

Published in final edited form as:

J Neurosci. 2011 March 30; 31(13): 4811–4820. doi:10.1523/JNEUROSCI.1452-10.2011.

The known unknowns: neural representation of second-order uncertainty, and ambiguity

Dominik R. Bach*, Oliver Hulme, William D. Penny, and Raymond J. Dolan

Wellcome Trust Centre for Neuroimaging, University College London, 12 Queen Square, London WC1N 3BG, United Kingdom

Abstract

Predictions provided by action-outcome probabilities entail a degree of (first-order) uncertainty. However, these probabilities themselves can be imprecise and embody second-order uncertainty. Tracking second-order uncertainty is important for optimal decision making and reinforcement learning. Previous functional magnetic resonance imaging investigations of second-order uncertainty in humans have drawn on an economic concept of ambiguity, where action-outcome associations in a gamble are either known (unambiguous) or completely unknown (ambiguous). Here, we relaxed the constraints associated with a purely categorical concept of ambiguity and varied the second-order uncertainty of gambles continuously, quantified as entropy over second-order probabilities. We show that second-order uncertainty influences decisions in a pessimistic way by biasing second-order probabilities, and that second-order uncertainty is negatively correlated with posterior cingulate cortex activity. The category of ambiguous (compared to non-ambiguous) gambles also biased choice in a similar direction, but was associated with distinct activation of a posterior parietal cortical area; an activation that we show reflects a different computational mechanism. Our findings indicate that behavioural and neural responses to second-order uncertainty are distinct from those associated with ambiguity and may call for a reappraisal of previous data.

Keywords

ambiguity; risk; uncertainty; entropy; information theory; ideal observer; Bayesian model comparison

Introduction

Probabilistic action-outcome associations provide uncertain predictions of the future, even when their probabilities are known precisely, which we refer to as *first-order uncertainty*. However, action-outcome probabilities themselves can be endowed with imprecision, that is, *second-order uncertainty*. Understanding how uncertainty is encoded is paramount for all probabilistic accounts of brain function (e.g. Friston, 2010), and tracking second-order uncertainty in particular is indispensable for optimal decision making (Daw et al., 2005), and associative learning (Pearce and Hall, 1980; Friston, 2010).

Economic theory furnishes a means to experimentally manipulate uncertainty in an explicit manner, where first-order uncertainty is usually conceptualised as *risk* (von Neumann and Morgenstern, 1944), and distinguished from second-order uncertainty (Knight, 1921),

*Address correspondence to: Wellcome Trust Centre for Neuroimaging 12 Queen Square London WC1N 3BG United Kingdom
Phone: +44 20 7833 7472 d.bach@fil.ion.ucl.ac.uk.

referred to as *ambiguity* (Ellsberg, 1961). The distinction arises from observations that when winning probabilities of lotteries are known (non-ambiguous) or completely unknown (ambiguous), individuals avoid ambiguous lotteries even when this avoidance is costly (see for a review Camerer, 1999). Furthermore, neural responses to ambiguity in posterior parietal cortex, posterior dorsolateral prefrontal cortex and anterior insula (Huettel et al., 2006; Bach et al., 2009) suggest a representation that is distinct from those supporting first-order uncertainty in economic gambles (Preuschoff et al., 2006; Tobler et al., 2007; Dreher et al., 2006).

However, ambiguous and non-ambiguous lotteries differ in several aspects other than second-order uncertainty. Therefore, it remains unclear whether ambiguity aversion and neural responses to ambiguity reflect second-order uncertainty. Ambiguity aversion is diminished when choices are made privately instead of publicly (Curley et al., 1986; Trautmann et al., 2008), and when unfavourable first-order probabilities are precluded (Larson, 1980; Keren and Gervitsen, 1999). Both ambiguity aversion (Chow and Sarin, 2002) and neural responses to ambiguity (Bach et al., 2009) are greatest when missing information about first-order probabilities is *potentially knowable* (i.e. known to the experimenter, and usable to infer the conditional first-order probabilities) (see also O'Neill and Kobayashi, 2009). These findings suggest that factors other than second-order uncertainty may confound experimental manipulations of ambiguity, such as beliefs about the availability of missing knowledge to oneself and to others, e.g. the idea that one is being cheated, or that one has to take action to get at the missing bit of information.

Hence, in this study we sought to quantify second-order uncertainty in ambiguous economic gambles and establish its specific neural correlates. Additionally, by contrasting ambiguous and non-ambiguous gambles, we could determine whether neuronal responses differed from those evoked specifically by second-order uncertainty. Any difference would argue that categorical manipulations of ambiguity may not serve as a 'pure insertion' of second-order uncertainty (Friston et al., 1996). Experimentally, we depart from standard economic experimental approaches that provide either all or no information about first-order probabilities, and constructed ambiguous gambles, where some information is given in the form of *second-order probabilities* that specify which conditional first-order probability will be realised. By showing participants two coloured balls on the edges of a screen, each denoting a first-order probability of getting electric shocks (see Figure 1), we could vary the amount of this information (the second-order uncertainty) continuously, and quantify it by its Shannon entropy (1948).

Materials and methods

Participants

24 healthy right-handed individuals took part in the behavioural experiment (12 male, 12 female, mean age \pm standard deviation: 25.3 ± 4.5 yrs), and an independent sample of 20 healthy right-handed individuals participated in the imaging experiment (12 male, 8 female, mean age \pm standard deviation: 22.9 ± 2.5 yrs). The absence of visual impairments and neurological or psychiatric symptoms/disorders was assessed via questionnaire. For the imaging experiment, handedness was controlled with the Edinburgh Handedness Inventory (Oldfield, 1971; mean \pm standard deviation: 85.3 ± 22.1). Participants were recruited from the general population by advertisement and were given a monetary compensation of £30 (behavioural experiment) or £40 (imaging experiment). All participants gave written informed consent, and the study was approved by a NHS research ethics committee.

Design

Our experiments took the following general form. Participants learned that the colour of a ball presented on a computer screen was associated with a certain *conditional first-order probability* of receiving three aversive electric shocks (figure 1A). On any trial, participants could either choose to accept a gamble on this ball, or receive a single electric shock of the same intensity. After learning these first-order probabilities, the gamble was rendered ambiguous by presenting two coloured balls at the opposite edges of the screen and a ball silhouette in between (see figure 1B).

Participants were instructed that the two coloured balls corresponded to two “bowling ball players”, one of which would play his ball onto the lane. The ball silhouette was the final position of the played ball, and was more likely to come from either of the two “players”, and thus represent either of the two coloured balls, the closer it was to this ball. To solve this problem optimally, one needs to infer the *second-order probabilities* that either coloured ball would be sampled on the basis of how close the silhouette was to them, then combine these with the two conditional first order probabilities, hence compute the *expected first-order probability* of being shocked, and thus decide whether to accept the gamble or to receive a single electric shock instead.

If the two balls at the edges of the screen have different colours, the second-order probabilities can be inferred when the distributions of the final ball positions are known. These distributions were fixed and indicated by colour bars with varying intensity (see figure 1B). The *second-order uncertainty* is then dependent on the silhouette position and was quantified as Shannon entropy of the second-order probabilities:

$$H = - \sum_{i=1}^2 p(\text{colour}_i | \text{position}) * \log [p(\text{colour}_i | \text{position})]$$

where, by Bayes’ rule and in the absence of prior knowledge

$$p(\text{colour} | \text{position}) = \frac{p(\text{position} | \text{colour})}{\sum_{i=1}^2 p(\text{position} | \text{colour}_i)}$$

One can see from figure 1D that the silhouette is more informative (i.e., has less entropy H) about the colour, the closer it is to one of the coloured balls on the edges of the screen. This manipulation incidentally also varies the expected first order probability which was decorrelated from entropy by using six different combinations/orderings of balls, and was additionally taken into account for analysis (see data analysis).

In this way, we were able to measure acceptance of gambles and associated evoked brain responses as a function of the entropy of the second-order probabilities. This design also allowed us to collapse ambiguous gambles into a factor level to allow a categorical contrast with non-ambiguous gambles where both colour options were the same (figure 1C). In this case, the colour of the silhouette can be inferred with certainty, and the prediction is non-ambiguous. The expected first order probability is then simply the conditional first order probability associated with the colour of both balls.

Thus, the experiment followed an incomplete factorial design with a discrete factor *ambiguity* (ambiguous vs. non-ambiguous) and a continuously varied factor *entropy* within the ambiguous condition. This design incidentally also varied the additional discrete factors

first-order probability (0.2, 0.5, 0.8) within the non-ambiguous and *range of first-order probabilities* (0.3, 0.6) within the ambiguous condition.

Stimuli and gambles

There was one trial for each of 3 ball colours, in combination with the other possible ball colours (*ambiguous trials*: 2 colours; *non-ambiguous trials*: 1 colour), originating position (left or right), and possible ball position (*ambiguous trials*: 23; *non-ambiguous trials*: 14), resulting in 276 ambiguous and 84 non-ambiguous trials, or 360 trials on the whole.

The three ball colours (*black, orange, and blue*) represented three different lotteries. All three lotteries had the same stake (win: nothing; lose: three electric shocks), but different (first-order) probabilities of being shocked (0.2, 0.5, and 0.8). The association of colour and shock probability was balanced across subjects and learnt in a preceding discrimination learning paradigm involving 144 learning trials (that is, 46 trials per ball colour, see procedure).

Each trial started with a variable interval of 2 – 4 s (behavioural experiment) or 1.5 – 3.5 s (imaging experiment), during which a lane and an empty feedback box were visible. Then, two coloured balls were shown. Above the balls, two colour bars with varying intensity indicated the probability that either ball would land on a given position on the lane if played. This followed a linear probability density function with a maximum at the position of this ball and reaching zero at the position of the opposite ball. A grey silhouette appeared somewhere on the lane 0.5 s later (behavioural experiment) or immediately (imaging experiment).

The position of the silhouette was determined as follows: the lane was divided into 23 (ambiguous trials) or 14 (non-ambiguous trials) intervals of equal cumulative landing probability. The interval in which the ball occurred on a given trial was pseudo-randomized and kept constant across all participants. This ensured that the entropy variation was approximately the same for each participant. We randomly drew the exact position within the landing interval from the corresponding probability density function, and this exact position was taken into account for analysis. The grey silhouette also served as a timer, fading away within 1.5 s (behavioural experiment) or 2 s (imaging experiment). Within this time, participants had to make a decision to either accept or reject the lottery by pressing an up or down arrow key on a standard computer keyboard.

The association of up and down button presses and reject or accept was balanced across subjects and kept constant for each participant. It was signalled on each trial (e.g. *up arrow: fixed, down arrow: gamble*), where the response option associated with the up key was always presented on top of the other option to avoid stimulus-response incongruencies. If participants did not make a choice, they received three electric shocks with certainty, and these trials were excluded from analysis. After the decision time was over, we signalled which colour was actually sampled (and thus, the conditional first-order probability) to emphasise that it was pre-determined and consistent (Bach et al., 2009). This was shown in a separate feedback box on the top half of the screen, together with an indication of the participant's choice (*fixed option or gamble*) and an indication of the electric stimulations they would receive (one or three lightning-style electric shock signs, or the words *no shock*). This feedback was visible for 1.5 s. At the same time, electric stimulations were delivered as a 500 Hz train of electrical pulses (square wave, individual pulse duration: 1 ms, total duration: 100 ms, 400 V, mean current \pm standard deviation, behavioural experiment: 0.076 mA \pm 0.044 mA, imaging experiment: 0.073 mA \pm 0.037 mA) via a pin-shaped electrode attached to the left forearm, repeated once (fixed option) or three times with 0.4 s inter stimulus interval (gamble). Stimulation intensity was the same for the fixed option and

gamble, and set in a pre-test assessment as slightly below the pain threshold (see procedure). In 12 of 20 participants in the imaging experiment, 14 random stimuli below the pain threshold were rated for intensity before and after the experiment, and there were no differences in the mean ratings ($p > .50$).

Procedure

Upon arrival in the laboratory, the procedure was explained to participants in detail by standardised written and additional oral instructions (see supplemental material). Then, discomfort and pain thresholds for the electric stimulation were assessed with ascending stimulation intensity. Shocks with randomly determined intensity below the pain threshold were then delivered and rated by the participants to establish a stimulus-perception relationship. Stimulus intensity was set slightly below the pain threshold to a clear discomfort level. The same stimulus intensity was used for the single fixed shock and the three shocks delivered after losing a gamble.

In a discrimination learning paradigm closely related to the bowling ball game, participants could explore the three different gambles. One bowling ball was presented at a time, and participants chose to accept or reject the gamble. Otherwise, the procedure was similar to the one described above. Participants played 5 training trials without electric shocks in order to habituate them to the task. Then, 144 trials followed in two 8 minute blocks. The last 24 trials were fully balanced and served to probe whether participants ordered their preference according to the ordering of objective outcome contingencies (figure 2). Nine (behavioural experiment) and 7 (imaging experiment) additional participants who did not have this ordering of preference were not included in the experiment proper. (Note that this procedure excludes both participants who did not learn the outcome contingencies and those who did not have a preference over one or three electric shocks.) The experiment proper was then explained with standardised written and additional oral instructions (see supplemental material), and participants had 5 training trials without electric shocks. In the imaging experiment, participants went into the scanner at this point and, while field maps were taken, engaged in another 24 trials of the learning part in order to habituate them to the scanner environment. Another 5 training trials for the experiment proper without electric shocks followed. Then, participants engaged in the experiment proper, containing 360 trials in 5 blocks of 8 minutes or 72 trials each. Breaks between blocks lasted at least 30 seconds and could be extended by participants. After the experiment, participants were asked to rate the (first-order) shock probability for each of the three ball colours (*How likely was it that shocks would be delivered after the following ball?*) on a horizontal visual analogue scale from 0% to 100% (figure 2).

Image acquisition

Images were acquired on a 3 T whole body scanner (Trio, Siemens Medical Systems, Erlangen, Germany) with a 12 channel head coil for RF transmission and signal reception. Field maps were acquired with the standard manufacturer's double echo gradient echo field map sequence (TE , 10.0 and 12.46 ms; TR , 1020 ms; matrix size, 64×64), using 64 slices covering the whole head (voxel size, $3 \times 3 \times 3$ mm).

Whole-brain T1-weighted scans were acquired using a modified driven equilibrium Fourier transform (MDEFT) sequence with optimised parameters as described previously (Deichmann et al., 2004). 176 sagittal partitions were acquired with an image matrix of 256×240 (read \times phase) and two-fold oversampling in read direction (head/foot direction) to prevent aliasing (isotropic spatial resolution 1 mm; α , 16° ; $TR/TE/TI$, 7.92 ms/2.48 ms/910 ms; band width, 195 Hz/Px). Special RF excitation pulses were used to compensate for B1 inhomogeneities of the transmit coil in superior/inferior and anterior/posterior directions.

Images were reconstructed by performing a standard 3D Fourier Transform, followed by modulus calculation. No data filtering was applied in k-space or in the image domain.

For functional images, we used BOLD signal-sensitive T_2^* -weighted transverse single-shot gradient-echo echo-planar imaging (EPI; flip angle α , 90° ; band width BW , 2298 Hz/Px; phase-encoding (PE) direction, anterior-posterior; TE , 30 ms; TR , 3264 ms). The manufacturer's standard automatic 3D-shim procedure was performed at the beginning of each experiment. Each volume contained 48 slices of 2 mm thickness (1mm gap between slices; field of view, $192 \times 192 \text{ mm}^2$; matrix size, 64×64). BOLD sensitivity losses in the orbitofrontal cortex and the amygdala due to susceptibility artefacts were minimized by applying a z-shim gradient moment of $-0.4 \text{ mT/m} \cdot \text{ms}$, a slice tilt of -30° and a positive PE gradient polarity (Weiskopf et al., 2007; Weiskopf et al., 2006). In each of five scanning sessions, 147 functional whole-brain volumes were acquired. The first 5 volumes, or 16.3 s, of each session were discarded to obtain steady-state longitudinal magnetization. Each session was concluded by 8 volumes, or 26.1 s, of rest without stimuli. During EPI imaging, we recorded respiration with a chest belt and heart rate with a pulse oxymeter on the left little finger.

Image analysis

EPI images were generated online using the scanner manufacturer's reconstruction method. Further processing was done using statistical parametric mapping (SPM8; Wellcome Trust Centre for Neuroimaging, London, UK; www.fil.ion.ucl.ac.uk/spm) on Matlab 7.4 (MathWorks, Natick, MA, USA). Images were corrected for geometric distortions caused by susceptibility-induced field inhomogeneities. A combined approach was used which corrects for both static distortions and changes in these distortions due to head motion (Andersson et al., 2001; Hutton et al., 2002). The static distortions were calculated for each subject from a field map that was processed using the FieldMap toolbox as implemented in SPM8. Using these parameters, the EPI images were then realigned and unwarped, a procedure that allows the measured static distortions to be included in the estimation of distortion changes associated with head motion. No participant moved more than 6 mm in any direction during scanning. The motion-corrected images were then co-registered to the individual's anatomical MDEFT image using a 12-parameter affine transformation, and normalised to the Montreal Neurological Institute (MNI) T1 reference brain template (re-sampled voxel size $2 \times 2 \times 2 \text{ mm}$). Normalized images were smoothed with an isotropic 8 mm full-width at half-maximum Gaussian kernel. The time series in each voxel were high-pass filtered at 1/128 Hz to remove low-frequency confounds as is standard in SPM8.

We modelled neuronal responses as a mixture of categorical (ambiguity vs. non-ambiguity) and parametric effects (linear effect of entropy within ambiguous gambles). The categorical effect involves that of entropy (entropy is always zero in non-ambiguous and non-zero in ambiguous gambles) but is confounded by a range of qualitative differences between the two conditions, whereas the parametric effect involves variation of entropy alone. Based on previous findings (Bach et al., 2009) we hypothesised that a categorical contrast of ambiguity vs. non-ambiguity would be driven by factors unrelated to entropy. Hence, different brain regions should be activated in this categorical contrast to that associated with a continuous variation of entropy. Crucially, because entropy and expected first-order probability are both a function of the second-order probabilities, expected first-order probability was taken into account for all analyses involving entropy, such that none of our results can be explained by an effect of expected first-order probability.

Specifically, we modelled trial onset and trial outcome across all blocks as stick functions convolved with a canonical haemodynamic response function (Friston et al., 1994). Trial onset was modelled separately for both trial types and for correct (i.e. making a choice

between accepting and rejecting the gamble) and missing responses (2.2% of all trials, range across participants 0% - 13%). Only correct responses were analysed. Range of first-order probabilities (only for ambiguous trials), expected first-order probability, second-order uncertainty, and choice were modelled as parametric modulators for each of the four trial onset regressors and were serially orthogonalised in this order such that the regressor pertaining to second-order uncertainty only explained variance not already explained by range of first-order probabilities and expected first-order probability. Trial outcome was modelled as a single stick function, parametrically modulated by the number of shocks received.

To account for serial acquisition of different slices in one volume and differences in the shape of the haemodynamic response function, time and dispersion derivatives for each regressor were included into the model. Regressors of no interest were movement parameters derived from the realignment procedure and regressors for variance caused by the cardiac (Glover et al., 2000) and respiratory cycle (Birn et al., 2006), and a separate constant for each block.

In additional analyses, to account for possible condition differences that might be explained by reaction time differences, these were z-transformed for each participant, multiplied with a stick function centred at trial onset, convolved with a hemodynamic response function and, together with time and dispersion derivatives, included in the model as separate regressors of no interest (i.e. not orthogonalised to any other regressor). This regressor will explain all variance in haemodynamic responses that relates to reaction times. If some variance in the other regressors can also be explained by reaction times, SPM will disregard variance that is shared between these regressors and thus discard brain responses that could be explained by reaction times.

Because trial onset and outcome were so close together in time (2 s), haemodynamic responses to these events will be correlated (for the canonical responses implemented in SPM, $R^2 = 0.56$). That means that we cannot effectively disambiguate whether an effect of trial type was caused by the presence of ambiguity, or by the resolution of ambiguity with the outcome. In order to do so, one would have had to forego presenting outcomes which was not possible within our study design (see discussion). Crucially however, within ambiguous trials, one can distinguish responses to second-order uncertainty, and response to the resolution of second-order uncertainty (i.e. second-order surprise, or self-information), because the latter depends on the actual outcome that varies on a trial-by-trial basis and is therefore only mildly correlated with entropy. Additional analyses therefore separated outcome into ambiguous/non-ambiguous trials and modelled entropy of the grey silhouette and surprise of the ball outcome (quantified as self-information, i.e. negative log of the probability of that ball outcome, given the second-order probabilities).

Statistical parametric maps were generated from linear contrasts of interest, involving only trials where a response was given (ambiguous > non-ambiguous trials, entropy, range of possible first-order probabilities in ambiguous trials, first-order probability in non-ambiguous trials) in each participant. A second level random effect analysis (RFX) was then performed using one sample *t*-tests on contrast images obtained in each participant for each comparison of interest ($df = 19$). We report clusters with a voxel-level threshold of $p < 0.001$ (uncorrected) and a cluster-level threshold of $p < .05$ (cluster-level corrected across the whole brain for family-wise error). In regions of interest for which we had prior hypotheses (posterior inferior frontal gyrus/sulcus, anterior insula, posterior parietal cortex, orbitofrontal and dorsomedial prefrontal cortex, amygdala), results were small volume corrected for family-wise error within a sphere of 15 mm diameter around peak coordinates as reported by

Bach et al. (2009), Huettel et al. (2006) and Hsu et al. (2005), and reported at a voxel-level threshold of $p < .05$. This is indicated in the results section and result tables.

Analysis of choice behaviour

Our analysis of choice behaviour had three goals: (1) to confirm that participants took the expected first-order probability conveyed by the ball position into account when making a decision (i.e. regardless of the second-order uncertainty also conveyed by the ball position). (2) To determine whether participants accounted for the categorical difference between ambiguity and non-ambiguity when making a decision and what computational algorithm they used in so doing. (3) To ascertain whether, and by which computational mechanism, participants accounted for entropy in their choice. Because the categorical difference between ambiguity and non-ambiguity also involves entropy, but encompasses additional factors that are inseparable, we needed to ensure that this categorical difference was already accounted for when analysing the effect of entropy within ambiguous trials. Therefore, in this step, we used the algorithm that best explained the categorical effect of ambiguity, and complemented this algorithm with additional attributes to account for entropy only within ambiguous trials.

The above aims were implemented in a step-wise Bayesian model comparison. To achieve goal (1), a *simple* decision algorithm was built, based on expected utility theory (von Neumann and Morgenstern, 1944) that uses the average of both conditional first-order probabilities as expected first-order probability (i.e. does not take the ball position into account), weights all possible outcomes (0, 1, or 3 electric shocks) by a standard power utility function with one free parameter, and computes a choice probability using a softmax rule with another free parameter. This was then compared against a more plausible *baseline* algorithm (ideal Bayesian or rational observer) where expected first-order probability is a function of second-order and conditional first-order probabilities (i.e. does take the ball position into account).

In addressing goal (2), note that the *baseline* model was built on the assumption that compound (i.e., hierarchical) lotteries can be reduced to single-stage bets (von Neumann and Morgenstern, 1944) by calculating an *expected first-order probability*, which would be incompatible with aversion to ambiguity or to second-order uncertainty (Halevy, 2007). We therefore compared the baseline model to a number of models from the economic literature that drop the compound lotteries axiom, all of which are listed in table 1, described in mathematical detail in the supplemental material and reviewed comprehensively in Camerer (1999).

As detailed, the model that best explained the categorical difference between ambiguous and non-ambiguous gambles, namely the SOP model (see table 1 for a description), was then complemented to accommodate for an effect of entropy within ambiguous trials (goal 3). The models took an analogous form as for step 2, with the difference that they multiplied the respective free parameter by mean-centred entropy for each ambiguous gamble (with the exception of the minimax model that cannot usefully be modified by entropy and was therefore dropped, see table 1).

Hence, model comparison was conducted in three steps, comparing (1) the simple with the baseline model, (2) the baseline with ambiguity models, and (3) the best ambiguity model with entropy models. This step-wise model comparison is feasible as all steps refer to independent variance components, and is more robust than simultaneous comparison of all possible model combinations. To make allowance for the expected explanatory power of more complex models in the absence of any true difference, model complexity was penalised using Akaike information criterion (AIC) and the Bayesian information criterion

(BIC) which led to similar results unless stated otherwise. Note these approaches are simply two different approximations to the true Bayesian model evidence; though not originally motivated from a Bayesian perspective, model comparisons based on AIC are asymptotically equivalent to those based on Bayes factors (Akaike, 1973; Penny et al., 2004), see also Burnham & Anderson (2004) for a Bayesian perspective on AIC. We assumed that different individuals might implement different models and performed a random-effects analysis where the structure of the model is assumed to be a random variable across participants as described previously (Stephan et al., 2009). This analysis computes the proportion of individuals that implement each model and an *exceedance probability* that states how likely it is that each model is more frequent than any other, in the population from which the sample of subjects were drawn. In the absence of conventions on exceedance probabilities, we defined a probability of $p > .95$ as decisive, to be comparable to standard model comparison. Also, we used family wise comparison of models as pointed out in the results section (Penny et al., 2010). All models used the objective conditional first-order probability. Additional analyses taking account of the subjective probability as stated in post-experimental ratings (see procedure) led to similar results. To enhance statistical power, data from both experiments was combined for choice analysis. Analysis of the individual experiments lead to very similar results.

Results

Reaction times

To control for possible effects of ambiguity on decision difficulty, we analysed reaction times (RT). Responses in ambiguous trials were considerably slower than in non-ambiguous trials (behavioural experiment: 987 ± 44 ms vs. 850 ± 50 ms; $p < .001$; imaging experiment: 1095 ± 47 ms vs. 958 ± 39 ms; $p < .001$). There was no significant correlation between RT and entropy.

Choice behaviour associated with ambiguity and entropy

Firstly, the baseline (ideal observer) model performed better than the simple model (exceedance probability $p > .999$), showing that participants took the ball position into account to compute expected first-order probability, as instructed. Secondly, we found that models which take account of a categorical difference between ambiguous and non-ambiguous gambles were better at explaining the data than the baseline model (exceedance probability $p > .999$). Figure 3 (left panel) shows the exceedance probabilities (i.e. probabilities that this model is more frequent in the population than any other) for individual models. The model that most frequently accounted for the difference between ambiguous and non-ambiguous lotteries was the SOP model (see table 1, and see supplementary material for all formulae). In this model, conditional expected outcomes are non-linearly weighted by exponentiating them with a constant – for both ambiguous and non-ambiguous trials – and then combined using the unbiased second order probabilities (Segal, 1987; Klibanoff et al., 2005). Hence, for a typical ambiguity-averse individual in an ambiguous trial, the combination of two weighted conditional expected outcomes will be less valuable than weighting the average expected outcome under non-ambiguous conditions and so explain a difference between ambiguous and non-ambiguous lotteries. In economic theory, this is analogous to a non-linear utility function explaining individual risk preferences. Thirdly, models which account for entropy within ambiguous gambles explained the data better than a model that did not when analysing AIC ($p > .95$), while BIC with its heavier penalisation of additional parameters preferred the simpler model. As shown in figure 3 (right panel), the “pessimistic” weighting model performed best where the amount and direction (optimistic vs. pessimistic) of a second-order probability weighting was dependent on the actual amount of entropy (i.e., pessimistic weighting with above-average entropy, and

vice versa). This is achieved, in the model, by multiplying the second-order probability of the worst outcome scenario with a variable that is greater than one if entropy is higher than average, and smaller than one if entropy is below average. There was even stronger evidence (i.e. for both AIC and BIC) that models which did not use SOP to account for the effect of entropy were more frequent than the model using SOP ($p > .99$), implying different mechanisms accounting for the effects of ambiguity and entropy.

Taken together, this suggests that both the categorical presence of ambiguity (as opposed to non-ambiguity) and the amount of entropy within these ambiguous gambles are taken into account when making a decision. However, there is strong evidence that the mechanism by which the decision is influenced is different for the mere presence of ambiguity and for variations in the associated degree of entropy. The algorithm that best explained choice differences between ambiguous and non-ambiguous lotteries leaves the second-order probabilities unbiased, whereas the algorithm accounting for entropy biases appears to bias second-order probabilities.

Brain regions responsive to entropy and ambiguity

Our modelling data pointed to distinct processes associated with ambiguity and its entropy. Thus, we used fMRI to test the validity of our model in so far as it predicts dissociable brain responses to categorical ambiguity on the one hand and its entropy on the other. Data are summarised in table 2. The main contrast of interest was the effect of entropy within ambiguous trials. Here, we found an association between entropy and BOLD responses in bilateral posterior midline areas, including the posterior cingulate and cuneus, extending laterally into adjacent parts of the parietal, occipital, and temporal cortices, and also into areas outside of the gray matter, possibly due to between-subjects anatomical variability and smoothing (figure 4A-D; estimated BOLD responses in figure S1, supplemental data). These BOLD responses were greater when entropy was lower with no responses observed for the opposite contrast. In order to exclude that these responses were caused by the surprise of outcomes (which is, on the average, related to entropy), we modelled outcomes separately for the two trial types, and added a regressor accounting for the instantaneous second-order surprise of the ball outcome in ambiguous trials. This revealed the same association of entropy with BOLD responses in a slightly smaller cluster (4925 voxels). Positive responses to surprise in the posterior parietal cortex failed to reach whole-brain level corrected significance ($p = .06$), whereas negative responses to surprise were found in the medial right cuneus (see table S1 in supplemental material). To exclude the possibility that choice might explain brain activations relating to second-order uncertainty, we switched the serial position of the regressors for second-order uncertainty and choice. This did not impact on our reported results and brain responses were similar in location and extent (4996 voxels).

In contrast, a categorical difference between ambiguous and non-ambiguous trials (table 2; figure 4E; estimated BOLD responses in figure S1, supplemental data) showed enhanced BOLD responses in left posterior parietal cortex (pPAR), overlapping with clusters found previously (Bach et al., 2009; Huettel et al., 2006). To exclude reaction time differences between conditions as explaining a difference in BOLD response, we co-varied these out in an additional analysis. This yielded a smaller (335 voxel) cluster in the same location. No brain region showed a greater response in non-ambiguous compared to ambiguous trials, in keeping with a previous observation (Bach et al., 2009). Notably, responses in the pPAR overlapped with those seen for the effect of entropy (compare figure 4D and E). However, the direction of the effect in this region was in diametrical opposition: the BOLD response in this region was negatively correlated with entropy, but positively correlated with ambiguity. Given the opposite sign of the association, one might conclude that responses to ambiguity are not driven by the difference in entropy between these two conditions, but other factors.

For some individuals, the presence of ambiguity is likely to be more relevant than for others. We explored inter-individual differences in the brain response to ambiguity (figure 4F) and showed that higher rejection rates for ambiguous gambles (hence: *ambiguity aversion*) were associated with enhanced responses to ambiguous > non-ambiguous gambles in the right inferior frontal gyrus. This cluster was not found in the main contrast ambiguous > non-ambiguous gambles but was in the vicinity of a cluster reported previously and survived small volume correction within a sphere around this location (Huettel et al., 2006). There was no association of ambiguity aversion with brain responses to entropy, again underlining the distinction between categorical ambiguity and entropy in ambiguity.

Our design incidentally also varied the range of possible outcome probabilities, and the expected risk probability. The difference between ambiguous trials with a high range of possible outcome probabilities (i.e. between 0.2 and 0.8) and those with a low range (i.e. 0.2 to 0.5 or 0.5 to 0.8), represented in the contrast low > high outcome probability range, showed a localised BOLD response in the left prefrontal cortex (see figure S2C, supplemental data). This cluster survived small volume correction around peak coordinates reported previously (Huettel et al., 2006). Note that our paradigm requires participants to form neural representations of three previously learned shock probabilities; and behavioural analyses indicated this was the case. In keeping with this, brain responses associated with higher risk probability were observed in the right posterior parietal cortex and bilateral dorsolateral prefrontal cortex (see figure S2A and B).

Discussion

Uncertainty about action-outcome associations pervades most real-life decisions and can be conceptualised in a hierarchical probabilistic model where uncertainty of future outcome predictions is termed *first-order uncertainty*, and uncertainty about the associations themselves *second-order uncertainty*. To investigate the neural representation of second-order uncertainty, previous studies have relied on the economic concept of ambiguity where second-order uncertainty can be either zero (no ambiguity) or maximal (ambiguity); however this categorical contrast is confounded with qualitative factors. Hence, in the present study we continuously varied the second-order uncertainty of ambiguous lotteries continuously as the entropy of second-order probabilities and contrasted this with non-ambiguous lotteries that did not contain second-order uncertainty.

Three main findings emerge. First, decisions on the lotteries appeared to depend on the amount of second-order entropy such that participants were biased towards avoiding a lottery with high second-order entropy, in keeping with a previous study that qualitatively varied second-order uncertainty in ambiguous gambles but fell short of a precise definition (Keren and Gerritsen, 1999). Bayesian comparison suggested that this effect on choice is achieved by biasing the second-order probabilities towards the worst possible outcome scenario when second-order entropy is high (and vice versa when second-order uncertainty is low). Secondly, BOLD responses to second-order entropy were found in a large posterior midline area encompassing posterior cingulate cortex (PCC) and precuneus, with stronger responses when entropy was low. Thirdly, we replicate previous findings that the mere presence of ambiguity is avoided when compared to non-ambiguous lotteries (Slovic and Tversky, 1974; Becker and Brownson, 1964; Yates and Zukowski, 1976; Curley et al., 1986; MacCrimmon and Larson, 1979; Keren and Gerritsen, 1999). Model comparison indicated that this was achieved by non-linearly weighting the conditional expected outcomes before combining them into an expected outcome, without biasing second order probabilities (Segal, 1987; Klibanoff et al., 2005). BOLD responses to this contrast were seen in posterior parietal cortex (pPAR), replicating two previous studies (Huettel et al., 2006; Bach et al., 2009). Additionally, within posterior prefrontal cortex BOLD responses in this same

contrast varied with the amount of ambiguity aversion on a between-subjects level, an area previously implicated in processing ambiguity (Huettel et al., 2006; Bach et al., 2009).

In interpreting these results, we first note that second-order entropy and the contrast ambiguity vs. non-ambiguity engage different brain areas and appear to evoke different computational processes. The latter contrast has previously been taken to reflect differences in second-order uncertainty (because non-ambiguous lotteries have zero second-order uncertainty), but this disregards confounds such as having to publicly reveal uncertainty attitudes (Curley et al., 1986; Trautmann et al., 2008), the possibility of very unfavourable first-order probabilities (Keren and Gershten, 1999; Larson, 1980), and search for potentially knowable information (Chow and Sarin, 2002; Bach et al., 2009). Also, the very logic of ambiguous lotteries seems more complicated than of non-ambiguous ones and might therefore be aversive. Our findings suggest that the neural representation of ambiguity may not be driven by second-order uncertainty but may reflect these confounding factors, a conclusion that qualifies previous findings (Huettel et al., 2006; Bach et al., 2009). At the same time, ambiguity aversion might not be driven by second-order uncertainty either, thus calling into question the conceptualisation of ambiguity as “missing information” (Camerer, 1995).

By manipulating second-order uncertainty in a continuous and unconfounded manner as second-order entropy, we demonstrate an encoding in a posterior midline network with greater responses when uncertainty is lower. This effect survives controlling for surprise at the time point of the outcome and is therefore driven by the presence of uncertainty rather than its resolution. Several explanations might account for this. First, our behavioural data indicate that lotteries with high second-order uncertainty are devaluated. Previous fMRI studies have found PCC responses representing reward size or subjective value in different paradigms (Smith et al., 2009; Kable and Glimcher, 2007; Ballard and Knutson, 2009; Peters and Buchel, 2009; Kable and Glimcher, 2010). Hence, the stronger signal with lower entropy could be interpreted as reflecting higher subjective value of less uncertain lotteries.

However, other areas commonly associated with subjective value (namely striatum and prefrontal/orbitofrontal regions) do not track second-order entropy in the present task. An alternative possibility therefore is that posterior midline areas implicated in the “default mode” network (Raichle et al., 2001; Buckner et al., 2008) are actively suppressed when processing demands (i.e. uncertainty) are high, an observation in keeping with findings that PCC is involved in retrieval of self-relevant memory (Summerfield et al., 2009), and that its activity predicts task errors (Li et al., 2007). Similar reasoning has been deployed to explain findings of smaller prefrontal deactivations in easy as compared to difficult perceptual decisions (Tosoni et al., 2008; Ho et al., 2009).

However, BOLD responses in PCC can also be conceived as directly related to second-order entropy. Greater responses with lower uncertainty would then point to the notion that precision (or amount of information) is encoded. How this encoding is achieved at a neuronal level remains to be determined although we note a proposal that allows for a representation in a large neural population (see for a review Friston, 2009; 2010). For future investigations it might be useful to reframe the decision making model that best explained the data in a (more general) Bayesian framework with “pessimistic” priors that enjoy more weight when the likelihood function has higher uncertainty. Thus, it would be possible to study the effect of pessimistic priors separately from the effect of the likelihood function in order to disambiguate effects of subjective value and second-order uncertainty.

With respect to the categorical effect of ambiguity we replicated activation in posterior parietal cortex (pPAR) (Bach et al., 2009; Huettel et al., 2006) which we have previously

interpreted as reflecting more demanding expected value calculations. Ambiguity might also engage behavioural mechanisms to find out an unknown bit of information, an explanation we previously used to explain responses in anterior insula/posterior prefrontal cortex. No such responses were found in the present study. We suggest that BOLD responses related to outcome search would possibly not occur with the very short anticipation times implemented in the present study (2 s) as opposed to longer times (4.5 to 6 s) used in previous studies. Note, however, that an effect was found in this area for participants with high ambiguity aversion such that differences between this and the previous study may also relate to sample differences in ambiguity aversion. We did not find any responses in brain areas identified in Hsu et al. (2005), in particular amygdala and orbitofrontal cortex (OFC). Although concerned with ambiguity, the latter study collapsed different kinds of situations involving a lack of information during analysis of fMRI data. Thus, in our view there is no convincing empirical evidence that the amygdala or OFC respond to ambiguity.

Model comparison demonstrated that a second order probability (SOP) model was most frequent in explaining the effect of ambiguity on choice, where conditional expected outcomes are non-linearly weighted before being combined with the unbiased SOP into an expected outcome (Segal, 1987; Klibanoff et al., 2005). A plethora of models accounting for ambiguity has been proposed in the economic literature, mainly based on theoretical/axiomatic considerations (see for a review Camerer, 1999); to the best of our knowledge we present the first Bayesian model comparison based on choice data.

Several potential limitations of the present study remain to be resolved. Firstly, although there is strong evidence that ambiguity and second-order entropy are accounted for by different computational mechanisms, we acknowledge that the evidence is weaker in determining the precise model by which second-order entropy is taken into account. Thus, it would be desirable to replicate our findings in an experimental setup specifically designed to distinguish between the most likely models from our study. Also, we did not test alternative formulations of second-order uncertainty (e.g. variance). A limitation concerning BOLD responses to ambiguity is that onset and offset of the lottery were close together in time such that it is not possible to disambiguate responses to the presence of ambiguity and its resolution. Note that this does not discount the BOLD effect of second-order uncertainty where responses to uncertainty itself, and to the surprise of the outcome, could be disambiguated.

Tracking second-order uncertainty is crucial for any hierarchical model of perception or learning, such as hierarchical Bayesian (and related predictive coding schemes) (e.g. Friston, 2010) and volatility models used to account for context-sensitive learning (Behrens et al., 2007). Indeed, recent findings suggest that the changes in precision during perceptual discrimination are a key determinant of neuronal activity (Hesselmann et al., 2010). We extend such ideas to explicit decisions by showing that second-order uncertainty in ambiguous lotteries influences choice, possibly by a mechanism that biases second-order probabilities. This second order uncertainty is negatively correlated with BOLD activity in a posterior midline network encompassing the PCC. However, ambiguity as defined within an economic perspective activates different brain regions and appears to exert its effect on choice via a different mechanism. Hence, we argue that behavioural and neural responses to ambiguity might not be driven by second-order uncertainty, calling for a reinterpretation of previous data, and for further research into the mechanisms by which second-order uncertainty acts to guide learning and decision-making.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors wish to thank Karl Friston for his comments on this manuscript. Rosalyn Moran and Mkael Symmonds contributed to the design of the experiments; and Peter Dayan provided constructive discussion of the results. This work was funded by a programme grant from the Wellcome Trust and a Max Planck Research Award to RJD, and by a personal grant from the Swiss National Science Foundation to DRB.

References

- Akaike H. Information measures and model selection. *Bulletin of the International Statistical Institute*. 1973; 20:277–290.
- Andersson JL, Hutton C, Ashburner J, Turner R, Friston K. Modeling geometric deformations in EPI time series. *Neuroimage*. 2001; 13:903–919. [PubMed: 11304086]
- Bach DR, Seymour B, Dolan RJ. Neural activity associated with the passive prediction of ambiguity and risk for aversive events. *J Neurosci*. 2009; 29:1648–1656. [PubMed: 19211872]
- Ballard K, Knutson B. Dissociable neural representations of future reward magnitude and delay during temporal discounting. *Neuroimage*. 2009; 45:143–150. [PubMed: 19071223]
- Becker SW, Brownson FO. What price ambiguity? Or the role of ambiguity decision making. *J Polit Econ*. 1964; 72:62–73.
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat Neurosci*. 2007; 10:1214–1221. [PubMed: 17676057]
- Birn RM, Diamond JB, Smith MA, Bandettini PA. Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *Neuroimage*. 2006; 31:1536–1548. [PubMed: 16632379]
- Buckner RL, Andrews-Hanna JR, Schacter DL. The brain's default network: anatomy, function, and relevance to disease. *Ann N Y Acad Sci*. 2008; 1124:1–38. [PubMed: 18400922]
- Burnham K, Anderson D. Multimodel Inference: Understanding AIC and BIC in model selection. *Sociological Methods Research*. 2004; 33:260–304.
- Camerer, C. Individual decision making. In: Kagel, JH.; Roth, AE., editors. *Handbook of experimental economics*. Princeton University Press; Princeton, NJ: 1995. p. 587-704.
- Camerer, C. Ambiguity aversion and non-additive probability: Experimental evidence, models and applications. In: Luini, L., editor. *Uncertain decisions: Bridging theory and experiments*. Kluwer; Boston: 1999. p. 53-80.
- Chow CC, Sarin RK. Known, unknown, and unknowable certainties. *Theory and Decision*. 2002; 52:127–138.
- Curley SP, Yates F, Abrams RA. Psychological sources of ambiguity avoidance. *Organ Behav Hum Decis Process*. 1986; 38:230–256.
- Daw ND, Niv Y, Dayan P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci*. 2005; 8:1704–1711. [PubMed: 16286932]
- Deichmann R, Schwarzbauer C, Turner R. Optimisation of the 3D MDEFT sequence for anatomical brain imaging: technical implications at 1.5 and 3 T. *Neuroimage*. 2004; 21:757–767. [PubMed: 14980579]
- Dreher JC, Kohn P, Berman KF. Neural coding of distinct statistical properties of reward information in humans. *Cereb Cortex*. 2006; 16:561–573. [PubMed: 16033924]
- Ellsberg D. Risk, ambiguity, and the Savage axioms. *Q J Econ*. 1961; 75:643–669.
- Friston K. The free-energy principle: a rough guide to the brain? *Trends Cogn Sci*. 2009; 13:293–301. [PubMed: 19559644]
- Friston K. The free-energy principle: a unified brain theory? *Nat Rev Neurosci*. 2010; 11:127–138. [PubMed: 20068583]
- Friston KJ, Jezzard P, Turner R. Analysis of functional MRI time-series. *Hum Brain Mapp*. 1994; 1:153–171.
- Friston KJ, Price CJ, Fletcher P, Moore C, Frackowiak RS, Dolan RJ. The trouble with cognitive subtraction. *Neuroimage*. 1996; 4:97–104. [PubMed: 9345501]

- Ghirardato P, Maccheroni F, Marinacci M. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*. 2004; 118:133–173.
- Glover GH, Li TQ, Ress D. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magn Reson Med*. 2000; 44:162–167. [PubMed: 10893535]
- Halevy Y. Ellsberg revisited: An experimental study. *Econometrica*. 2007; 75:503–536.
- Hesselmann G, Sadaghiani S, Friston KJ, Kleinschmidt A. Predictive coding or evidence accumulation? False inference and neuronal fluctuations. *PLoS ONE*. 2010; 5:e9926. [PubMed: 20369004]
- Ho TC, Brown S, Serences JT. Domain general mechanisms of perceptual decision making in human cortex. *J Neurosci*. 2009; 29:8675–8687. [PubMed: 19587274]
- Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF. Neural systems responding to degrees of uncertainty in human decision-making. *Science*. 2005; 310:1680–1683. [PubMed: 16339445]
- Huettel SA, Stowe CJ, Gordon EM, Warner BT, Platt ML. Neural signatures of economic preferences for risk and ambiguity. *Neuron*. 2006; 49:765–775. [PubMed: 16504951]
- Hutton C, Bork A, Josephs O, Deichmann R, Ashburner J, Turner R. Image distortion correction in fMRI: A quantitative evaluation. *Neuroimage*. 2002; 16:217–240. [PubMed: 11969330]
- Kable JW, Glimcher PW. The neural correlates of subjective value during intertemporal choice. *Nat Neurosci*. 2007; 10:1625–1633. [PubMed: 17982449]
- Kable JW, Glimcher PW. An “As Soon As Possible” Effect in Human Inter-Temporal Decision Making: Behavioral Evidence and Neural Mechanisms. *J Neurophysiol*. 2010
- Keren G, Gershten LEM. On the robustness and possible accounts of ambiguity aversion. *Acta Psychol (Amst)*. 1999; 103:149–172.
- Klibanoff P, Marinacci M, Mukerji S. A smooth model of decision making under ambiguity. *Econometrica*. 2005; 73:1849–1892.
- Knight, FH. *Risk, Uncertainty and Profit*. Houghton Mifflin; Boston: 1921.
- Larson JR. Exploring the external validity of a subjectively weighted utility model of decision making. *Organ Behav Hum Perform*. 1980; 26:293–304.
- Li CS, Yan P, Bergquist KL, Sinha R. Greater activation of the “default” brain regions predicts stop signal errors. *Neuroimage*. 2007; 38:640–648. [PubMed: 17884586]
- MacCrimmon, KR.; Larson, S. Utility theory: Axioms versus ‘Paradoxes’. In: Allais, M.; Hagen, O., editors. *Expected utility hypotheses and the Allais Paradox*. D. Reidel; Dordrecht: 1979. p. 333-410.
- O’Neill M, Kobayashi S. Risky business: disambiguating ambiguity-related responses in the brain. *J Neurophysiol*. 2009; 102:645–647. [PubMed: 19515950]
- Oldfield RC. The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*. 1971; 9:97–113. [PubMed: 5146491]
- Pearce JM, Hall G. A Model for Pavlovian Learning - Variations in the Effectiveness of Conditioned But Not of Unconditioned Stimuli. *Psychol Rev*. 1980; 87:532–552. [PubMed: 7443916]
- Penny WD, Stephan KE, Daunizeau J, Rosa MJ, Friston KJ, Schofield TM, Leff AP. Comparing families of dynamic causal models. *PLoS Comput Biol*. 2010; 6:e1000709. [PubMed: 20300649]
- Penny WD, Stephan KE, Mechelli A, Friston KJ. Comparing dynamic causal models. *Neuroimage*. 2004; 22:1157–1172. [PubMed: 15219588]
- Peters J, Buchel C. Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. *J Neurosci*. 2009; 29:15727–15734. [PubMed: 20016088]
- Preuschoff K, Bossaerts P, Quartz SR. Neural differentiation of expected reward and risk in human subcortical structures. *Neuron*. 2006; 51:381–390. [PubMed: 16880132]
- Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL. A default mode of brain function. *Proc Natl Acad Sci U S A*. 2001; 98:676–682. [PubMed: 11209064]
- Segal U. The Ellsberg Paradox and Risk-Aversion - An Anticipated Utility Approach. *International Economic Review*. 1987; 28:175–202.
- Shannon CE. A mathematical theory of communication. *The Bell System Technical Journal*. 1948; 27:379–423. 623–656.

- Slovic P, Tversky A. Who accepts Savage's axiom? *Behav Sci.* 1974; 19:368–373.
- Smith BW, Mitchell DG, Hardin MG, Jazbec S, Fridberg D, Blair RJ, Ernst M. Neural substrates of reward magnitude, probability, and risk during a wheel of fortune decision-making task. *Neuroimage.* 2009; 44:600–609. [PubMed: 18804540]
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ. Bayesian model selection for group studies. *Neuroimage.* 2009
- Summerfield JJ, Hassabis D, Maguire EA. Cortical midline involvement in autobiographical memory. *Neuroimage.* 2009; 44:1188–1200. [PubMed: 18973817]
- Tobler PN, O'Doherty JP, Dolan RJ, Schultz W. Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *J Neurophysiol.* 2007; 97:1621–1632. [PubMed: 17122317]
- Tosoni A, Galati G, Romani GL, Corbetta M. Sensory-motor mechanisms in human parietal cortex underlie arbitrary visual decisions. *Nat Neurosci.* 2008; 11:1446–1453. [PubMed: 18997791]
- Trautmann ST, Vieider FM, Wakker PP. Causes of ambiguity aversion: Known versus unknown preferences. *Journal of Risk and Uncertainty.* 2008; 36:225–243.
- von Neumann, J.; Morgenstern, O. *Theory of games and economic behavior.* Princeton University Press; Princeton, NJ: 1944.
- Weiskopf N, Hutton C, Josephs O, Deichmann R. Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *Neuroimage.* 2006; 33:493–504. [PubMed: 16959495]
- Weiskopf N, Hutton C, Josephs O, Turner R, Deichmann R. Optimized EPI for fMRI studies of the orbitofrontal cortex: compensation of susceptibility-induced gradients in the readout direction. *MAGMA.* 2007; 20:39–49. [PubMed: 17268781]
- Yates JF, Zukowski LG. Characterization of Ambiguity in Decision-Making. *Behav Sci.* 1976; 21:19–25.

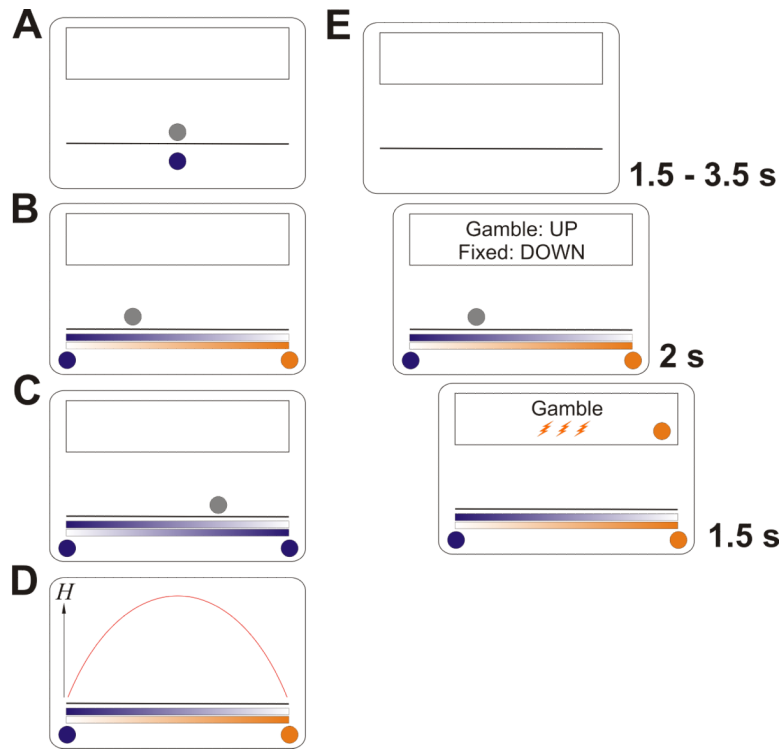


Figure 1.

Experimental setup. (A) Preceding learning task. In each of 144 trials, participants saw a coloured ball. Each of three ball colours represented a gamble (framed as “lottery ticket”) with a different first-order probability (0.2, 0.5, 0.8) of getting three electric shocks, or nothing. Participants chose between that gamble and one certain electric shock of the same magnitude. A fading grey silhouette directly above it indicated the time that was allowed for the decision. (B) Experiment proper. On each of 276 ambiguous trials, two coloured balls appeared, indicating two possible first-order probabilities. It was explained that these balls represented two bowling ball players, one of which would play his ball that would be shown as a silhouette. Thus, the closer the silhouette appeared to one of the two balls, the more likely it was to represent this ball. Again, a decision had to be made between the gamble and one certain shock. (C) In 84 additional non-ambiguous trials, the two balls had the same colour. The position of the grey silhouette is uninformative in this case. (D) Second-order uncertainty, quantified as Shannon entropy H of the second-order probabilities of which conditional first-order probability would be realised, given the silhouette position. It can be seen that this quantity is highest in the middle between the two balls. (E) Timeline. After a variable inter-trial interval, the gamble was presented, together with an indication of which key to press (this was held constant within participants and varied between). While the silhouette faded out, a decision had to be made within 1.5 s (behavioural experiment), or 2 s (imaging experiment). The choice, the ball that was being played, and the ultimate outcome, were then shown in the feedback box above the bowling ball lane.

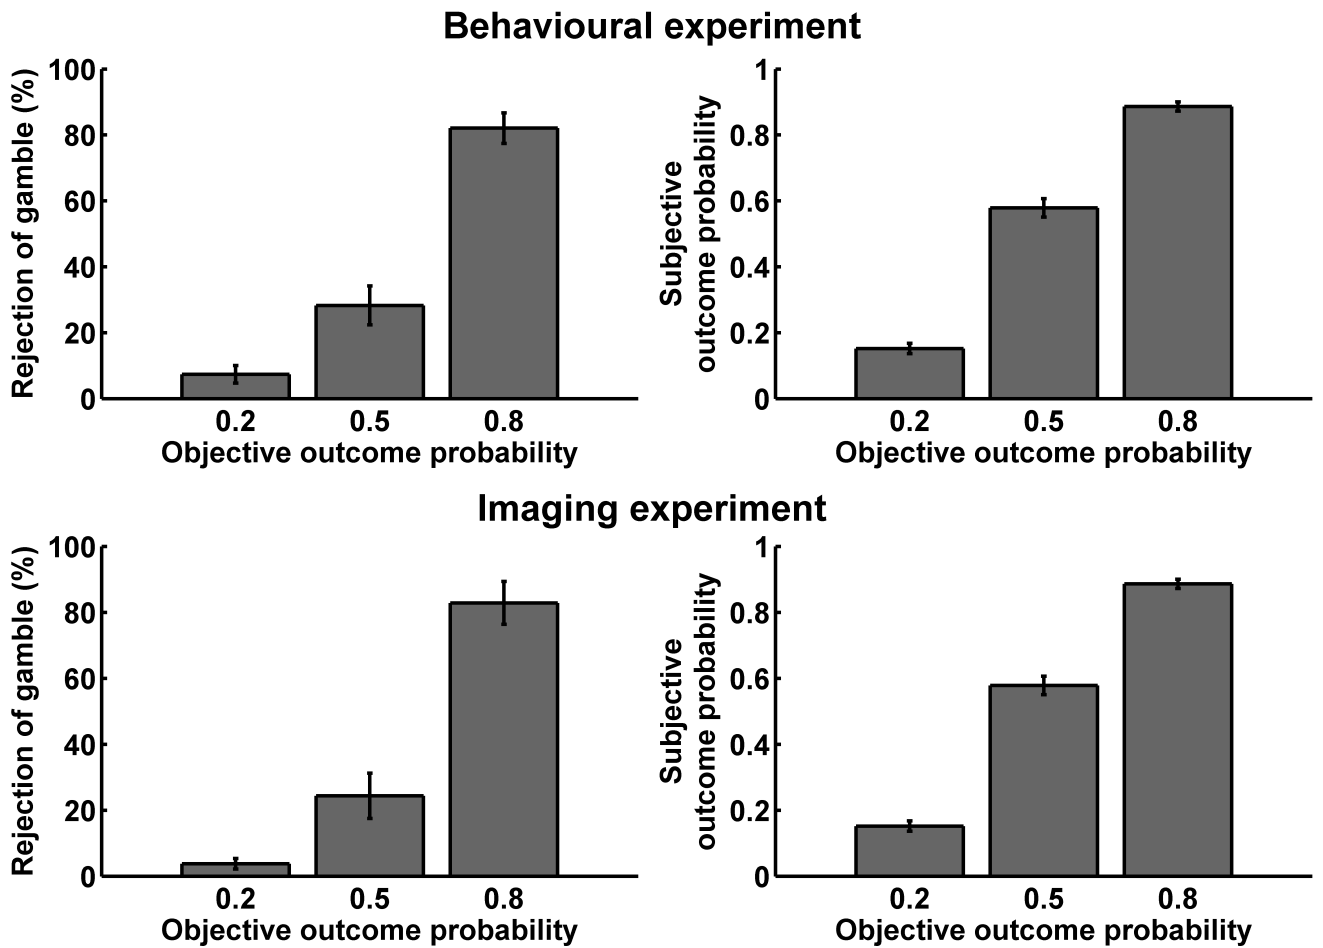


Figure 2.

In a preceding learning task consisting of 144 trials, participants learned the outcome probabilities of the three coloured balls. Choices for the last 24 balanced trials of this part are shown on the left. After the experiment proper, participants were asked for explicit estimates of the outcome probabilities of the three coloured balls, as shown on the right. All results are mean \pm standard error for participants that were included into the analysis (behavioural experiment: $N = 24$, imaging experiment, $N = 20$). Participants whose choice proportions and subjective probability estimates did not monotonically vary with outcome probability were excluded from scanning/analysis.

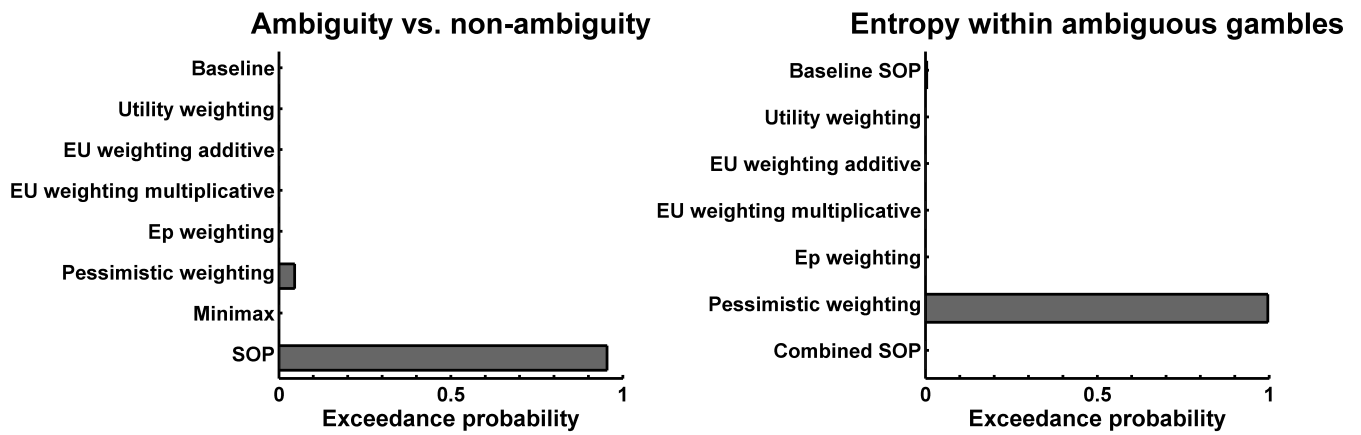


Figure 3.

Exceedance probabilities (i.e. probability that a given model is more frequent in the population than the other models) of a random-effects Bayesian model comparison. Results are shown for the Akaike information criterion (AIC); see text for Bayesian information criterion (BIC). Exceedance probabilities of the winning models are decisive ($p > .95$). The models are detailed in table 1 and in the supplemental material.




Figure 4.

A-D: BOLD responses to low second-order uncertainty (i.e. Shannon entropy of second-order probabilities) within ambiguous trials. E: Different responses were found for the contrast ambiguity vs. non-ambiguity (note that responses in the posterior parietal cortex, panels D/E, are high for ambiguity where entropy is higher than in non-ambiguity, but also high for low entropy within ambiguous trials, suggesting that responses to ambiguity are not driven by entropy). F: Responses to ambiguity vs. non-ambiguity in right inferior frontal gyrus covaries with ambiguity aversion on a between-subject level.

All contrasts are overlaid on an average normalised T1 image from the whole sample and are in neurological convention. Responses in panel A-E survive whole brain correction for family-wise error at a cluster-level threshold of $p < .05$ and a voxel-level threshold of $p < .001$; responses in panel F are small volume corrected for family-wise error at a voxel-level threshold of $p < .05$.

Table 1

Summary of models analysed to account for the influence of ambiguity and second-order uncertainty on choice (see also figure 3).

Models accounting for ambiguity		Models accounting for ambiguity and second-order uncertainty*	
Baseline	Baseline (ideal observer) model, assuming the reduction of compound lotteries axiom and taking into account the silhouette position for the calculation of expected outcome probability	Baseline SOP	Baseline second order probability (SOP) weighting model without additional parameters to account for entropy
Utility weighting	Utility weighting model, using different utility functions for ambiguous and non-ambiguous decisions	Utility weighting	Utility weighting model, assuming modulation of the utility function in ambiguous trials
EU weighting additive	Expected utility additive weighting model, adding a constant to the expected utility of ambiguous choices	EU weighting additive	Expected utility additive weighting model, adding a variable amount to the expected utility of ambiguous choices
EU weighting multiplicative	Expected utility multiplicative weighting model, multiplying the expected utility of ambiguous choices with a constant.	EU weighting multiplicative	Expected utility multiplicative weighting model, multiplying the expected utility of ambiguous choices with a variable
Ep weighting	Expected probability weighting model, non-linearly weighting the expected probabilities in ambiguous choices by exponentiating them with a constant. (Hsu et al., 2005)	Ep weighting	Expected probability weighting model, exponentiating expected outcome probability of ambiguous choices with a variable
Pessimistic weighting	Pessimistic weighting model, biasing the second order probabilities towards the worse scenario in ambiguous choices (Ghirardato et al., 2004; Huettel et al., 2006)	Pessimistic weighting	“Pessimistic” weighting model, overestimating/underestimating the second order probability of the worse scenario
Minimax	Minimax model, where choice in ambiguous trials is based on the worse scenario only		not applicable
SOP	Second order probability (SOP) model, where the conditional expected outcomes are non-linearly weighted before being combined with the unbiased second-order probabilities into an expected outcome (Segal, 1987; Klibanoff et al., 2005)	Combined SOP	Combined second order probability (SOP) model, where the second order probabilities are additionally weighted by entropy.

* Based on the second order probability model, additional models were formed that accounted for entropy within ambiguous choices by having an additional free parameter that is modulated by mean-centred entropy.

Table 2

Blood oxygen level dependent (BOLD) responses to contrasts of interest. All clusters are reported cluster-level corrected with $p < .05$ for family-wise error (FWE). Additional clusters from regions of interest are reported at a voxel-level significance threshold of $p < .001$, small volume corrected with $p < .05$ for family-wise error (FWE) in regions of interest for a sphere with 15 mm radius around peak coordinates as reported by Bach et al. (2009) and Huetzel et al. (2006) (marked with *). No significant BOLD responses were observed for any other contrast of interest.

Brain regions	Brodman area of local maxima	Hemi-sphere	Voxel number	Voxel t score	Montreal Neurological Institute brain template coordinates of local maxima
<i>Responses to low second-order uncertainty (i.e. entropy) within ambiguous gambles</i>					
Posterior cingulate, superior parietal lobule, postcentral gyrus, parahippocampal gyrus, precuneus, middle temporal gyrus, angular gyrus, superior & middle occipital gyrus,	30 & 31	Bilateral	5383	7.67	16, -58, 6; -20, -62, 8; 18, -66, 24
Superior parietal lobule, precuneus, angular gyrus	7	Left	193	5.05	-18, -54, 54
<i>Ambiguous > non-ambiguous gambles</i>					
Precuneus, superior & inferior parietal lobule, postcentral gyrus	7	Left	516	5.6	-16, -54, 46; -22, -54, 52; -10, -62, 56
<i>Ambiguous > non-ambiguous gambles: correlation with individual measures of ambiguity aversion</i>					
<i>Low > high probability range within ambiguous gambles</i>					
Inferior frontal gyrus*	47	Right	61	5.21	44, 32, 2
Middle frontal gyrus*	45 & 46	Left	63	4.58	-40, 36, 16; -34, 28, 12
<i>First-order probability (non-ambiguous gambles) †</i>					
Inferior parietal lobule, supramarginal gyrus, angular gyrus	40 & 39	Right	404	6.77	44, -44, 32; 38, -50, 26; 48, -64, 34
Superior & middle frontal gyrus	9 & 8	Left	154	5.78	-22, 30, 36; -28, 26, 40; -28, 18, 48
Middle frontal gyrus	9	Right	126	5.22	38, 22, 42

† Note that in non-ambiguous trials, normalised shock probability equals normalised expected value, or utility, as value/utility was not varied throughout the experiment.