the cue exerts no effect on pain experience, namely, no effect of expectations of any kind. Next, the same data were analysed with an elaboration of the winning fixed-parameters model, now allowing the parameters *β, ρ, η, μ, and ν,* and to vary across individuals (the 'Variable full Bayesian model', Model 5). We took the values of these parameters in individual *i* (respectively denoted by β_i ρ_i, η_i, μ_i, and ν_i,) to be independently drawn from corresponding population-

level normal distributions, with hyperparameters independently drawn from flat hyperprior distributions, resulting in a fully-Bayesian, hierarchical (multilevel) model. This model allows us to describe the variability of the individual-specific posterior estimates for parameters, and therefore to infer how the parameters vary across the population. Finally, we fitted these models to data from Experiment 2, to provide validation that the findings from the first dataset would generalise to an independent dataset.

As regards parameter estimation, the 'Full Bayesian model', when combined with data, gives rise to a Bayesian posterior distribution for the unknown model parameters. We sampled this posterior by using the Hamiltonian dynamics Markov chain Monte Carlo (HDMCMC) techniques (Metropolis, Rosenbluth, Rosenbluth, Teller, & Teller, 1953; Neal, 2012) incorporated in the program Stan (Stan Development Team, 2014a, 2014b). Any summary of the marginal posterior for any parameters of inferential interest can be reconstructed on the basis of the sampled value for that parameter, generated by the Markov chain after a suitable number of "warm-up" iterations. Initial values for the Markov chains were obtained via variational inference techniques (Wainwright, Jordan, Wainwright, & Jordan, 2008) based on a parametric approximation of the model posterior.

3. Results

The average performance during the cue selection task (i.e. proportion of targets selected) was 96% in Experiment 1 and 97% in Experiment 2, suggesting that participants understood the task well. In both experiments the mean ratings given during the trials where the cue provided 100% chance of a particular stimulation intensity increased linearly as a function of increased stimulation levels, confirming that participants were able to distinguish between the stimulation levels (Supplementary Table 4).

Table 1 summarizes the DIC values for the fixed-parameter models. Lower values indicate higher likelihood of the data being generated by the model and thus better fit. A difference between the DICs for two models that are greater than 10 are considered strong evidence against the model with the higher score. For Experiment 1, Model 1 fitted the data less well than the other models, which have all assumed

that the cue has some effect on the rating. The best-fitting fixed-parameter model was the 'Full Bayesian model' (Model 4). This finding supports the hypothesis that both cue-dependent and cue-independent biases influence pain perception. When the same models were tested using the data from Experiment 2, the results of the model comparisons were the same (Table1).

Table 1:DIC values for fixed and variable-parameter models.

			Experiment 1	Experiment 2
Fixed	1	Baseline model	17007	17304
	2	Mean-only model	16178	16179
	3	Mean and variance model	16138	16107
	4	Full Bayesian model	15937	15714
Variable	5	Full Bayesian model	15588	15047

Table 2: Estimates of the model parameters arising from fixed-parameter model.

	Experiment 1		**Experiment 2**	
Parameter	**Mean**	**Standard Error**	**Mean**	**Standard Error**
β	8.97	0.18	9.45	0.21
ρ	6.331	0.37	8.20	0.45
η	0.12	0.020	0.09	0.024
μ	46.05	0.57	45.61	0.42

v	12.04	0.49	10.67	0.28

Note to Table 2: *Parameter values from the 'Full Bayesian model' (Model 4), where the parameters are assumed to be identical across the participants. The values in the table represent the mean and standard error (the standard deviation of the marginal posterior distribution) of each estimated parameter, respectively. The corresponding credible interval, defined for each parameter by the mean plus/minus 1.96 times the standard deviation, can be interpreted to contain the true value of the parameter with probability 0.95. The model parameters can be described as follows.*

β: The standard deviation of the likelihood function that represents the actual, delivered pain stimulation. Controls how likely is a value of R_{ij} which is distant from X_{ij}, compared to a value of R_{ij} close to X_{ij}.

ρ: The increase in the standard deviation of prior distribution of R_{ij} due to a unit increase in the mean of the cue q_{ij}. This parameter controls how likely is a value of R_{ij} which is distant from q_{ij}, compared to a value of R_{ij} which is close to q_{ij}.

η: The increase in the standard deviation of the prior distribution of R_{ij} due to a unit increase in the standard deviation of the cue (sd_{ij}). This parameter captures an effect of the disparity in possible pain levels described in the cue on the precision of the prior.

μ: The mean of the distribution of the predicted pain intensity before incorporating the likelihood of the cue.

v: The standard deviation of the distribution of predicted pain intensity before incorporating the likelihood of the cue.

Table 2 summarizes the parameter estimates achieved from fitting Model 4, the winning 'Full Bayesian model'. Within this perspective, the finding that the average value of ρ is of smaller magnitude than the average value of β suggests that that the prior distribution for R_{ij} is biased towards the mean of the cue.

These parameters are inversely proportional to the strength of the cue and stimulus information when they are combined to form the posterior in this model. Thus, the parameter values suggest that an increase in the mean of the cue will, on average, and all other variables remaining equal, correspond to an increase in the mean of the rating R_{ij}. The positive average value of η supports the hypothesis that an increase in the cue variance will, on average, correspond to an increase in the variance of the prior distribution for R_{ij} and thus a smaller effect of the cue on the perceived intensity, in accordance with HPP.

Table 1 also contains the DIC scores for Model 5, the 'Variable full Bayesian model'. This model fitted significantly better than all fixed-parameter models (models 1-4), suggesting that individual differences are relevant for understanding how expectation influences perceived pain. Model 5 allows us to dissect a number of aspects of the psychological construct of pain expectations, illustrated in Figures 3-5. Figure 3 illustrates the way in which Model 5 explains the responses of two different participants in Experiment 1 to a trial where these participants were exposed to the same cue. The charts on the top row illustrate their cue-independent expectations before either the cue or stimulus is delivered. In the middle row, the likelihood of the cue (red dashed curve) is incorporated to create a posterior (black curve) which will later serve as the 'prior distribution' (blue curve), $P\ (R_{ij}|Z_{ij})$. In the last row the 'prior distribution' (now the blue curve) is incorporated with the likelihood of the stimulus (red dashed curve) leading to the posterior distribution (black curve), $P(R_{ij}|\ X_{ij}, Z_{ij})$. The left column of the figure describes a participant (participant 3) with parameters β = 11.41, ρ = 13.05, η = 0.105, μ = 44.82, and v = 10.00. The right column displays another participant (participant 4) with parameters β = 11.85, ρ = 4.54, η = 0.106, μ = 51.23, and v = 11.70. The left participant's value for ρ is markedly higher than that of the right participant. This causes the 'prior distribution' in the second row to be markedly less precise. In the third row, this results in the posterior being closer to the actual intensity of the delivered stimulus. We might characterize the behaviour displayed in the left diagram in terms of the cue exerting little influence on the participant's expectation and rating. The right participant's value for ρ is markedly lower, which is reflected by the sharper 'prior distribution', centred around the cue mean, q_{ij}. This causes the strong pull of the posterior towards the cue mean.

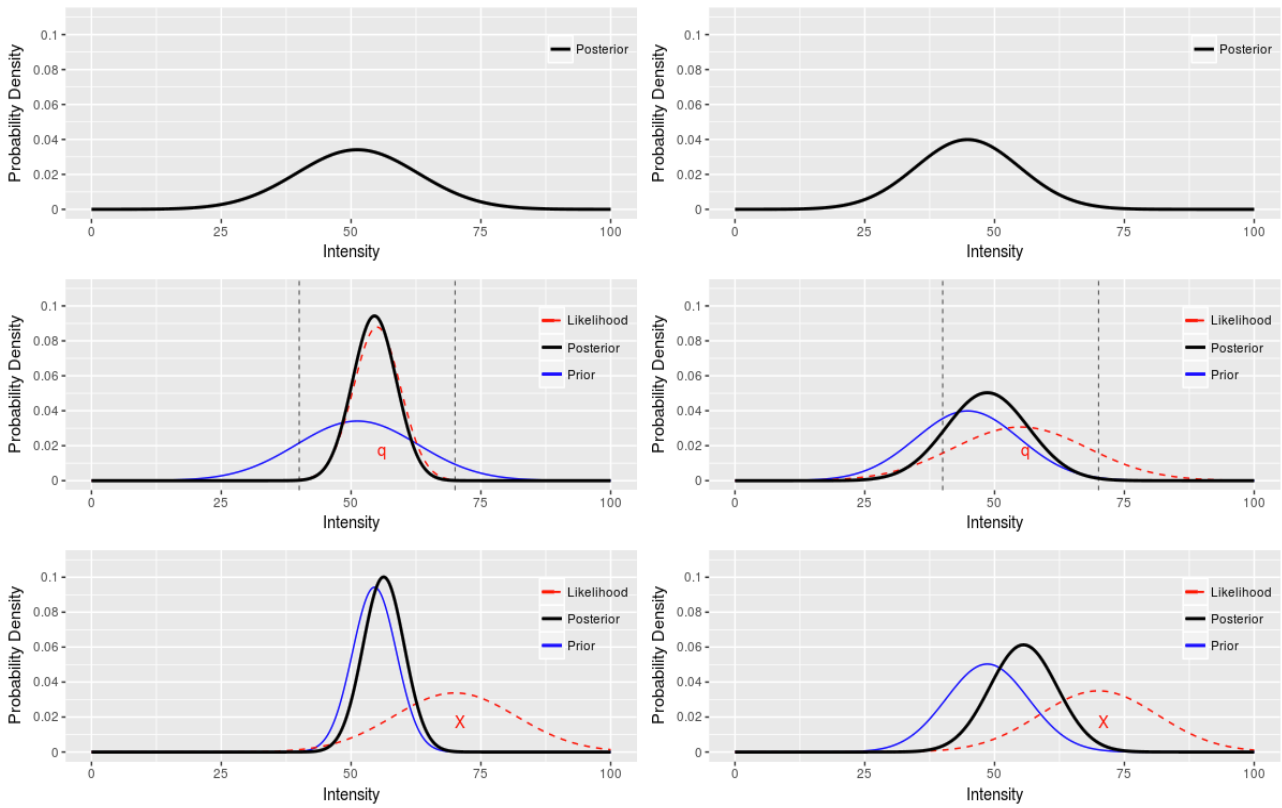

Figure 3: Illustration of the way the model explains the responses of two different participants in Experiment 1 after exposure to the same cue. *We selected a trial in which the cue consisted of two possible pain levels (4 and 7), each with 50% probability. The horizontal axis describes pain levels on a percent scale, from 0 to 100. The two vertical bars indicate the possible levels of pain according to the cue (40 and 70). The symbol "q" marks the mean of the expected pain, based on the cue (here, 55). The symbol "X" marks the intensity* X_{ij} *of the delivered stimulus. Participants 3 and 4 are respectively represented in the left and right columns. The top row illustrates cue-independent expectation. In the middle row, the likelihood of the cue (red dashed curve) is incorporated to create a posterior (black curve) which will later serve as the 'prior distribution'* $P\ (R_{ij}|Z_{ij})$*. In the last row the 'prior distribution' (now the blue curve) is incorporated with the likelihood of the stimulus (red dashed curve) leading to the posterior distribution (black curve),* $P(R_{ij}|\ X_{ij}, Z_{ij})$.

Figure 4 illustrates the variability in cue-independent priors among participants in Experiment 1. Here the mean and standard deviation of the curves correspond with the individual parameters μ and ν

respectively. We see that the cue-independent priors differ in location and strength leading to differences in how the cue shapes the 'Prior Distribution'.

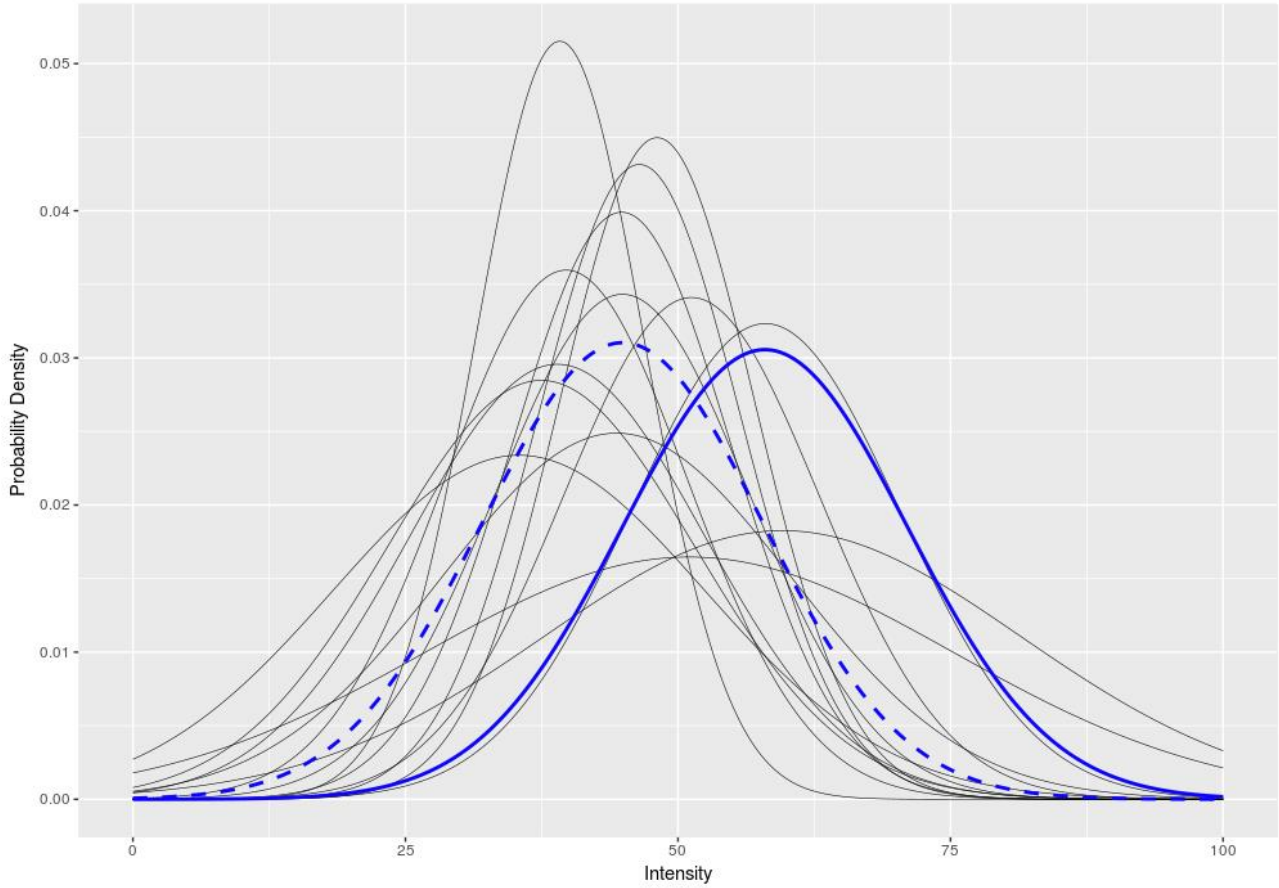

Figure 4: Individual-level, cue-independent priors representing trait-like bias*. The graphs represent the cue-independent priors for participants in Experiment 1 under the 'Variable full Bayesian model' (Model 5). For example, the bold blue distribution represents a participant who is more pessimistic about pain than the dashed blue distribution in terms of the mean pain that they expect. The distributions vary in both their mean and variance with the variance determining how strong the posterior will be biased towards the mean.*

Figure 5 plots the individual-specific estimates of the parameters obtained in the 'Variable full Bayesian model' (Model 5) for each experiment. A clear-cut ranking of the individuals by sensitivity is evident in the plots. The estimates of parameter η appears on average to be less precise than the remaining parameters, as evident from their larger credible intervals. The individual-level estimates of the remaining parameters

μ, ν, β and ρ, have narrower credible intervals, and less overlap across individuals, which suggests these parameters could provide a good basis for classifying individuals. This suggests that individuals can be accurately ranked on the basis of the latter set of parameters.

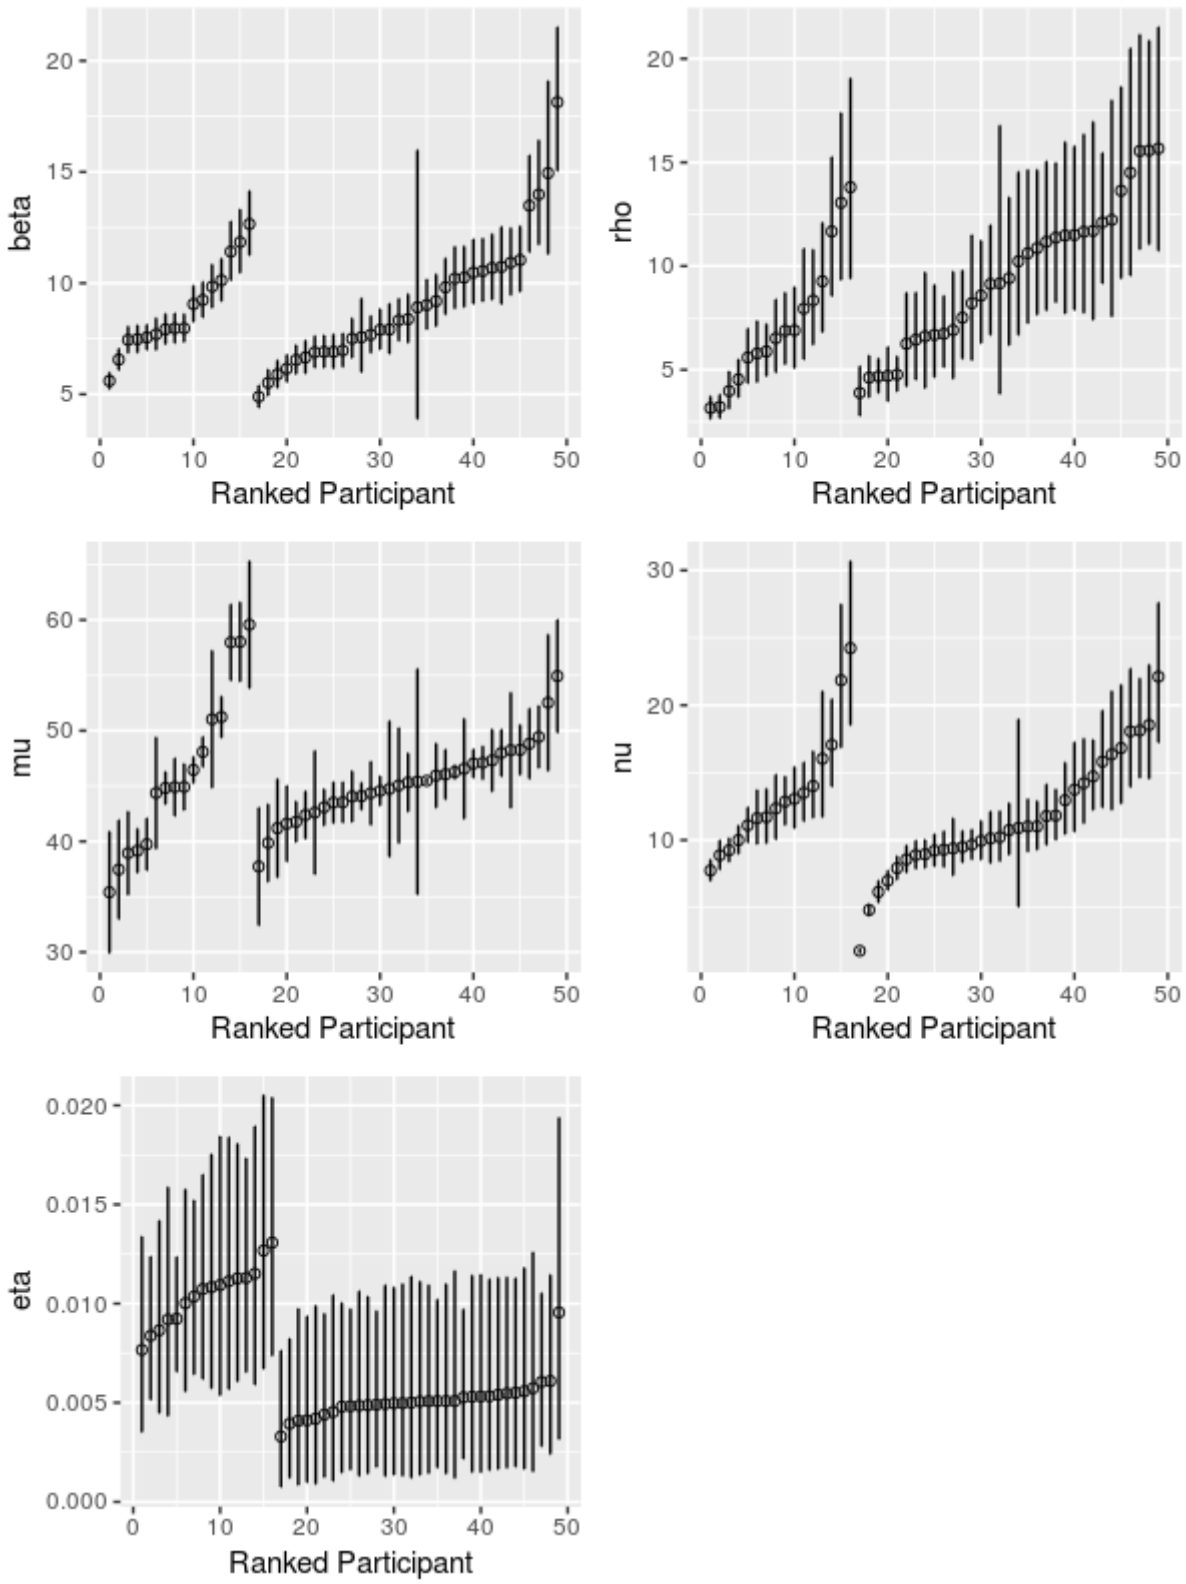

Figure 5: Individual-level estimates for parameters in the Variable full Bayesian model (model 5). *There is one plot per parameter (see vertical axis label). Each plot is divided in two sections, the left section reporting the ranked estimates of the corresponding parameters for the participants in Experiment 1, and the right section reporting the ranked estimates for the participants in Experiment 2. These estimates were obtained by separately analysing the data from Experiments 1 and 2 on the basis of Model 5. Each individual in each plot is represented by a circle, whose y-coordinate corresponds to the point estimate of the parameter for that individual. Superimposed to each circle is the 95% Bayesian credible interval for the corresponding estimate. The left (respectively, right) portion of each plot refers to Experiment 1 (respectively, 2). The ordering of the individuals along the horizontal axis corresponds to their ranking by increasing value of the parameter estimate.*

4. Discussion

Five models of pain perception were tested using data collected from two experiments that utilised a short, novel and deception-free experimental paradigm. A model comparison performed using data from the first experiment supported a number of hypotheses. Firstly, it was found that increased variance in pain expectations decreases the influence that those pain expectations exert on the perception of pain. Secondly, it was shown that adding parameters that describe cue-independent biases in pain perception significantly improved model fit, suggesting that such biases produce an effect on pain perception that is independent of cued pain expectations. Thirdly, as the best fitting 'Full Bayesian model' was influenced by the Hierarchical Predictive Processing (HPP) framework, the results of this study lend support for the HPP model of perception. Finally, this winning model produced a significantly better fit of the data when its parameters were allowed to vary at an individual level, suggesting that there are substantial individual differences in sensitivity to aspects of pain expectations. We also show that this individual-level model can be used to distinguish individuals based on the characteristics of their pain expectations. Importantly, the

above model comparison results were replicated when the analysis was repeated using data from a second experiment which involved an independent sample.

The comparison between the fits achieved by the models presented in this paper allowed an assessment of the viability of the HPP framework to be made. The fact that Model 2 outperformed Model 1 replicated the typical finding the expectation can alter pain perception (e.g. Atlas & Wager, 2012; Bingel et al., 2011; Colloca, & Benedetti, 2006; Peerdeman et al., 2016; Tracey, 2010). More interestingly, the fact that Model 3 outperformed Model 2 confirms that the uncertainty around the expected pain value modulates its influence on pain perception (Brown, Seymour, Boyle, El-Deredy, & Jones, 2008b; Colloca, Petrovic, Wager, Ingvar, & Benedetti, 2010). Given that the model parameter which captures the effect of expectation uncertainty (η) was found to have a positive value, it can be concluded that greater uncertainty in the pain expectation decreases its effect on perception. This finding aligns with the predictions of HPP, but is in contrast to previous work which has suggested that greater expectation uncertainty may increase the influence of expectation (Watkinson et al., 2013) or may increase the level of perceived pain independent of the expectation (Yoshida et al., 2013). The reason for these divergent findings needs to be investigated in future research. It is possible that the different types of pain stimuli used (pressure by Watkinson et al., and temperature by Yoshida et al.,) may generate different results to those achieved via the electrical pain stimulation used in the current study. Indeed, although expectations of electric pain and laser pain influence ERP markers of pain delivery similarly (Hird, Jones, Talmi, & El-Deredy, 2018), their effect on neural markers of pain *anticipation* differs (Babiloni et al., 2007; Hird et al., 2018). Alternatively, as Yoshida et al. used a cue which involved seeing ratings in the form of a number of 'ticks' on a visual-analogue scale, such that greater uncertainty was operationalised in ticks that were more spread out, there is a possibility that participants may have misinterpreted the meaning of this display, perhaps thinking that a larger spread implies greater pain intensity. It is worth noting that of all the parameters in the 'Full Bayesian model', the estimates of parameter η appeared to be less precise than the other parameters when calculated at an individual level (Figure 5.). The values of η also differed noticeable between the two experiments. This may suggest that further research is needed to understand exactly how expectation

uncertainty maps onto pain experience. Nevertheless, our results corroborate HPP in the domain of pain perception, in accordance with previous claims in favour of the HPP theory in the same and in other modalities (Clark et al., 2008; Krol & El-Deredy, 2011; Krol & El-Deredy, 2015).

Another finding of interest in relation to the cognitive aspects of pain expectation was that Model 4 outperformed Model 3, a result which suggests that cue-independent (i.e. trait-like) differences in pain expectation have a significant effect on pain perception over and above the effect of cue-dependent pain expectations. The way we modelled the effect of cue-independent expectations, as shaping the prior distribution that was informed by the cue, was inspired by work on cognitive expectations where optimism bias modulates participants' ability to learn from the information they are given (Garrett & Sharot, 2017). Experimental studies of the pain perception process have often ignored the influence of expectations that exist outside of those generated by the experimental paradigm, which in clinical settings are known to contribute greatly to pain experience. In the case of chronic pain patients such expectations are likely to contribute more to perceived pain than more cue-dependent expectations. Because pessimistic expectations about procedure pain could prevent patients from taking up critical preventative treatments, such as colonoscopy (Trevisani, Zelante, & Sartori, 2014) future work on pain expectation therefore needs to consider the influence that individual differences in cue-independent pain expectations have on pain perception.

While much previous research has supported HPP's account of pain perception (e.g. Brown et al., 2008a; Brown et al., 2008b; Brown, El-Deredy, & Jones, 2014; Clark, 2013) most of this previous research has been carried out at a population level, without paying attention to individual-level differences in the effects of expectation. To counter this shortcoming a version of 'Full Bayesian model' was created which allowed the model parameters to vary from one individual to another. This individual-variant model was found to fit the experimental data significantly better than its fixed-parameter equivalent, suggesting that the model and its parameters were able to characterise each individual participant in terms of their sensitivity to pain expectations. This finding further suggests that these parameters capture important aspects of each individual's pattern of response to pain. For example parameters ρ, η and β quantify the individual's

sensitivity to aspects of the cue (expectation precision, its variance, and the precision of their representation of the pain stimulation). Individuals with a low value of ρ and a high value of β may be more sensitive to the specific information they receive about upcoming pain, so they may be more inclined to dismiss their own physical sensations in favour of expectations generated by the cue. There are hints that patients with medically-unexplained pain complaints may belong to this group (Van den Bergh, Witthöft, Petersen, & Brown, 2017). In a clinical setting, such individuals might benefit most from psychological pain therapies that concentrate on information processing biases relating to pain, as well as from training methods to increase their somatosensory sensitivity (Huque, Poliakoff, & Brown, 2017). In contrast, parameters μ, *and* v, capture trait-like differences in pain expectation, which are independent of contextual information (e.g. the cue information). The value of parameter μ specifies the mean pain expected before the effect of any contextual information is considered (i.e. the cue in the experimental paradigm used here) while parameter v allows us to set apart individuals with exceptionally low or high confidence in this cue-independent expectation. These cue-independent parameters can be thought of as analogous to an optimism/pessimism trait relating to future pain. In a clinical setting, individuals with a high value of μ and a low value of v may be more likely to benefit from psychological pain therapies that concentrate on general beliefs about pain.

The individual differences that the methodology presented in this paper quantifies go beyond what is available through existing self-report instruments, such as the well-established pain catastrophizing scale (Sullivan, Bishop, & Pivik, 1995). In chronic pain, high PCS scores are associated with reports of more severe pain and reduced efficacy of physical and psychological therapies (Keefe et al., 1989). PCS and other self-report measures offer a fast, reasonably accurate, but ultimately gross and subjective, description of pain expectation and experience. They cannot explain what aspects of expectations drive changes in pain experience. Specifically, the same score on a pain scale, given by two patients, could arise through very different psychological processes. Arguably, the method described here to characterise individuals' sensitivity to expectation goes further to reveal the underlying cognitive mechanism which drive the impact of expectations on pain experience.

The ability of the methodology described in this paper to characterise the effect of pain expectations at the individual level could provide useful information in assigning treatment plans if it were applied in a clinical setting. As patients are likely to be differentially sensitive to expectations (Brown, El-Deredy, & Jones, 2014) a method, such as the one presented in this paper, which can potentially identify which patients are most sensitive to pain expectations, and which facets of pain expectations they are sensitive to, could help inform clinical decisions. For example, pain expectations are a key target in non-pharmacological therapies for pain, such as Cognitive Behavioural Therapy (CBT). CBT is effective in reducing chronic pain and pain-related disability partly because it changes patients' pain expectations (Bingel et al., 2011). Unfortunately, CBT is not as effective as we would wish: more than half of patients with chronic pain do not benefit from CBT (Peerdeman et al., 2016; Smith et al., 2016; de Williams, Eccleston, & Morley, 2012). Put simply, if a patient can be identified as being biased to report stronger pain they may be a better candidate for CBT than a patient who reports the delivered pain more accurately; and their specific expectation sensitivity profile can further improve the specific therapeutic intervention. It is hoped that the method described in this paper may ultimately prove useful in selecting suitable candidates for pain management treatment, adding quantitative information that is less affected by introspective limitations to existing self-report questionnaires that assess pain expectations. This could facilitate the personalisation of treatment selection, an approach that is particularly important in the current climate when health services are stretched, and access to effective treatment is limited (Enck et al., 2013). In addition to providing input to clinical decisions, the methodology presented here could also be used in clinical trial research to attempt to control for placebo effects when selecting control and experimental groups. To further the goal of providing a clinical tool using our approach, the current research needs to be repeated in a variety of clinical pain populations. Such work would be required in order to establish the clinical relevance of the individual parameters of pain expectations that have been established in this study.

One limitation of our design is that the pain intensity participants received was, on average, not very different from the pain intensity that they had expected. This may be important, because HPP resembles the assimilative part of assimilation-contrast theory (Hoyland, Harvey, & Sherif, 1957), which predicts that

the distance between prior expectations and the physical stimulus is crucial. Both theories agree that when the distance is small, namely, when the input is not too discrepant from our existing anchor, we will assimilate it and our experience will reflect a 'compromise' between the reality and our expectations. However, when the reality is very different to our expectations assimilation-contrast theory predicts that we will experience the real, discrepant stimulus as even more discrepant compared to our expectations. Our data does not currently allow us to decide whether pain experiences show such contrast effects. An additional limitation of our design is that we are currently not able to distinguish between participants who rated what they 'really felt' and participants who rated their pain as closer to the values of the cue because of demand characteristics. The possible influence of demand characteristics is prevalent in studies of pain experience across paradigms, including placebo designs and cue-based designs, so this caveat is not unique to our design. Future experiments using our experimental paradigm could explore this issue further, for example by manipulating the working memory load of participants, which reduce goal-directed behaviour (Otto, Gershman, Markman, & Daw, 2013).

In summary, model comparison and validation yielded a number of important findings. The HPP framework was shown to provide an accurate description of pain perception. In particular the framework's predictions relating to the effect of expectation uncertainty were supported. Trait-like pain expectations that exist independent of any cued pain information were shown to have an influence on pain perception that was separate and distinguishable from those provided by cued pain expectations. Finally the methodology presented in this paper was able to characterise an individual's sensitivity to pain expectations, potentially providing the basis for a clinical tool that could guide practitioners who deliver individual psychological pain management interventions.

Acknowledgements

CB was partially supported by the FP7 MIMOmics European Collaborative Project, as part of the HEALTH-2012-INNOVATION scheme and by the Medical Research Council. RH and DT were partially supported by ESRC First Grant ES/I010424/1. WED acknowledges the support of CONICYT, Chile, FONDECYT project 1161378 and Basal project FB0008. We thank A. Jones for inspiring this project, N. Begum and E. Hird for helpful suggestions, and A. Sparkes for collecting some of the data reported in this paper.

All authors declare no conflict of interest.

References

Anchisi, D., & Zanon, M. (2015). A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. *PloS one*, *10*(2), e0117270.

Atlas, L. Y., Bolger, N., Lindquist, M. A., & Wager, T. D. (2010). Brain mediators of predictive cue effects on perceived pain. *Journal of Neuroscience*, *30*(39), 12964-12977.

Atlas, L. Y., & Wager, T. D. (2012). How expectations shape pain. *Neuroscience Letters, 520*(2), 140-148. doi: DOI 10.1016/j.neulet.2012.03.039

Babiloni, C., Brancucci, A., Capotosto, P., Del Percio, C., Romani, G. L., Arendt-Nielsen, L., & Rossini, P. M. (2007). Different modalities of painful somatosensory stimulations affect anticipatory cortical processes: a high-resolution EEG study. *Brain research bulletin*, *71*(5), 475-484.

Bhat, A. A., DeWalt, D. A., Zimmer, C. R., Fried, B. J., & Callahan, L. F. (2010). The role of helplessness, outcome expectation for exercise and literacy in predicting disability and symptoms in older adults with arthritis. *Patient education and counselling*, *81*(1), 73-78.

Bingel, U., Wanigasekera, V., Wiech, K., Ni Mhuircheartaigh, R., Lee, M. C., Ploner, M., & Tracey, I. (2011). The effect of treatment expectation on drug efficacy: imaging the analgesic benefit of the opioid remifentanil. [Clinical Trial]. *Sci Transl Med, 3*(70), 70ra14. doi: 10.1126/scitranslmed.3001244

Brown, C. A., Seymour, B., El-Deredy, W., & Jones, A. K. (2008a). Confidence in beliefs about pain predicts expectancy effects on pain perception and anticipatory processing in right anterior insula. *Pain, 139*(2), 324-332. doi: 10.1016/j.pain.2008.04.028

Brown, C. A., Seymour, B., Boyle, Y., El-Deredy, W., & Jones, A. K. (2008b). Modulation of pain ratings by expectation and uncertainty: Behavioral characteristics and anticipatory neural correlates. *Pain, 135*(3), 240-250. doi: 10.1016/j.pain.2007.05.022

Brown, C. A., El-Deredy, W., & Jones, A. K. (2014). When the brain expects pain: common neural responses to pain anticipation are related to clinical pain and distress in fibromyalgia and osteoarthritis. *European Journal of Neuroscience*, *39*(4), 663-672.

Buchel, C., Geuter, S., Sprenger, C., & Eippert, F. (2014). Placebo Analgesia: A Predictive Coding Perspective. *Neuron, 81*(6), 1223-1239. doi: DOI 10.1016/j.neuron.2014.02.042

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci, 36*(3), 181-204. doi: 10.1017/S0140525X12000477

Clark, J. A., Brown, C. A., Jones, A. K., & El-Deredy, W. (2008). Dissociating nociceptive modulation by the duration of pain anticipation from unpredictability in the timing of pain. *Clin Neurophysiol, 119*(12), 2870-2878. doi: 10.1016/j.clinph.2008.09.022

Colloca, L., & Benedetti, F. (2006). How prior experience shapes placebo analgesia. *Pain, 124*(1-2), 126-133. doi: DOI 10.1016/j.pain.2006.04.005

Colloca, L., Petrovic, P., Wager, T. D., Ingvar, M., & Benedetti, F. (2010). How the number of learning trials affects placebo and nocebo responses. *Pain®*, *151*(2), 430-439.

Enck, P., Bingel, U., Schedlowski, M., & Rief, W. (2013). OPINION The placebo response in medicine: minimize, maximize or personalize? *Nature Reviews Drug Discovery, 12*(3), 191-204. doi: Doi 10.1038/Nrd3923

Friston, K. (2003). Learning and inference in the brain. *Neural Netw, 16*(9), 1325-1352. doi: 10.1016/j.neunet.2003.06.005

Garrett, N., & Sharot, T. (2017). Optimistic update bias holds firm: Three tests of robustness following Shah et al. *Consciousness and cognition*, *50*, 12-22.

Geuter, S., Boll, S., Eippert, F., & Büchel, C. (2017). Functional dissociation of stimulus intensity encoding and predictive coding of pain in the insula. *eLife*, *6*, e24770.

Hird, E. J., El-Deredy, W., Jones, A., & Talmi, D. (2017). Temporal dissociation of salience and prediction error responses to appetitive and aversive taste. *Psychophysiology*.

Hird, E. J., Jones, A. K. P., Talmi, D., & El-Deredy, W. (2018). A comparison between the neural correlates of laser and electric pain stimulation and their modulation by expectation. *Journal of Neuroscience Methods*, *293*, 117-127.

Hjermstad, M.J., Fayers, P.M., Haugen, D.F., Caraceni, A., Hanks, G.W., Loge, J.H., Fainsinger, R., Aass, N., Kaasa, S. and European Palliative Care Research Collaborative (EPCRC), (2011). Studies comparing Numerical Rating Scales, Verbal Rating Scales, and Visual Analogue Scales for assessment of pain intensity in adults: a systematic literature review. *Journal of pain and symptom management*, *41*(6), pp.1073-1093.

Hovland, C. I., Harvey, O. J., & Sherif, M. (1957). Assimilation and contrast effects in reactions to communication and attitude change. *The Journal of Abnormal and Social Psychology*, *55*(2), 244.

Huque, A., Poliakoff, E., & Brown, R. J. (2017). Effects of learning on somatosensory decision-making and experiences. *Journal of Experimental Psychology: General*, *146*(11), 1631.

Jung, W. M., Lee, Y. S., Wallraven, C., & Chae, Y. (2017). Bayesian prediction of placebo analgesia in an instrumental learning model. *PloS one*, *12*(2), e0172609.

Keefe, F. J., Brown, G. K., Wallston, K. A., & Caldwell, D. S. (1989). Coping with rheumatoid arthritis pain: catastrophizing as a maladaptive strategy. *Pain*, *37*(1), 51-56

Krol, M., & El-Deredy, W. (2015). The clash of expectancies: Does the P300 amplitude reflect both passive and active expectations? *Q J Exp Psychol (Hove)*, 1-12. doi: 10.1080/17470218.2014.996166

Krol, M. E., & El-Deredy, W. (2011). When believing is seeing: The role of predictions in shaping visual perception. *Quarterly Journal of Experimental Psychology, 64*(9), 1743-1771. Doi 10.1080/17470218.2011.559587

Maloney, L. T., & Zhang, H. (2010). Decision-theoretic models of visual perception and action. *Vision research*, *50*(23), 2362-2374.

Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, M. & Teller, E. (1953) Equations of state calculations by fast computing machines. *Journal of Chemical Physics,* 21:1087–1092.

Neal. R. (2011) MCMC using hamiltonian dynamics. In Gelman A. Jones G. L. Brooks, S. and X.L. Meng, editors, *Handbook of Markov Chain Monte Carlo,* pages 116–162. Chapman and Hall/CRC.

Otto, A. R., Gershman, S. J., Markman, A. B., & Daw, N. D. (2013). The curse of planning: dissecting multiple reinforcement-learning systems by taxing the central executive. *Psychological science*, *24*(5), 751-761.

Parducci, A. (1963). Range-frequency compromise in judgment. *Psychological Monographs: General and Applied*, *77*(2), 1.

Peerdeman, K. J., van Laarhoven, A. I., Keij, S. M., Vase, L., Rovers, M. M., Peters, M. L., & Evers, A. W. (2016). Relieving patients' pain with expectation interventions: a meta-analysis. *Pain*, *157*(6), 1179-1191.

Ploghaus, A., Tracey, I., Gati, J. S., Clare, S., Menon, R. S., Matthews, P. M., & Rawlins, J. N. (1999). Dissociating pain from its anticipation in the human brain. *Science, 284*(5422), 1979-1981.

Schrooten, M. G., Vlaeyen, J. W., & Morley, S. (2012). Psychological interventions for chronic pain: reviewed within the context of goal pursuit. *Pain management*, *2*(2), 141-150.

Seymour, B., O'Doherty, J.P., Dayan, P., Koltzenburg, M., Jones, A.K., Dolan, R.J., Friston, K.J. and Frackowiak, R.S., (2004). Temporal difference models describe higher-order learning in humans. *Nature*, *429*(6992), pp.664-667.

Smith, J. G., Knight, L., Stewart, A., Smith, E. L., & McCracken, L. M. (2016). Clinical effectiveness of a residential pain management programme–comparing a large recent sample with previously published outcome data. *British Journal of Pain*, *10*(1), 46-58.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(4), 583-639.

Stan Development Team (2014a). RStan, version 2.2. http://mc-stan.org/rstan.html.

Stan Development Team (2014b). Stan: A C++ library for probability and sampling, version 2.2. http://mc-stan.org/,2014.

Sullivan, M. J., Bishop, S. R., & Pivik, J. (1995). The pain catastrophizing scale: development and validation. *Psychological assessment*, *7*(4), 524.

Tabor A, Thacker MA, Moseley GL, Körding KP (2017) Pain: A Statistical Account. PLoS Comput Biol13(1): e1005142. https://doi.org/10.1371/journal.pcbi.1005142

Talmi, D., Atkinson, R., & El-Deredy, W. (2013). The Feedback-Related Negativity Signals Salience Prediction Errors, Not Reward Prediction Errors. *Journal of Neuroscience, 33*(19), 8264-8269. Doi 10.1523/Jneurosci.5695-12.2013

Tracey, I. (2010). Getting the pain you expect: mechanisms of placebo, nocebo and reappraisal effects in humans. *Nat Med, 16*(11), 1277-1283. doi: 10.1038/nm.2229

Trevisani, L., Zelante, A., & Sartori, S. (2014). Colonoscopy, pain and fears: Is it an indissoluble trinomial? *World J Gastrointest Endosc, 6*(6), 227-233. doi: 10.4253/wjge.v6.i6.227

Van den Bergh, O., Witthöft, M., Petersen, S., & Brown, R. J. (2017). Symptoms and the body: Taking the inferential leap. *Neuroscience & Biobehavioral Reviews*.

Wager, T.D., Rilling, J.K., Smith, E.E., Sokolik, A., Casey, K.L., Davidson, R.J., Kosslyn, S.M., Rose, R.M. and Cohen, J.D., (2004). Placebo-induced changes in FMRI in the anticipation and experience of pain. *Science*, *303*(5661), pp.1162-1167.

Wainwright, M.J. & Jordan, M.I. (2008) Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305.

Watkinson, P., Wood, A. M., Lloyd, D. M., & Brown, G. D. (2013). Pain ratings reflect cognitive context: A range frequency model of pain perception. *PAIN*, *154*(5), 743-749.

Watson, A., El-Deredy, W., Iannetti, G. D., Lloyd, D., Tracey, I., Vogt, B. A., . . . Jones, A. K. (2009). Placebo conditioning and placebo analgesia modulate a common brain network during pain anticipation and perception. *Pain, 145*(1-2), 24-30. doi: 10.1016/j.pain.2009.04.003

de Williams, A. C., Eccleston, C., & Morley, S. (2012). Psychological therapies for the management of chronic pain (excluding headache) in adults. *Cochrane Database of Systematic Reviews*.

Yoshida, W., Seymour, B., Koltzenburg, M., & Dolan, R. J. (2013). Uncertainty increases pain: evidence for a novel mechanism of pain modulation involving the periaqueductal gray. *The Journal of Neuroscience*, *33*(13), 5638-5646.

Supplementary material

Supplementary Table 1: Voltage delivered for each subjective pain level in Experiments 1 and 2.

Voltage (in Volts) delivered for each subjective pain level		
Pain Level	**Experiment 1: Mean(SD)**	**Experiment 2: Mean(SD)**
Pain Level 3	1.34 (.46)	1.59 (.17)
Pain Level 4	1.69 (.51)	1.90 (.20)
Pain Level 5	2.03 (.61)	2.23 (.22)
Pain Level 6	2.35 (.67)	2.56 (.25)
Pain Level 7	2.74 (.84)	2.90 (.29)

Supplementary Table 2: Trial types in Experiment 1.

Possible pain levels, depicted in the cue (their probabilities)	Mean expected pain	Delivered pain levels	Size of Prediction error at pain delivery	Number of repetitions per block
Cues with 1 option				
3 (100%)	3	3	0	1
4 (100%)	4	4	0	1
5 (100%)	5	5	0	1
6 (100%)	6	6	0	1
Cues with 2 options				
4 (50%) 5 (50%)	4.5	4	-0.5	1
		5	+0.5	1
3 (25%) 5 (75%)	4.5	3	-1.5	1
		5	+0.5	3
4 (75%) 6 (25%)	4.5	4	-0.5	3
		6	+1.5	1
5 (50%) 6 (50%)	5.5	5	-0.5	1
		6	+0.5	1

4 (25%) 6 (75%)	5.5	4	-1.5	1
		6	+0.5	3
5 (75%) 7 (25%)	5.5	5	-0.5	3
		7	+1.5	1
Cues with 4 options				
3 (25%) 4 (25%) 5 (25%) 6 (25%)	4.5	3	-1.5	1
		4	-0.5	1
		5	+0.5	1
		6	+1.5	1
4 (25%) 5 (25%) 6 (25%) 7 (25%)	5.5	4	-1.5	1
		5	-0.5	1
		6	+0.5	1
		7	+1.5	1

Note. 16 trial types correspond to a 2 (Expected pain level: 4.5, 5.5) x 2 (Discrepancy of expectations from delivered pain: 0.5,1.5) x 2 (Direction of discrepancy: positive, negative) x 2 (The number of possible pain levels participants could expect: 2, 4). In order to implement this design, 4 additional trial types were included; those are presented in light shaded boxes. For the purpose of a manipulation check, 4 further trial types with 100% probability of receiving a particular pain level were also included, presented in dark shaded boxes.

Supplementary Table 3. Trial types in Experiment 2.

Possible pain levels, depicted in the cue (their probabilities)	Mean expected pain	Delivered pain level	Size of prediction error at pain delivery
Cues with 1 option			
3 (100%)	3	3	0
4 (100%)	4	4	0
5 (100%)	5	5	0
6 (100%)	6	6	0
Cues with 2 options			
4 (50%) 5 (50%)	4.5	4	-0.5
		5	+0.5
3 (50%) 6 (50%)	4.5	3	-1.5
		6	+1.5
5 (50%) 6 (50%)	5.5	5	-0.5
		6	+0.5
4 (50%) 7 (50%)	5.5	4	-1.5
		7	+1.5

4 (50%) 6 (50%)	5	4	-1
		6	+1
3 (50%) 7 (50%)	5	3	+2
		7	-2

Note. 8 trial types correspond to a 2 (Expected pain level: 4.5, 5.5) x 2 (Discrepancy of expectation from delivered pain: 0.5, 1.5) x 2 (Direction of discrepancy: positive, negative). 4 trial types presented in shaded boxes correspond to a 2 (discrepancy of expectations from delivered pain) x 2 (direction of discrepancy: positive, negative). 4 further trial types where cues offered 100% certainty about the impending pain level were included for the purpose of a manipulation check. All trial types were delivered once per block.

Supplementary Table 4: Rated pain experience for cues with 100% certainty in Experiments 1 and 2.

Descriptive statistics for ratings of 100% chance trials at each pain level		
Pain level	**Experiment 1: Mean (SD)**	**Experiment 2: Mean (SD)**
Pain level 3	24.30 (8.47)	26.43 (15.92)
Pain level 4	36.96 (10.37)	35.25 (15.16)
Pain level 5	49.96 (13.47)	44.59 (14.80)
Pain level 6	60.99 (16.76)	52.72 (15.96)

Note: Mean (SD) of ratings on a scale of 1-100

Appendix

Derivation 1: Derivation of the 'prior distribution'

Here we derive $P(R_{ij}|\, Z_{ij})$ as resulting from a Bayesian updating process and motivate its use as the 'prior distribution'. Following from Bayes' theorem, the 'prior distribution' is proportional to the product of a cue-independent prior and a likelihood function of the cue:

$$P(R_{ij}|\, Z_{ij}) \propto P\,(Z_{ij}|\, R_{ij}) \times P(R_{ij})$$

We choose the cue-independent prior to take the form:

$$P(R_{ij}) = (2\pi v)^{-\frac{1}{2}} exp(-\frac{1}{2\upsilon^2}\,(R_{ij} - \mu)^2)$$

This prior is independent of the cue, and represents a stable, trait-like bias to expect pain of a certain intensity. Higher values of parameter μ (the mean of the cue-independent prior) reflect the expectation of higher levels of pain, independent of any cue information. Higher values of parameter υ (the standard deviation of the cue-independent prior) reflect greater uncertainty (and therefore lesser influence) of this cue-independent prior. Participants with high μ and low υ can be thought of as pessimistic about pain.

We choose the likelihood function of the cue to take the form:

$$P\,(Z_{ij}|\, R_{ij}) \propto (2\pi * (\rho + \eta \times sd_{ij}))^{-\frac{1}{2}} exp(-\frac{1}{2 * (\rho + \eta \times sd_{ij})^2}(q_{ij} - R_{ij})^2)$$

This likelihood function incorporates two parameters which will eventually modulate the effect of the cue information Z_{ij} on the likelihood of the rating R_{ij}. The parameter ρ controls the variance of the likelihood function generated by the cue. In effect it represents the precision of this prior. In contracts η dictates the influence the cue's variance has on the sharpness of the likelihood function. Together these parameters determine the extent to which the cue information modulates the prior distribution, and thus the subsequent experience of pain. The influence of the cue increases with lower values of ρ and η.

Plugging these functions back into the original equation and now conditioning on Z_{ij} we can express the prior distribution as:

$$P(R_{ij}|\, Z_{ij}) \propto Normal(q_{ij}, \rho + \eta \times sd_{ij}) \times Normal(\mu, \upsilon)$$

Derivation 2: Incorporation of the Likelihood of the Stimulation

The posterior is given, up to a proportionality constant, by the product of the prior and the likelihood of the stimulation:

$$P(\, R_{ij}|\, X_{ij}, Z_{ij}) \propto P\,(X_{ij}|R_{ij}, Z_{ij}) \times P\,(R_{ij}|Z_{ij})$$

We take the likelihood of the stimulation to factor into two components, $P\,(X_{ij}|R_{ij}, Z_{ij}) = f_1\,(X_{ij}, R_{ij}) \times f_2\,(X_{ij}, Z_{ij})$. Since we consider the posterior only up to a proportionality constant, we are only concerned with the $f_1\,(X_{ij}, R_{ij})$ component which we take to be normally distributed:

$$f_1 = (2\pi\beta)^{-\frac{1}{2}} exp(-\frac{1}{2\beta^2}(X_{ij} - R_{ij})^2)$$

with $\beta > 0$. Each predicted value of R_{ij} is compared with the external sensory information, associating each such value with a measure of 'surprise', which our model takes to increase with an increasing distance between R_{ij} and the actual intensity of the stimulation, X_{ij}. The value of β determines the extent that the perceptual mechanism will tolerate a given disparity between the predicted experience and the actual sensation, with smaller values of β implying greater sharpness of the likelihood function around X_{ij} and hence a reduced tolerance.

Plugging $f_1$back into the original equation and now conditioning on X_{ij} we can express the posterior distribution as:

$$P(\, R_{ij}|\, X_{ij}, Z_{ij}) \propto Normal(X_{ij}, \beta) \times P\,(R_{ij}|Z_{ij})$$