Article:
Lu, Y, Stafford, T and Fox, C (2016) Maximum saliency bias in binocular fusion. Connection Science, 28 (3). pp. 258-269. ISSN: 0954-0091

https://doi.org/10.1080/09540091.2016.1159181

To appear in *Connection Science*
Vol. 00, No. 00, Month 20XX, 1–11

Maximum saliency bias in binocular fusion

(*Received 00 Month 20XX; accepted 00 Month 20XX*)

Subjective experience at any instant consists of a single ('unitary'), coherent interpretation of sense data rather than a 'Bayesian blur' of alternatives. However, computation of Bayes-optimal actions has no role for unitary perception, instead being required to integrate over *every* possible action-percept pair to maximise expected utility. So what is the role of unitary coherent percepts, and how are they computed? Recent work provided objective evidence for non-Bayes-optimal, unitary coherent, perception and action in humans; and further suggested that the percept selected is not the maximum *a posteriori* (MAP) percept but is instead affected by utility. The present study uses a binocular fusion task first to reproduce the same effect in a new domain, and second, to test multiple hypotheses about exactly *how* utility may affect the percept. After accounting for high experimental noise, it finds that both Bayes optimality (MEU) and the previously proposed maximum-utility (MU) hypothesis are outperformed in fitting the data by a modified maximum-salience (MS) hypothesis, using unsigned utility magnitudes in place of signed utilities in the bias function.

Keywords: Bayesian brain; psychophysics; utility; bias; perception; binocular fusion

1. Introduction

Our experience of the world is 'unitary coherent'. *Unitary* means we only perceive one interpretation at a time (e.g. either a face *or* a vase in the Rubin Vase illusion) rather than a blur of multiple possible interpretations (never the face *and* vase together). *Coherent* means that we almost always perceive scenes comprised of parts which do not contradict one another (e.g we do not see part face and part vase). While perception's unitary coherent (UC) nature may seem *prima facie* obvious from subjective experience, such experience is notoriously unreliable and ill-defined (No, Pessoa, & Thompson, 2000). Further, UC perception is not required at all by optimal decision making, which must consider *all* possible interpretations of sensory data (Bernardo & Smith, 2000) to maximize *expected* utility (MEU strategy), without reference to any single privileged unitary state,

$$a_{MEU} = \arg_a \max \int_s U(a,s)P(s|d)ds, \tag{1}$$

where d is sensory data, s are states, U is utility and a are actions. [1] So what are our UC percepts for? MEU inference is generally NP-hard (Cooper, 1990) and so unitary coherent perception seems likely to be involved in some approximation to optimal action selection. What kind of approximation is this? In machine perception, such as machine

[1] Sampling approximations to MEU ignore some improbable interpretations (Spiegelhalter, Thomas, Best, & Gilks, 2008), but still sum over multiple samples of possible world states s rather than privileging any particular unitary state.

vision (Felzenszwalb & Huttenlocher, 2005) or speech recognition (Young et al., 2006), the unitary coherent maximum a posteriori (MAP) state is often reported as final output,

$$s_{MAP} = \arg_s \max P(s|d). \tag{2}$$

However agents usually have percepts to help make actions rather than for their own sake. In these cases, perceiving the MAP state does not necessarily lead to the best action if the following *MAP strategy* is used,

$$a_{MAP} = \arg_a \max U(a, s_{MAP}). \tag{3}$$

A recent study (Fox & Stafford, 2012) provided objective evidence of the existence of UC perception, removing the need for reference to 'subjective experience' in the UC debate. It further proposed a particular form of UC perception and action, differing from both MEU action and MAP action, called *maximum utility perception (MU).* In MU a UC state and action are chosen together,

$$(s_{MU}, a_{MU}) = \arg_{s,a} \max U(a, s)P(s|d). \tag{4}$$

For example, suppose you see your neighbour, whose state (s) may be happy or sad. This state is not observable directly but data (d) from his facial expression gives uncertain evidence about it. Approaching the neighbour to start a conversation, avoiding him, or hedging by just saying hello, are three possible actions (a), and you may assign estimates of the utility (U) of each action in each state as in Table 1. MEU, MAP and MU will each choose different actions under these utilities due to their differences in maximizing or integrating over the possible states.

Fox & Stafford (2012) used the ambiguous states (s) of a Necker cube image (d) to produce multiple perceptual possibilities and showed MU actions gave higher utility than MAP actions, and were a better fit to human actions than either MEU or MAP actions. However the experiment left open the possibilities that the effect might be specific to Necker cube percepts and that further action selection strategies could fit the data even better than MU.

The present study seeks to address these issues by using a different type of visual ambiguity – binocular fusion (Maier, Panagiotaropoulos, Tsuchiya, & Keliris, 2012)[2]. Corroborative evidence of UC perception and action in this new visual task, when taken together with the previous study (Fox & Stafford, 2012), would suggest its generality across many visual tasks. The binocular task has a simpler action space than the Necker cube task to allow more precise noise modelling, and is further designed to be able to differentiate between the MU form of UC and a new related hypothesis – maximum salience – by including negative utilities.

It is important that as in (Fox & Stafford, 2012), findings concern only the 'everyday' system of immediate perception and action: of perception of complex objects *on the scale of fractions of seconds*, such as glancing at a new scene. In contrast it is well-known that MEU-like behaviour occurs both for simpler, stimulus-response type tasks on shorter times scales in specific environments (Körding & Wolpert, 2004); and also in complex percept-action on longer timescales, during which multiple stable percepts and actions can be perceived and deliberated over sequentially (Griffiths, Chater, Kemp, Perfors, &

[2]Binocular 'rivalry' is a special case in which a UC percept is seen that matches one of the eye inputs exactly. Pre-supposing the presence of rivalry would not allow MEU to be tested, so the more general term 'fusion' is used throughout this paper.

Tenenbaum, 2010; Chater, Tenenbaum, & Yuille, 2006; Lepora et al., 2012; Eisenfhr, Weber, & Langer, 2010). The present study does not dispute the existence of these other systems but corroborates the existence of an additional system as in (Fox & Stafford, 2012; Lengyel & Dayan, 2008). Similar findings were reported in an fMRI study (Fleming, Whiteley, Hulme, Sahani, & Dolan, 2010) but without explanation via specific strategy models such as MU and MS. As in (Fox & Stafford, 2012) and psychophysics in general, experimental noise is very high due to the need of very brief stimuli. New noise models are used to analyze this, requiring fewer assumptions that the method of (Fox & Stafford, 2012).

2. Varieties of UC perception and action

MAP and MU are not the only possible types of UC perception, and we may consider a new, generalized class of strategies which replace the utility term of eqn. 4 with other bias functions $B(a, s)$ of action and state:

$$(s_{MB}, a_{MB}) = \arg_{s,a} \max B(a, s)P(s|d). \quad (5)$$

The simplest $B(a, s) = 1$ gives the MAP strategy. $B(a, s) = U(a, s)$ gives the MU strategy. Choosing instead $B(s) = |U(a, s)|$ we may define a new *maximum saliency, MS* strategy,

$$(s_{MS}, a_{MS}) = \arg_{s,a} \max |U(a, s)|P(s|d). \quad (6)$$

This might be more useful than MU in cases where negative utilities are present. For example, consider deciding to reach for an object which may be a sleeping predator or prey with equal probability. MAP would perceive one percept, either predator or prey with a 50% chance each. MU would perceive prey, because it is part of the best-case scenario. But MS is more prudent, seeing the predator, as the large negative utility from waking it and being eaten is more salient than the positive utility of reducing our hunger. Negative utilities were not present in the previous study (Fox & Stafford, 2012) so it was unable to distinguish between MU and MS: they give identical predictions in tasks where all utilities are positive. By introducing negative utilities in the present study we can newly discriminate between MU and MS. The present study tests for MS against MAP, MEU and MU in a controlled environment to determine which best resembles observed human actions. No experiment can measure human perceptions directly, but by measuring actions we can make inferences about the accompanying content of perception within these hypotheses.

3. Methods

Overview. The present study uses controlled stimuli and forced-choice actions in an artificial binocular fusion task to simplify presentation and analysis. The four hypotheses to be tested against one another are MAP, MEU, MU and MS as in equations 2,3,4 and 6. Data, d, are image pixels, states are $s \in \{happy, sad\}$ of faces generating the pixels, and actions, a, are button presses. A system of training and examination introduced by (Fox & Stafford, 2012) is used to guarantee that subjects understand the utility function and that their use of it is close to optimal, before collecting data.

(a) unmasked happy and sad faces. (b) full-masked stimulus. (c) half-masked happy and sad.

Figure 1. Three types of stimulus used in the experiment.

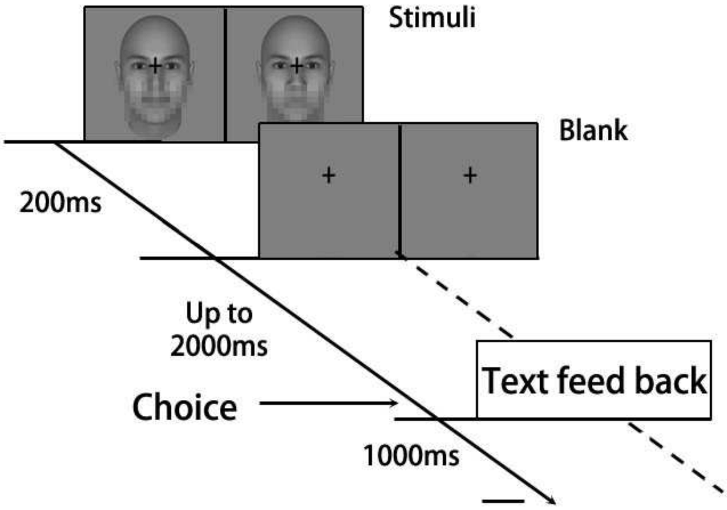

Figure 2. Trial structure. A stereo pair is presented, then blanked, then the user's choice is recorded.

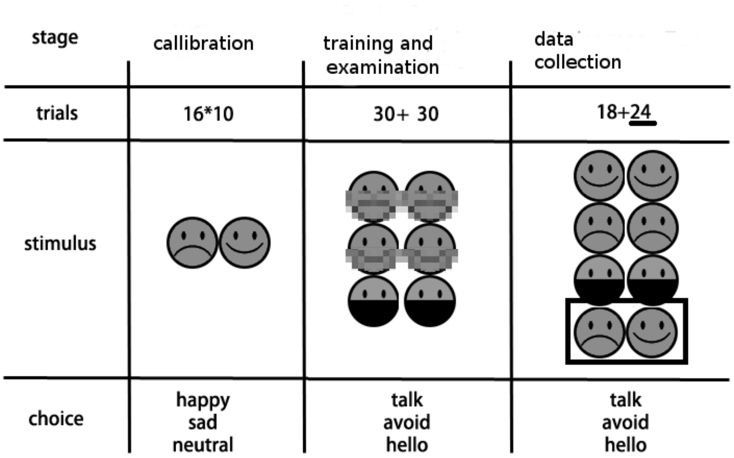

Figure 3. Structure of whole study, only data from boxed/underlined trials are analysed. Blacked-out mouth=full mask; pixellated mouth=half-masked. The three stages proceed sequentially: calibration, training and examination, then data collection. Training and examination is repeated until the examination is passed.

Stimuli. Stimuli used in all experiments were pairs of faces, as in fig. 1, which may be unmasked, half-masked or full-masked. The faces in the pair may be the same or different states ('rivalrous') but always have the same degree of masking. *Full masking* is a low-resolution pixellation of the mouth makes the happy/sad state completely unknown; *half-masking* is a higher resolution pixellation which makes it visible with some visual effort, typically a few seconds of full attention to become certain of the happy or sad mouth direction. Binocular rivalry is well-known to lead to reports of subjective experiences of seeing one face or the other.[3] Subjects viewed the stimuli via a pair of dichoptic displays from 70cm through red-green anaglyphic glasses, seeing one face in each eye. Two faces of neutral gender generated by Facegen Modeller (www.facegen.com) were used to generate the stimuli (fig. 2(a)). Stimuli were presented against a uniform gray background using the MATLAB Psychophysics Toolbox (www.psychtoolbox.org) on an LCD monitor (1024×768 pixels, 60Hz).

Trials. A *trial* means one presentation of a face pair and the collection of one button press from the subject. Fig. 2 shows the process for a trial. A fixation cross ($0.5^\circ \times 0.5^\circ$) was presented to each eye. Test images subtend $12^\circ \times 8.5^\circ$ visual angle. Distance from the fixation cross to the mouse of the face stimulus was constantly 4.5° visual angle. Subjects were instructed to stare at the fixation cross constantly. Face pairs are presented to left and right eye simultaneously, centred on the cross, for 200ms, followed by a maximum 2000ms blank screen, during which subjects made a decision to press a button according to their latest instructions. After a button press, the program waited for another 1000ms then the next trail begun.

Subjects. 23 subjects were recruited, all college students in the Department of Psychology, University of Sheffield. Ages ranged 18-28 years except one 50 year old female, all with normal or corrected-to-normal vision.

Calibration. To simplify the analysis, we would like all rivalrous stimuli to have $P(s) = 0.5$ for both of the coherent states $s \in \{happy, sad\}$. To achieve this, stimuli were adjusted per subject so that in the absence of any utility bias, each subject perceived roughly equal numbers of happy and sad faces when confronted with half-masked rivalrous pairs. An iterative process tailored the stimuli to achieve this by countering biases due to eye dominance, color dominance, or emotional preference. Calibration trails used conflicting pairs of happy and sad faces only, and subjects reported their percepts directly via two buttons marked HAPPY and SAD. Images were adjusted until approximately half each of happy and sad faces were reported. We calibrated both for any possible eye dominance and for emotional preference. An annealed gradient descent algorithm was used to adjust image parameters. Calibration contained 16 blocks, each block contained 10 trials, comprising 5 'left happy / right sad' trials and 5 'right happy / left sad' trials. During each block, if information on one eye had more dominant influence on subject's choice, the image contrast on this eye would be impaired. Similarly, if a subject saw one emotion (happy or sad) more frequently, the image contrast on this emotion would be impaired. The adjustable parameters were: eye bias DB and emotional bias EB. Another 2 parameters $C1$ and $C2$ are used to record subject's behaviour in one block of the calibration. $C1$ record the 5 left happy right sad trials. In each LHRS trial, if subject recognised the contradictory visual stimuli as an unitary coherent happy face, subject would choose 'happy' and C1=C1+1; if subject see unitary coherent sad face, he/she would choose 'sad' and C1=C1-1; if subject see neutral or contradictory or uncertain

[3]We make no use of potentially unreliable subjective reports, only of directly observed, utility-driven actions. Percepts can then be inferred from the actions under the various models being tested. If subjects had difficulty seeing the mouth direction in half-masks, the degree of masking was adjusted to make it just possible for them at the start of the experiment, before calibration.

Percept/action	APPROACH	HELLO	AVOID
:-)	10	1	-4
:-(	-20	1	3
Average	-5	1	-0.5

Table 1. Payoff matrix used in action-report experiment.

about the emotion, he/she would choose 'uncertain' and C1 stays the same. $C2$ record the same in right happy left sad trials (RHLS trials). The following annealed gradient descent updates where N is the number of the trial within the block,

$$\triangle DB = e^{\frac{C1-C2}{3N}}, \triangle EB = e^{\frac{C1+C2}{3N}}. \tag{7}$$

Training and examination. After calibration, subjects are provided with three new buttons and given the following text:

Every day your neighbour is in his garden. Sometimes he is in a good mood and sometimes a bad mood. If he is in a good mood then it is good to start a conversation with him (press TALK button). If he is in a bad mood it is best to avoid him (press AVOID button). Or you can just say hello (press HELLO button), this is quite good if he is in a good mood but not too bad if he is in a bad mood.

The scoring system for percept-action pairs in table 1, was carefully designed to discriminate between our different hypotheses, such that the different strategies under consideration lead to different actions for rivalrous percepts. As in (Fox & Stafford, 2012) subjects were required train then pass an 'examination' to prove they have fully internalized this scoring system. In training blocks of 30 trials, stimuli on left eye and right eye are the same. 10 are half-masked happy faces, 10 are half-masked sad faces and 10 are full masked random faces.[4] Scores for each trial are printed on the screen immediately after each action. Each training *round* has two blocks, score maximization and score minimisation, where subjects were asked by the experimenter to obtain highest and lowest scores possible. Subjects perform repeated training rounds until they achieve threshold scores of 98 in maximisation and negative 200 (-200) in minimisation. These thresholds are set such that they can only be achieved by doing the optimal action in *every* trial. Therefore we can be certain that any subject passing this 'examination' understands, and can act optimally using, the scoring system. Importantly this includes use of the hedging *HELLO* action when faced with the rivalrous full-masked stimuli. (All subjects passed the exam eventually).

Data collection. Following successful completion of the exam, subjects were presented with a further 42 similar trials. The buttons and scoring were identical to the training rounds. However subjects are asked only to maximise their total score, and the face pairs used are different. Here we use the following image pairs: (*happy,happy*), (*sad,sad*), (*happy,sad*), (*sad,happy*), (*full-mask, full-mask)*, where *happy* and *sad* are half-masked. Only the 24 rivalrous pairs are required in this paper's analysis, the other 18 may be considered as distractors.[5]

[4]Half-masking was used here because in pilot experiments, subjects trained on unmasked tasks tended to notice something was different in the data collection phase and perhaps realise that a 'trick' is being played on them by the experimenter, but half-masking was found to remove this effect because it accustoms subjects to the presence of noise in general.

[5]Criticism of the experimental logic might argue that the training and experiment tasks are different, because 'bait and switch' is applied to the stimuli between them, and that subjects have therefore been trained to rely on a unimodal distribution and are unfairly switched to a task relying on a bimodal one. However this is precisely the point of the experiment. The underlying probabilistic *model* required to solve the tasks does not change, only

(a) single faces with impaired happy and sad emotion.

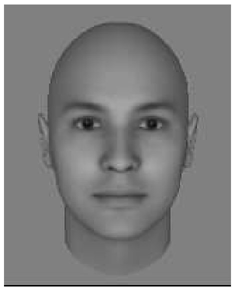

(b) neutral emotion.

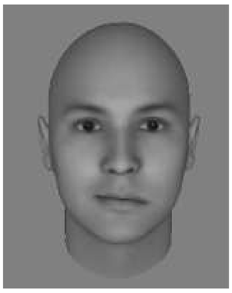

(c) One of many examples that subjective experience of half happy and half sad.

(d) The example of subjective experience of contradictory face with smiling and sad mouth overlapping each other.

Figure 4. Mockups of reconstructed unitary coherent subjective perception after binocular fusion

Debriefing. Subjects were asked for informal reports of their subjective experiences at the end of the whole experiment.

A summary of the trial types used in the experimental stages is shown in fig. 3.

4. Results

Debriefing. Binocular rivalry is sometimes assumed to be a perfect form of perceptual ambiguity, however subjects did report experiences of other types of fusion than rivalry, disclosed here: sometimes a less-happy or less-sad face, see fig. 4;(a); sometimes a neutral face, see fig.4(b); sometimes a reconstruction into other combinations of happy and sad faces, see fig.4(c); in very rare cases, subjects might see a unitary but non-coherent face, see fig.4(d). This was reported by 3 of 23 subjects. Reports of 'no percept' were common due to the short presentation time (shorter than in the examination). Also the long data collection period may have induced boredom in some subjects leading to giving up on the task and to completely random actions. All subjects were found naive to the purpose of the study, and unaware of binocular fusion theory.

Testing hypotheses against the data must take account of experimental noise. This is because for each of the pure hypotheses, there is at least one action it would never make, but the results do contains mixtures of all the actions. From debriefing it was clear that some stimuli were not perceived at all or perceived in ways not included in

the shape of the distribution *over* this model. The fact that unimodal distributions *can* be used is sufficient to accept the existence of UC perception in at least these tasks.

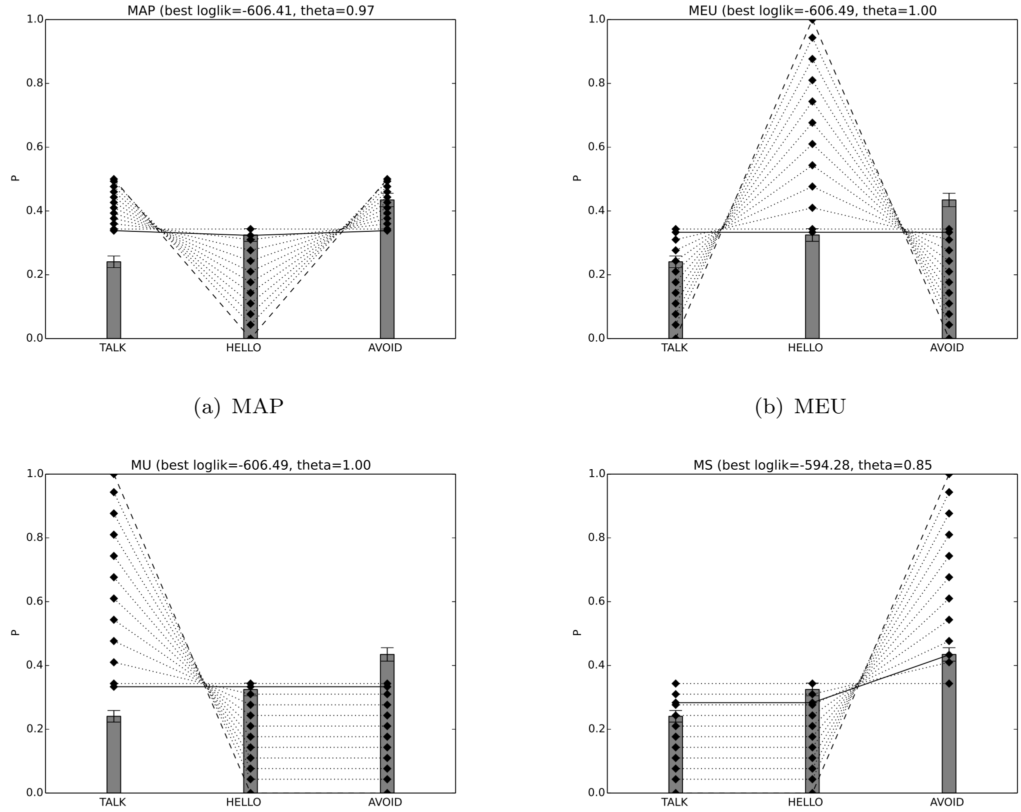

(a) MAP

(b) MEU

(c) MU

(d) MS

Figure 5. Fitting four hypotheses to the Action Report task: MAP, MEU, MU and MS. Observed responses from 552 trials are plotted as gray bars, as beta posteriors of their generators with 1 standard deviation bars. Each hypothesis has one free parameter, θ. Choosing a value of θ specifies a model with the hypothesis. Dotted lines show models with $P(\text{noise})$ = 10%:10%:100%; dashes show the noiseless model $\theta = 0$. The best fitting θ model within each hypothesis is shown by a bold line; its value and log likelihood are above the graph. All models are discrete, each describing only three numerical probabilities over the three actions as shown by the diamonds – the continuous dotted, dashed, and thick lines are used only to indicate which predictions come from which model.

the hypotheses. Such experimental noise is very common in forced-choice psychophysics experiments, and is easily accounted for.

We therefore consider each *hypothesis* ($h \in \text{MAP}, \text{MEU}, \text{MU}, \text{MS}$) to be set of *models*, $\{h(\theta)\}_\theta$, having continuous values of a free parameter θ. The meaning of $\theta = P(\text{noise})$ is the probability that a trial's action is a random button press rather than the optimal action a_h of the hypothesis. For each model in a hypothesis' set, the likelihood $P(data|h(\theta))$ can be computed by

$$P(data|h(\theta)) = \prod_a P(a|h(\theta))^{|a|}, \tag{8}$$

where $|a|$ is the number of trials in which action a was taken in the observed *data*. We consider the set of all rivalrous trials as *data*, other trials are distractors. Making minimal assumptions, a flat prior is taken over hypotheses and noise values.

Response prediction from a selection of these models are shown in fig. 5. Dotted lines show predictions using ten values of θ from 10% to 100% in steps of 10%. The pure,

$\theta = 0$ model shown by the dashed line: for MEU, MU and MS these predict different single actions, and for MAP predict a mix of two actions. For each hypothesis, the optimal θ is found by numeric search. Its value and corresponding log likelihood – the best log likelihood for the hypothesis whole model set – is shown above the graph, and its prediction as a thick line.

We seek the single best model from the four hypotheses' combined sets of models. To find it, it suffices to find the best-fitting models within each of the four hypotheses, according to their log-likelihoods, then compare these four best models against one another. It can be seen from the log-likelihoods in fig. 5 that the overall best model is MS($\theta = 0.85$). This has a log-likelihood of 12.2 greater than any of the other hypotheses' models, so is $e^{12.2} \approx 200,000$ times more likely than any of them.

In addition to the four models under test, we considered two further models which may help to reduce or give insight into the noise levels and alternative future model classes.

Perfect two-parameter model. The four models tested each have one free parameter, θ, and we note that as there are only three possible actions, a perfect fit can be made by a two-parameter model (2PM) which simply assigns frequentist probabilities to two actions $P(TALK)$ and $P(HELLO)$ according to their observed counts, and thus $P(AVOID) = 1 - P(TALK) - P(HELLO)$. This model is obviously over-fitted to the data, but may be analyzed using Bayesian Information Criteria scores, $BIC = -2 \ln L + k \ln n$. The log likelihood $\ln L$ of 2PM on the data was found to be -585.71, compared with -594.28 for the MS model, a difference of 5.13 in favour of 2PM. BIC with $n = 552$ data points and k free parameters yields penalty terms (Occam factors) of $1 \times \ln 552 = 6.3$ and $2 \times 6.3 = 12.4$ respectively. This implies that the best possible two parameter model, and thus *any* two parameter model – such as any weighted mixture of strategies – has a worse BIC score than MS so can make no claim to be a better model than MS. However the BIC difference is not significant so this does not rule out two parameter models, it only shows they are not better than MS. For now we argue that MAP,MEU,MU and MS should have higher *subjective* priors because they are theory-driven rather than just fitted to data, and we exclude 2PM for this reason, but future larger data could resolve this without the need for subjective experimenter assumptions.

Per-subject model. The present study was not intended to explore per-subject effects, and the data sizes were selected to obtain significant population results at minimal financial experiment cost. However we can examine what, if any, power the present data has to inform per-subject models. While we cannot test every possible per-subject model, a reasonable model might be a Per Subject Switching (PSS) model which assigns to each subject the best fitting choice of MAP,MEU,MU or MS. In particular this would separate out the effects of some bored subjects abandoning the task from less noisy subjects. Assigning one indicator variable to each of 23 subjects, and assuming a single free noise parameter θ shared over all of MAP,MEU,MU and MS, gives gives $k = 24$. Optimizing over θ, the best log likelihood of PSS on the data was found to be -508.87 (at $\theta = 0.50$), a difference of 85.4 in favour of PSS over MS. However as with 2PM, this gain is outweighed by a large penalty term $k \ln n = 151$ which makes MS very strongly better than PSS. It remains possible that with larger data, per-subject effects may be resolvable, but the present data is not able, or intended, to do this.

5. Discussion

The maximum likelihood model is MS with a $\theta = 0.85$, i.e. most (85%) of trails in the test phase are responded by a random action. Importantly – the clearly significant victory of MS (200,000 more likely than its competitors) in the hypothesis contest is obtained despite this very high experimental noise level. This is possible because of the large number ($N = 552$) of trials conducted – a large number of very noisy trials give rise to a non-noisy, significant finding. The noise level was higher than expected during the experimental design, and possible causes include the absence of any percept at all due to the short (200ms) presentation time; bored subjects acting randomly to reach the end of the long experiment without regard for winning the small prize; and the occasional incoherent percepts reported during debriefing. These are all similar issues to the noise sources in (Fox & Stafford, 2012): unitary coherent perception is very difficult to isolate experimentally, because as is common in forced-choice psychophysics, too short a presentation results in no perception and random action, whilst too long a presentation allows time for other forms of cognition to occur. Psychophysics is famously difficult for this reason, and it remains a challenge to design more precise experiments to reduce the noise. High noise is expected because each presentation is short and barely perceivable, in order to avoid subjects seeing multiple interpretations in sequence, and so as in forced-choice blindsight experiments we expect each trial to be noisy but the aggregate of many trials to obtain significance. We stress that *despite the high noise, we have still shown MS to be 200,000 times more likely than the other hypotheses* because the large data set compensates for the per-trial noise.

As in any data modelling study, the possibility remains of newer, alternative hypothesis better fitting the data, than any of those tested so far. However, our hypothesis set includes the two most popular views of perception and action – MAP and MEU – which are not straw men. And they are clearly falsified in the study's domain by the success of MS. This does not mean that MS is 'true', just that it is a better model that was previously available. (This is standard Bayesian logic as detailed in (Jaynes & Bretthorst, 2003)). The contribution of this paper is in falsifying both MEU and MU for the face task, by proposing a new best hypothesis MS which beats them. Any future invention of new models giving even better fits than MS would be additional scientific progress – just as the present study refines and replaces the previous MU model of (Fox & Stafford, 2012) with MS. Future models might include *mixtures* of the current models. For example, new models might be completely non-probabilistic, or not involve assumptions about the existence of any mental representation of s. Or mixtures of strategies might occur within or between subjects. Mixtures made only of the hypotheses tested in the present study must still have MS as the strongest component because it has been shown to fit the data better than the others. Mixtures and other complex models will incur high Occam factor penalties for any extra parameters, and while the present study does not rule them out, it suggests that larger data sets would be needed to resolve competitions of mixture models against MS as the present data was not designed to test them.

The victory of MS in the present hypotheses contest is in agreement with the Necker cube study of (Fox & Stafford, 2012), fMRI work of (Fleming et al., 2010), and the unitary coherent hippocampal perception model of (Fox & Prescott, 2010). Together these findings suggest the existence of unitary coherent perception and action based upon it in more general domains. Given the current available hypothesis set, we have shown as in (Fleming et al., 2010), (Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012) that *utilities of available actions can affect what you see* in a new domain, and refined a novel, quantitative MS explanation which advances from the previous MU explanation. Utility biases unitary coherent perception, as is immediately clear from the difference between observed TALK and AVOID frequencies (recalling that calibration set the corresponding utility-neutral frequencies to be equal). The particular form of this bias now appears to be Maximum Salience – exactly as in Orwell's story '1984', whose hero not only reports but actually *sees* five fingers instead of four to avoid pain.

References

Bernardo, J., & Smith, A. (2000). Bayesian Theory. Wiley.

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, *10*(7), 287 - 291. (Special issue: Probabilistic models of cognition)

Cooper, G. F. (1990). The computational complexity of probabilistic inference. *Artif. Intell.*, *42*, 393–405.

Eisenfhr, F., Weber, M., & Langer, T. (2010). *Rational decision making.* Springer.

Felzenszwalb, P. F., & Huttenlocher, D. P. (2005). Pictorial structures for object recognition. *International Journal of Computer Vision*, *61*(1), 55-79.

Fleming, S. M., Whiteley, L., Hulme, O. J., Sahani, M., & Dolan, R. J. (2010). Effects of category-specific costs on neural systems for perceptual decision-making. *Journal of neurophysiology*, *103*(6), 3238.

Fox, C., & Prescott, T. (2010). Hippocampus as unitary coherent particle filter. In *International Joint Conference on Neural Networks (IJCANN)* (pp. 1–8).

Fox, C., & Stafford, T. (2012). Maximum utility unitary coherent perception vs. the Bayesian brain. *Proc. Ann. Conf. Cognitive Science Soc.*.

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010, June 23). Probabilistic models of cognition: exploring representations and inductive biases. *Trends in Cognitive Sciences*, *14*(8), 357-364.

Jaynes, E., & Bretthorst, G. (2003). *Probability theory: The logic of science.* Cambridge University Press.

Körding, K., & Wolpert, D. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.

Lengyel, M., & Dayan, P. (2008). Hippocampal contributions to control: the third way. In *NIPS 2008.* MIT Press.

Lepora, N. F., Sullivan, J. C., Mitchinson, B., Pearson, M., Gurney, K., & Prescott, T. J. (2012). Brain-inspired Bayesian perception for biomimetic robot touch. In *Robotics and automation (icra), 2012 ieee international conference on* (pp. 5111–5116).

Maier, A., Panagiotaropoulos, T. I., Tsuchiya, N., & Keliris, G. A. (2012). Binocular rivalry: a gateway to studying consciousness. *Frontiers in Human Neuroscience*, *6*(263).

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience*, *32*(7), 2335–2343.

No, A., Pessoa, L., & Thompson, E. (2000). Beyond the grand illusion: what change blindness really teaches us about vision. *Visual Cognition*, *7*, 93–106.

Spiegelhalter, D., Thomas, A., Best, N., & Gilks, W. (2008). *BUGS: Bayesian inference using Gibbs sampling.*

Young, S., Evermann, G., Gales, M., Hain, T., Kershaw, D., Liu, X., ... others (2006). The HTK book.