Supplemental Material

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

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

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

Supplemental Methods I: Instructions

a. Instructions for the behavioural experiment

Part 1: printout instructions

Hi

Thanks for taking part in this experiment. It's an easy, two-level computer game.

In each round, you will make a decision:

EITHER take a GAMBLE.

This option is called a gamble because you don't know what you will get. You might be lucky or unlucky. If you are unlucky, you get three electric shocks. If you are lucky, you get nothing. It's like a lottery.

OR take one FIXED SHOCK. You will always get it.

So, which decision should you take? Well, it's easy if the gamble is very risky. For example, when you are always unlucky and get the three shocks, many people would choose one fixed shock instead. It's also easy when you are almost always lucky, because then you would want to choose the gamble. When your risk is inbetween, it's more tricky, and it's up to you to decide: is it worth going for the gamble or not?

But we won't tell you what your risk is. You need to find out! And you can only find out when you choose the gamble a couple of times.

In the first level of the computer game, there will be three balls with different colours. Each colour means a different risk. And your task is to find out, for each of the colours, whether you want to go for the three-shock gamble or for the one fixed shock.

There will be more instructions on the screen to explain how you are supposed to make your choice.

Part 2: Instructions on the screen for the learning part

Screen 1

You know already what the game is about. Now we will explain what you will see, and what you are supposed to do.

On each screen, a bowling ball will appear. There are three ball colours, and each ball gives you a different gamble with a different risk. So you need to decide whether you want the gamble or not.

Then someone plays the ball. You will see a grey shadow on the screen. This shadow is a bit like a TIMER, or COUNTDOWN: it will disappear after 1.5 seconds. This is exactly how much time you have to make your decision.

If you want the gamble, press the … key. If you want the fixed option, press the … key.

Press SPACE to see what it will look like.

Example screen.

Screen 2

We will start with 5 training rounds without any electric shocks, so you can get used to the task.

Press SPACE to start.

Screen 3

Now, level one of this game will start. There will be 2 blocks with 72 decisions in each block, lasting around 8 minutes. Between the blocks, there is a break. Afterwards you will enter level two.

Do you have any more questions?

If not, press SPACE to start.

Part 3: Instructions on the screen for the experiment proper

Screen 1

You have done very well in level one. Now the game will change a little.

Imagine these balls are played by bowling ball players. The grey shadow always appeared above the ball. Imagine this is a player always playing his ball straight on.

Now, new players enter the game. They are VERY BAD players. Their balls can appear anywhere on the lane.

Press SPACE to see what it will look like.

Example screen.

Screen 2

This doesn"t bother you because both balls have the same colour. So it doesn"t matter whether it was the left or the right ball that was being played. They give you exactly the same gambles as in level one.

But often, two players with different balls will be present at each end of the lane - and only one will play.

By seeing the grey shadow, you won"t know for sure which ball is being played, and which gamble you are deciding upon.

Press SPACE to see what it will look like.

Example screen.

Screen 3

But at least you can guess which ball is being played. Have you noticed the colour bars? They show that the ball is likely to land close to the players, where the bar colour is very intense. Where it is fading out, far away from the player, it is unlikely to land. Of course, once in a while it will still land here.

Press SPACE to see it again.

Example screen.

Screen 4

We will start with 5 training rounds without any electric shocks, so you can get used to the task.
Press SPACE to start.

Example screen.

Screen 5

Now, level two of this game will start. There will be 5 blocks with 72 decisions in each block, lasting around 8 minutes. Between the blocks, there will be breaks.
Do you have any more questions?
If not, press SPACE to start.

b: Instructions for the imaging experiment

All instructions were presented as printout

Part 1: Instructions for the learning part

Thanks for taking part in this experiment. It's an easy, two-level computer game.

In this game, you will be asked to make decisions about lotteries. These lotteries involve electric shocks.

Here is how it works.

This is a lottery ticket!

Take it!

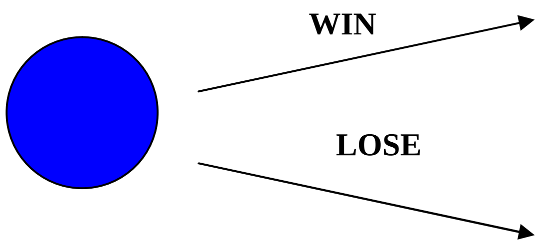

You’re lucky. Nothing happens.

You’re unlucky. You get three electric shocks.

There are three different lotteries, with blue, orange and black tickets. They all have the same payoffs - either you win and you avoid any electric shocks or you lose and you receive 3 electric shocks. But your chances of winning and losing are different with each lottery.

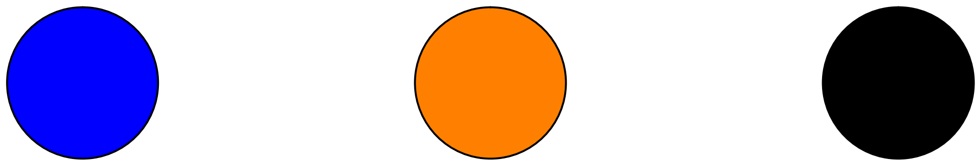

Remember, even on a good lottery you can lose. Sometimes even a few times in a row. Remember: you can flip a coin five times and get five heads – it's still a coinflip. But in the long run, your payoffs even out. And then, a good lottery gives you less shocks than a bad lottery.

Why should I want a lottery where I can only lose and win nothing?

You get this lottery for free! You don't pay for it!

But I don't want it! Can I get rid of it?

Yes you can. But it costs you one shock to give it away. You'll have to decide on each round of this game whether you want to keep the lottery or to give it away.

So how do I make my decision?

Now let's start with a very bad lottery: if you (almost) always lose and get three shocks, then its best to give it away and take one shock instead. With a good lottery you might (almost) always win, therefore its best to keep the lottery.

And when it's inbetween?

Well – thats up to you. You have to use your own judgements to try to minimise the number of shocks you recive

Do they tell me my risk?

No. Its like real life. You have to learn by experience by taking the lotteries a couple of times. And it is really important, and also in your best interest, that you find out properly, because level two of this game depends on this.

But I don't want to decide. I won't press a key.

Ok. But if you don't press the key in time, the computer will always give you three shocks.

... hmm ...!

Now let's have a look at how this works on the computer screen.

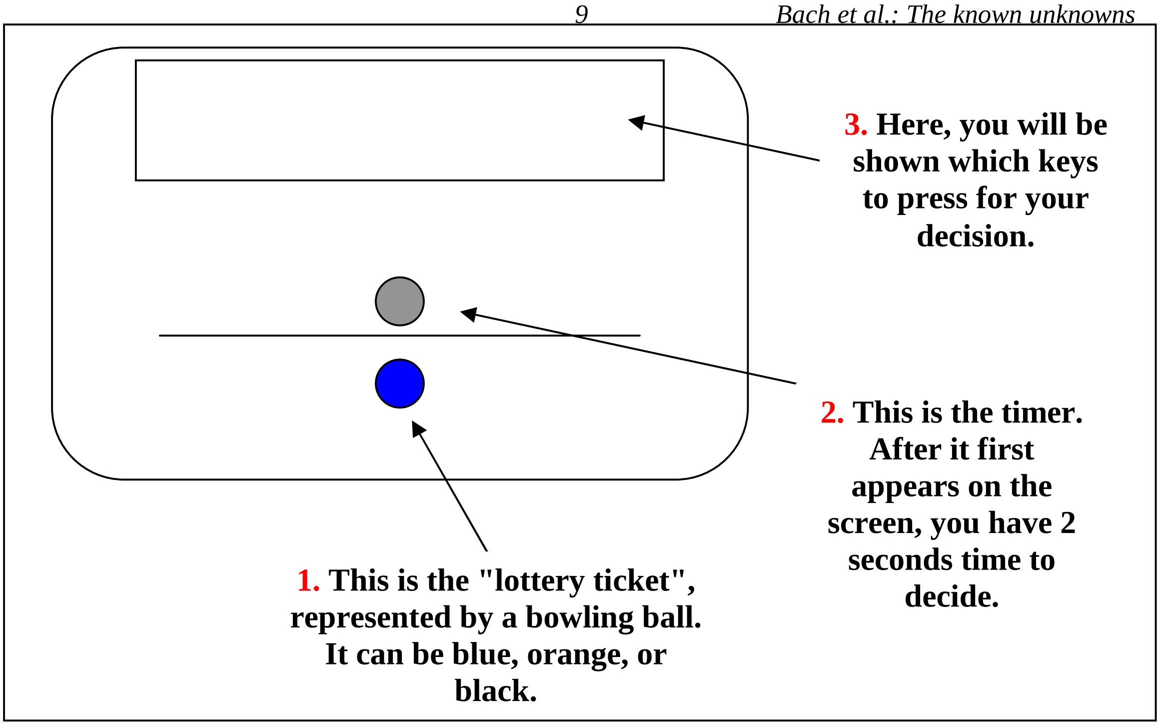

If you have any questions, please ask. It's really important that you understand the instructions properly. In this level one, there will be two sessions of 72 rounds each. Between the sessions, there is a break. Afterwards, it gets a bit more complicated. However, the odds of the three lotteries will always be constant throughout the whole game.

Part 2: Instructions for the experiment proper

It seems you have done quite well in level one of this game. Now you will proceed to level two, and it gets a bit more complicated. You will play level two in the scanner, but before that, you can practise a bit.

This is how it works.

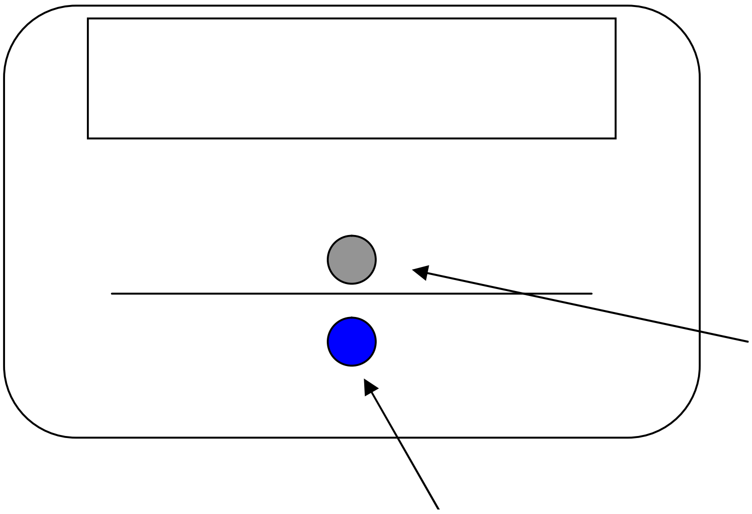

1. This shows you what colour bowling ball is being played at the moment.

2. Imagine you are looking down a bowling lane. This silhouette of the ball shows where the ball has landed at the end of the lane. Here it has directly in the middle. He's a good player!

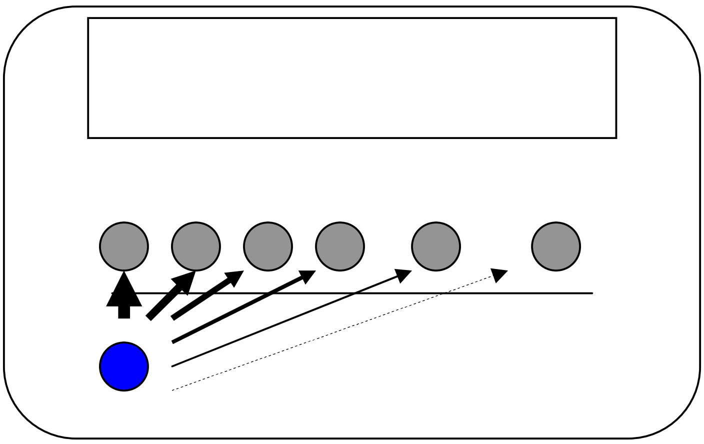

This would be a bad player. His ball can go almost anywhere. Well, it's most likely to land close to the player, but sometimes it goes right to the other side of the lane. We call this player Mr Lefty because he stands on the left of the lane, and that's where his ball lands most often.

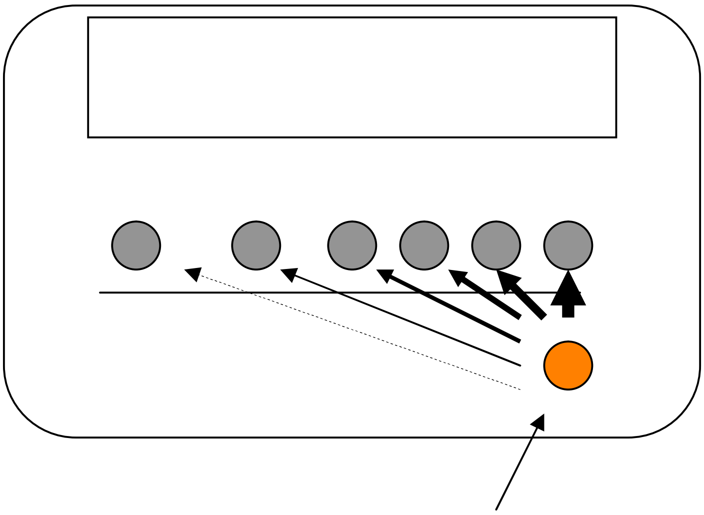

This is another bad player. We call this player Mr Righty because he stands on the right.

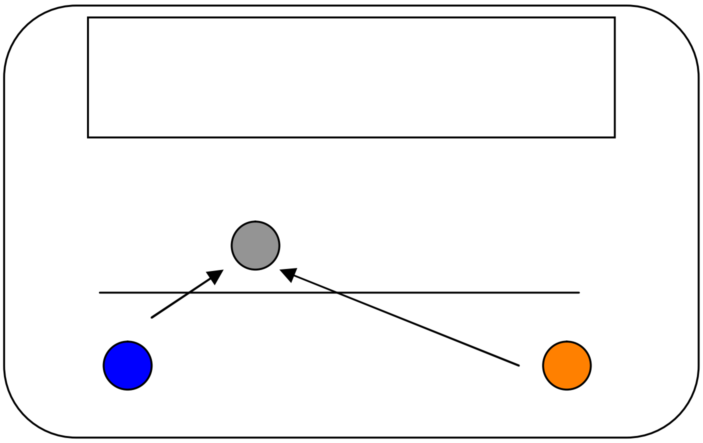

Now both players line up to play– but only one plays. You see the silhouette, but you can't see the true colour. Still you have to decide whether you want this – unknown – lottery ball or not. How can you decide? Remember it's the colour of the ball that determines your chances of winning or losing

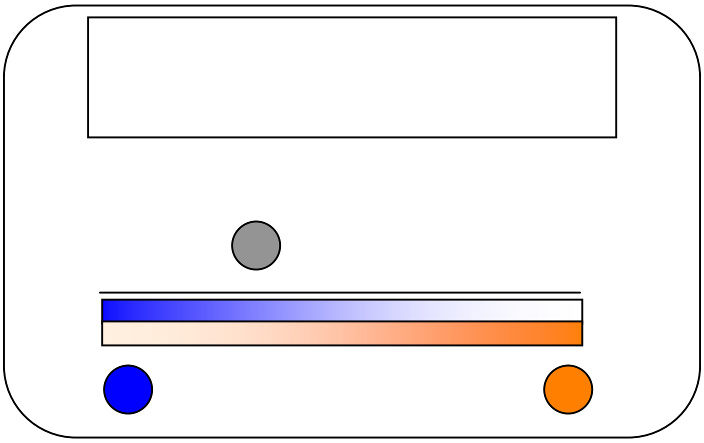

It's not difficult. The closer the shadow is to one of the balls, the more likely it is to be this ball. You see the colour bars? The more intense the colour is, the more likely it is that the ball causing the shadow at this place has this colour. When the shadow is exactly in the middle – you can't know for sure.

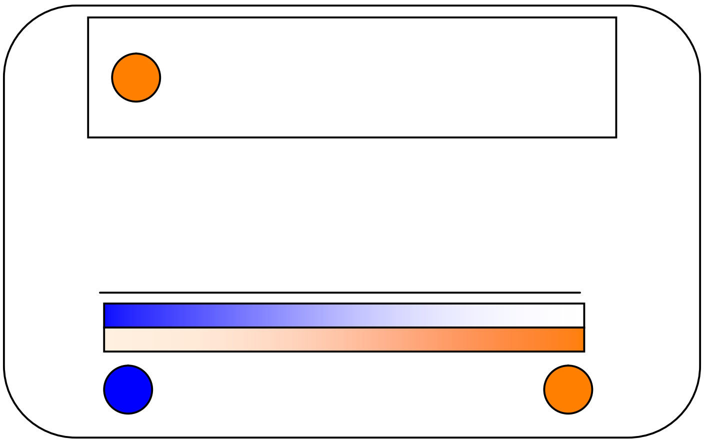

After you have made your decision (take the lottery or not), you will be shown which ball it really was. This helps you in the long run to figure out how these players play.

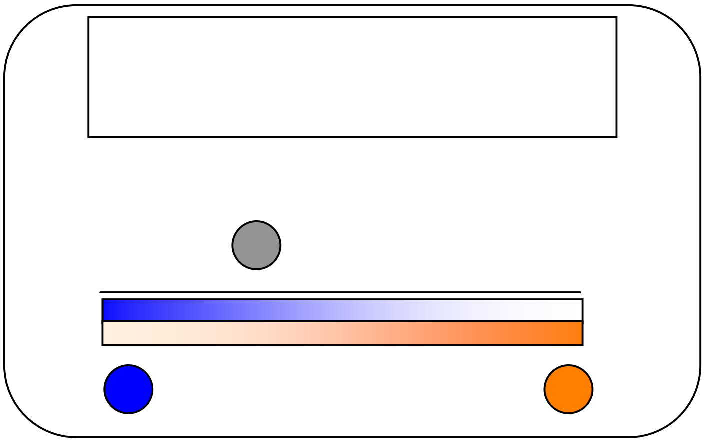

So remember: you DON'T decide between the two balls. You DON'T bet which ball it is. You simply say whether you want this lottery or not. And to make this decision, it will help you to figure out which lottery it might be.

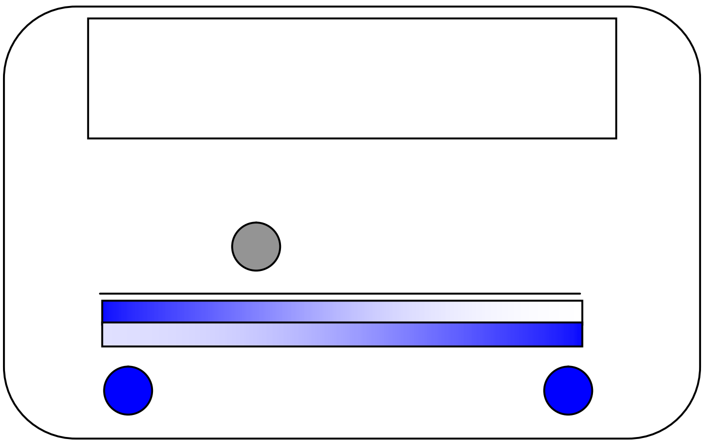

Each of the two players can play any of the three ball colours. So they might both have the same ball colour. Then it's really easy: it's just the same as in level one. That's because the odds of the lotteries never change!

Now you can practise a bit before we go into the scanner. During these five training rounds, you won't get any electric shocks.

Supplemental Methods II

a: Model comparison

(A) Assumptions

The following assumptions were made in this model comparison: (a) the value, and utility, of not getting an electric shock, was assumed to be zero. (b) unless otherwise stated, we assume the reduction of compound lotteries axiom according to expected utility theory (von Neumann and Morgenstern, 1944), that is, two-stage bets can be collapsed into a single stage bet. In this section we adopt the economic terminology where a non-ambiguous gamble is called *risky*.

(B) Model comparison

Model parameters were individually optimised for each participant and each model, using a maximum likelihood criterion and gradient search. The ensuing loglikelihood was then penalised for model complexity by using the Akaike information criterion (AIC) and Bayes information criterion (BIC). These result in two different approximations to the Bayesian model evidence. Model evidence was then compared on the group level in a random effects analysis, assuming that the true model might be different between individuals. This was implemented using group level Bayesian model selection (Stephan et al., 2009).

(C) Models accounting for a categorical effect of ambiguity

All of the following models are similar in how a choice is calculated from the net shock probability. They differ in how the net shock probability is derived in ambiguous and non-ambiguous trials, thus accounting for a categorical effect of ambiguity.

(0) Simple model, based on EU theory

This is a general model to describe how individuals compute expected utilities and transform them into choices. It uses a power rule to compute utilities and a logit function to generate choice probabilities. It does not take into account the position of the ball and averages across the two possible outcome probabilities for ambiguous trials, while for risky trials, this collapses to the unambiguously indicated outcome probability. Thus, the model does not treat risky and ambiguous gambles differently and follows the reduction of compound lotteries axiom.

$$p(choice) = \frac{1}{1 + e^{\lambda \cdot (EU(gamble) - EU(fixed))}}$$
$$EU = p(shock) \cdot V(shock)^{\beta}$$
$$p(shock) = \tfrac{1}{2} \cdot \left[p(shock \mid ball_1) + p(shock \mid ball_2) \right]$$

with

EU: expected utility

V: value ($V = -1$ for the fixed option and $V = -3$ for the gamble)

$\lambda > 0$, $\beta > 0$: subject-specific constants

(1) Baseline model, based on EU theory

As does the simple model, the baseline model does not treat risky and ambiguous gambles differently and follows the reduction of compound lotteries axiom. In contrast to the simple model, it takes shade positions into account. Outcome probability calculations are based on the true, linear density distribution of shade positions.

$$p(choice) = \frac{1}{1 + e^{\lambda \cdot (EU(gamble) - EU(fixed))}} \qquad (1)$$

$$EU = p(shock \mid position) \cdot V(shock)^{\beta} \qquad (2)$$

$$p(shock \mid position) = p(colour_1 \mid position) \cdot p(shock \mid colour_1) + p(colour_2 \mid position) \cdot p(shock \mid colour_2) \qquad (3)$$

$$p(colour_n \mid position) = \frac{p(position \mid colour_n)}{\sum_i p(position \mid colour_i)} \qquad (4)$$

with

EU: expected utility

V: value ($V = -1$ for the fixed option and $V = -3$ for the gamble)

$\lambda > 0, \beta > 0$: subject-specific constants

(2) Utility weighting model: A simple formulation of ambiguity aversion can be derived by keeping the reduction of compound lotteries, but assuming that the value, or utility, of choice options are different under ambiguity and risk (see Camerer, 1999 for an overview). With only one gamble value, this can be formulated by using different risk parameters under ambiguity and risk in equation (2). Ambiguity aversion would predict $\beta_1 < \beta_2$.

$$EU = p(shock \mid position) * V(shock)^{\beta_n}$$

with

$\beta_{1,2} > 0$: subject specific constants for risky, and ambiguous gambles

(3) Expected utility weighting model, additive: Ambiguity influences the expected utility (EU) of the gamble by adding a subject-specific constant in equation (2). Different from model (2), the effect of ambiguity does not depend on outcome probability. Ambiguity aversion would predict $a_2 < 0$.

$$EU = p(shock \mid position) \cdot V(shock)^{\beta} + a_n$$

with

$a_1 = 0$ (for risky gambles)

a_2: subject specific constant (for ambiguous gambles)

(4) Expected utility weighting model, multiplicative: Similar to model (3), ambiguity influences expected utility (EU) of the gamble by multiplying with a subject-specific constant in equation (3). Given negative values and utilities, ambiguity aversion would predict $m_2 > 1$.

$$EU = p(shock \mid position) \cdot V(shock)^{\beta} \cdot m_n$$

with

$m_1 = 1$ (for risky gambles)

m_2: subject specific constant (for ambiguous gambles)

(5) Expected probability weighting model: In this model, the expected probability is weighted in a non-linear manner (see Camerer, 1999 for an overview). The particular functional form used here was previously applied to data from a neuroimaging experiment by Hsu et al. (2005). In this formulation, ambiguity modifies expected probability by exponentiating with a subject-specific constant in equation (2). Given negative values and utilities, ambiguity aversion would predict $pc_2 < 1$.

$$EU = p(shock \mid position)^{pc_n} \cdot V(shock)^{\beta}$$

with

$pc_1 = 1$ (for risky trials)

pc_2: subject specific constant (for ambiguous trials)

(6) Pessimistic weighting model: This model assumes that in calculating expected utility from several possible outcome scenarios, some weight is put on the minimum expected outcome, as prescribed by the α-Maxmin Utility function (Ghirardato et al., 2004; applied to data from a neuroimaging experiment in Huettel et al., 2006). The original α-Maxmin Utility function does not assign crisp second order probabilities [SOPs] and assumes equal weight to worst and best outcome scenario under ambiguity-neutral conditions. However, in the present experiment, SOPs were imposed by the ball position and we showed in the first place that people do take this into account when calculating expected probabilities (comparison of baseline model against simple model). Therefore, the functional form of this model was adjusted so that the second order probability imposed by the ball position was weighted by an ambiguity constant. We call this a pessimistic weighting model, as ambiguity aversion would predict a "pessimistic" underweighting of the better outcome scenario (and ambiguity preference an "optimistic" underweighting of the worse outcome scenario). Instead of equation (3), this model prescribes, with $0 \leq \alpha \leq 2$

$$p(shock \mid position) = [\alpha \cdot p(colour_{better} \mid position)] \cdot p(shock \mid colour_{better}) + \\ [1 - \alpha \cdot p(colour_{better} \mid position)]) \cdot p(shock \mid colour_{worse})$$

(i.e. ambiguity aversion) for $\alpha \leq 1$, and

$$p(shock \mid position) = [(2-\alpha) \cdot p(colour_{worse} \mid position)] \cdot p(shock \mid colour_{worse}) + \\ [1 - (2-\alpha) \cdot p(colour_{worse} \mid position)]) \cdot p(shock \mid colour_{better})$$

(i.e. ambiguity preference) for $1 < \alpha \leq 2$

(7) Minimax model: This is a simpler form of model (6) and assumes that only the minimum expected outcome is taken into account. This model necessarily predicts ambiguity aversion. Thus, equation (3) collapses to

$$p(shock) = p(shock \mid colour_{worse})$$

(8) SOP model: This model assumes unique SOPs but relaxes the reduction of compound lottery axiom, as proposed by Segal (1987). Recall that EU theory asserts that given the choice between a certain and a risky option with the same expected value, a risk-averse decision maker will choose the certain option because his utility function is concave and the expected utility is therefore lower for the risky option. A similar idea applies in this model, where EU is computed separately for all possible values of the first order probabilities $p_{1\ldots n}$, which are then separately weighted by a mapping similar to an EU function, and these mappings are combined using the assigned SOPs for each p. The particular functional form used here draws on the model proposed by Klibanoff et al. (2005) which was simplified in the sense that it only works on probabilities, not utilities, which in our experiment is equivalent to the original model as only one non-zero utility value was used in the gamble. Thus, equation (3) from the baseline model is modified, and – given negative values and utilities - predicts ambiguity aversion if α < 0.

$$\begin{aligned} p(shock \mid position) = {} & p(colour_1 \mid position) \cdot \varphi(shock \mid colour_1) + \\ & p(colour_2 \mid position) \cdot \varphi(shock \mid colour_2) \end{aligned} \qquad (5)$$

$$\varphi(shock \mid colour_n) = \frac{1 - e^{-\alpha \cdot p(shock \mid colour_n)}}{1 - e^{-\alpha}} \qquad if \alpha \neq 0$$

$$\varphi(shock \mid colour_n) = p(shock \mid colour_n) \qquad if \alpha = 0$$

(D) Models accounting for the effect of entropy

The following models all account for categorical ambiguity with the SOP mechanism described in model C8. They differ in whether and how they additionally account for entropy.

Premise for the analysis was that medium entropy should have no effect on decisions other than the general effect of ambiguity as opposed to risk. Thus, we drew on the winning model from (C), i.e. SOP model (8), and modified it to accommodate the effect of *H*. The zero effect of a medium entropy was ensured by using the quantity

$$H' = H - mean(H)$$

which was set to zero for risky gambles.

(1) SOP model without accounting for entropy: Model C8 was used as a reference for the model comparison.

(2) Utility weighting model: This model assumed that the risk preference parameter in equation (2) was varied on ambiguous trials by entropy. Given negative values and utilities, aversion of entropy would predict $\beta_1 > 0$.

$$\beta = \beta_0 + \beta_1 \cdot H'$$

$\beta_0 < 0$: individual risk parameter

β_1: scaling factor of the effect of entropy.

(3) Expected utility weighting model, additive: This model assumed that entropy changes the EU in equation (2) by adding a variable amount. Given negative values and utilities, entropy aversion would predict a < 0.

$$EU = p(shock \mid position) \cdot V(shock)^{\beta} + H' \cdot a$$

(4) Expected utility weighting model, multiplicative: This model assumed that entropy changes the EU in equation (2) by multiplying with a variable amount. Given negative values and utilities, entropy aversion would predict m > 0.

$$EU = p(shock \mid position) \cdot V(shock)^{\beta} \cdot (1 + H' \cdot m)$$

(5) Expected probability weighting model: In this model, the expected probability is weighted in a non-linear manner. Given negative values and utilities, entropy aversion would predict *pc* > 1.

$$EU = p(shock \mid position)^{1+(H' \cdot pc)} \cdot V(shock)^{\beta}$$

(6) Pessimistic weighting model: Based on the pessimistic weighting model from C, this model assumes that the second order probability imposed by the ball position was weighted by entropy, such that when entropy is above average, the worse outcome is overweighted, and vice versa when entropy is below average.

Instead of equation (3), this model prescribes:

$$\begin{aligned} p(shock \mid position) = {} & [(1 + H' + \alpha) \cdot p(colour_{better} \mid position)] \cdot \varphi(shock \mid colour_{better}) + \\ & [1 - (1 + H' + \alpha) \cdot p(colour_{better} \mid position)]) \cdot \varphi(shock \mid ball_{worse}) \end{aligned}$$

for H'α ≤ 1, and

$$\begin{aligned} p(shock \mid position) = {} & [(1 - H' + \alpha) \cdot p(colour_{worse} \mid position)] \cdot \varphi(shock \mid colour_{worse}) + \\ & [1 - (1 - H' + \alpha) p(colour_{worse} \mid position)]) \cdot \varphi(shock \mid ball_{better}) \end{aligned}$$

for H'α > 1, and

(7) SOP model: This model assumes that the first order probability weighting function (C8) accounts for both ambiguity aversion, and entropy aversion. The probability weighting function is modified by entropy:

$$p(shock \mid position) = p(colour_1 \mid position) \cdot \varphi(shock \mid colour_1) + p(colour_2 \mid position) \cdot \varphi(shock \mid colour_2) \qquad (5)$$

$$\varphi(shock \mid colour_n) = \frac{1 - e^{-\alpha' \cdot p(outcome \mid colour_n)}}{1 - e^{-\alpha'}} \qquad if \alpha' \neq 0$$

$$\varphi(shock \mid colour_n) = p(shock \mid colour_n) \qquad if \alpha' = 0$$

$$\alpha' = \alpha_1 + H' \cdot \alpha_2$$

Supplemental Results

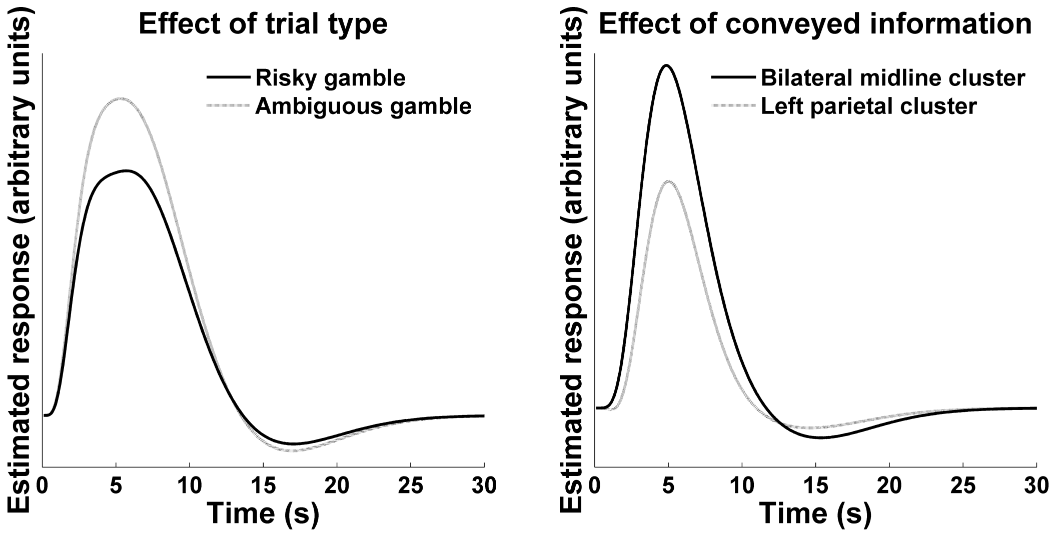

Figure S1

Estimated BOLD responses for the clusters shown in figure 3 of the main text. Left: Estimated responses in the left posterior parietal cortex to non-ambiguous and to ambiguous gambles. Right: Estimated effect of negative entropy on BOLD responses, for a large bilateral midline cluster as well as for a separate cluster in the left parietal cortex.

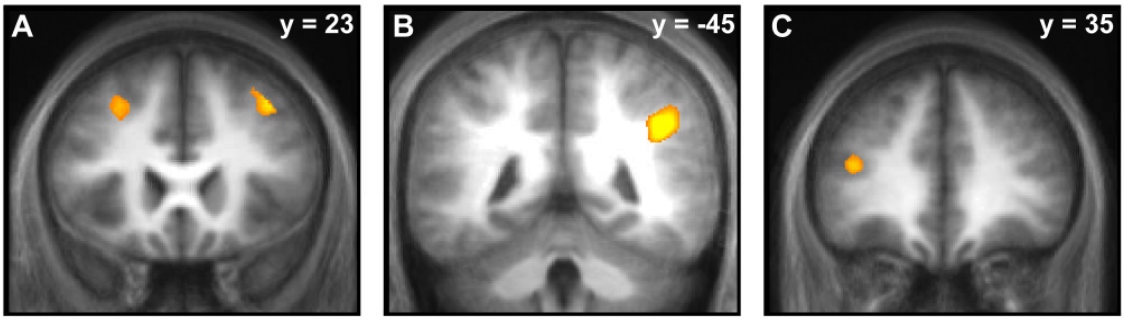

Figure S2

In the imaging experiment, we measured BOLD responses to outcome probability (and thus, negative utility and value) in the non-ambiguous trials (panels A and B). With three objective outcome probabilities (0.2, 0.5, 0.8), the possible outcome probabilities in ambiguous trials could either have a range of 0.3 (0.2 – 0.5 and 0.5 – 0.8) or 0.6 (0.2 – 0.8). Stronger BOLD responses during trials with lower outcome probability range than during trials with higher outcome probability range were observed in the cluster shown in panel C. The clusters shown in panels A and B survived whole-brain correction at a cluster-level threshold of $p < .05$ and a voxel-level threshold of $p < .001$. The cluster shown in panel C survived small volume correction at $p < .05$ around a sphere with 15 mm radius around peak coordinates reported previously (Huettel et al., 2006).

Table S1

Blood oxygen level dependent (BOLD) responses to entropy of the ball silhouette, and surprise of the ball outcome within ambiguous trials. All clusters are reported cluster-level corrected for family-wise error (*FWE*) and with $p < .05$, one cluster at $p = 0.06$ (marked with *).

Brain regions	Brodmann area of local maxima	Hemi-sphere	Voxel number	Voxel *t* score	Montreal Neurological Institute brain template coordinates of local maxima
Responses to low second-order uncertainty (i.e. entropy) within ambiguous gambles after controlling for the surprise of the outcome					
Posterior cingulate, superior parietal lobule, postcentral gyrus, parahippocampal gyrus, precuneus, middle temporal gyrus, angular gyrus, superior & middle occipital gyrus,	30 & 31	Bilateral	4925	7.39	16, -58, 6; -20, -62, 8; 24, -56, 10
Responses to high surprise at the time point of the outcome					
Superior parietal lobule	7 & 12	Left	191	5.63	-24, -60, 42; -32, -64, 50; -36, -54, 42
Responses to low surprise at the time point of the outcome					
Parahippocampal gyrus	19	Right	358	7.17	24, -58, -4; 4, -62, -4; 12, -70, -4
Posterior cingulate, caudate tail	29	Right	130	6.00	18, -38, 20; 8, -38, 14
Caudate nucleus		Right	96	5.18	34, -38, 2

Reference List

1. Camerer C (1999) Ambiguity aversion and non-additive probability: Experimental evidence, models and applications. In: Uncertain decisions: Bridging theory and experiments (Luini L, ed), pp 53-80. Boston: Kluwer.

2. Ghirardato P, Maccheroni F, Marinacci M (2004) Differentiating ambiguity and ambiguity attitude. Journal of Economic Theory 118: 133-173.

3. Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF (2005) Neural systems responding to degrees of uncertainty in human decision-making. Science 310: 1680-1683.

4. Huettel SA, Stowe CJ, Gordon EM, Warner BT, Platt ML (2006) Neural signatures of economic preferences for risk and ambiguity. Neuron 49: 765-775.

5. Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. Econometrica 73: 1849-1892.

6. Segal U (1987) The Ellsberg Paradox and Risk-Aversion - An Anticipated Utility Approach. International Economic Review 28: 175-202.

7. Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. Neuroimage 46: 1004-1017.

8. von Neumann J, Morgenstern O (1944) Theory of games and economic behavior. Princeton, NJ: Princeton University Press.