

Opinion

The Misestimation of Uncertainty in Affective Disorders

Erdem Pulcu¹ and Michael Browning ^{1,2,*}

Our knowledge about the state of the world is often incomplete, making it difficult to select the best course of action. One strategy that can be used to improve our ability to make decisions is to identify the causes of our ignorance (i.e., why an unexpected event might have occurred) and use estimates of the uncertainty induced by these causes to guide our learning. Here, we explain the logic behind this process and describe the evidence that human learners use estimates of uncertainty to sculpt their learning. Finally, we describe recent work suggesting that misestimation of uncertainty is involved in the development of anxiety and depression and describe how these ideas may be advanced.

Learning Why We Don't Know

When we observe our environment, we observe it incompletely and intermittently. Our senses report aspects of the particular location we find ourselves in, but many of the most important processes we want to know about cannot be directly perceived. For example, if I stroke my cat, will it scratch me or not (Figure 1A, Key Figure)? If we knew whether the cat was feeling angry or friendly, then we would know whether it was likely to scratch us, but the internal state of the cat is hidden and so, if we want to know it, it must be inferred from our previous experience of how it has behaved. To complicate matters further, the **hidden processes** (see [Glossary](#)) we estimate will often change over time, so even if the cat was friendly the last time we stroked it, it may now be angry. Therefore, the challenge is how to make reasonable inferences about hidden, changeable processes when faced with such fundamental uncertainty about these processes. Failure to meet this challenge reduces our ability to accurately gauge the opportunities and costs afforded by the environment and, ultimately, to select the best course of action.

One approach to mitigating the impact of uncertainty on our decisions is to identify its potential causes and use estimates of the extent of these causes to sculpt how we learn [1]. In other words, we need to learn why we do not know what we are interested in. In this opinion article, we focus on how estimates of two key types of uncertainty, **expected uncertainty** and **unexpected uncertainty** (see [Box 1](#) for detailed descriptions of these and other forms of uncertainty), may be used to improve learning, summarise the evidence that human learners utilise such estimates, and finally argue that difficulties with uncertainty estimation may be one mechanism underlying symptoms of psychiatric difficulties, such as anxiety and depression.

Ways of Not Knowing: Expected and Unexpected Uncertainty

'I thought that the cat was in a good mood, but it scratched me when I stroked it.' Surprising events occur when our understanding of the environment is less than perfect [2]. As introduced earlier, the potential causes of a **surprising** event may be useful to know. In the example of the cat that scratches you when you stroked it, it may be that the cat is in fact feeling friendly, but decided to scratch on this occasion, or that the cat is no longer feeling friendly and is now angry.

Highlights

Maintaining estimates of why unexpected events happen, of why we are uncertain, allows humans to adjust their inference and learn as efficiently as possible.

Humans adjust their learning in response to the levels of different types of uncertainty, including expected uncertainty, caused by probabilistic associations between observed events and what we are learning, and unexpected uncertainty, caused by changes in the thing we are learning.

People who experience anxiety are less sensitive to changes in levels of unexpected uncertainty.

Misestimation of the levels of uncertainty may lead to a tendency to be more influenced by negative events, which is believed to cause depression and anxiety.

Understanding how the estimation of uncertainty is related to symptoms of anxiety and depression promises to enhance our understanding of these illnesses and may suggest novel therapeutic interventions.

¹Department of Psychiatry, University of Oxford, Oxford, UK

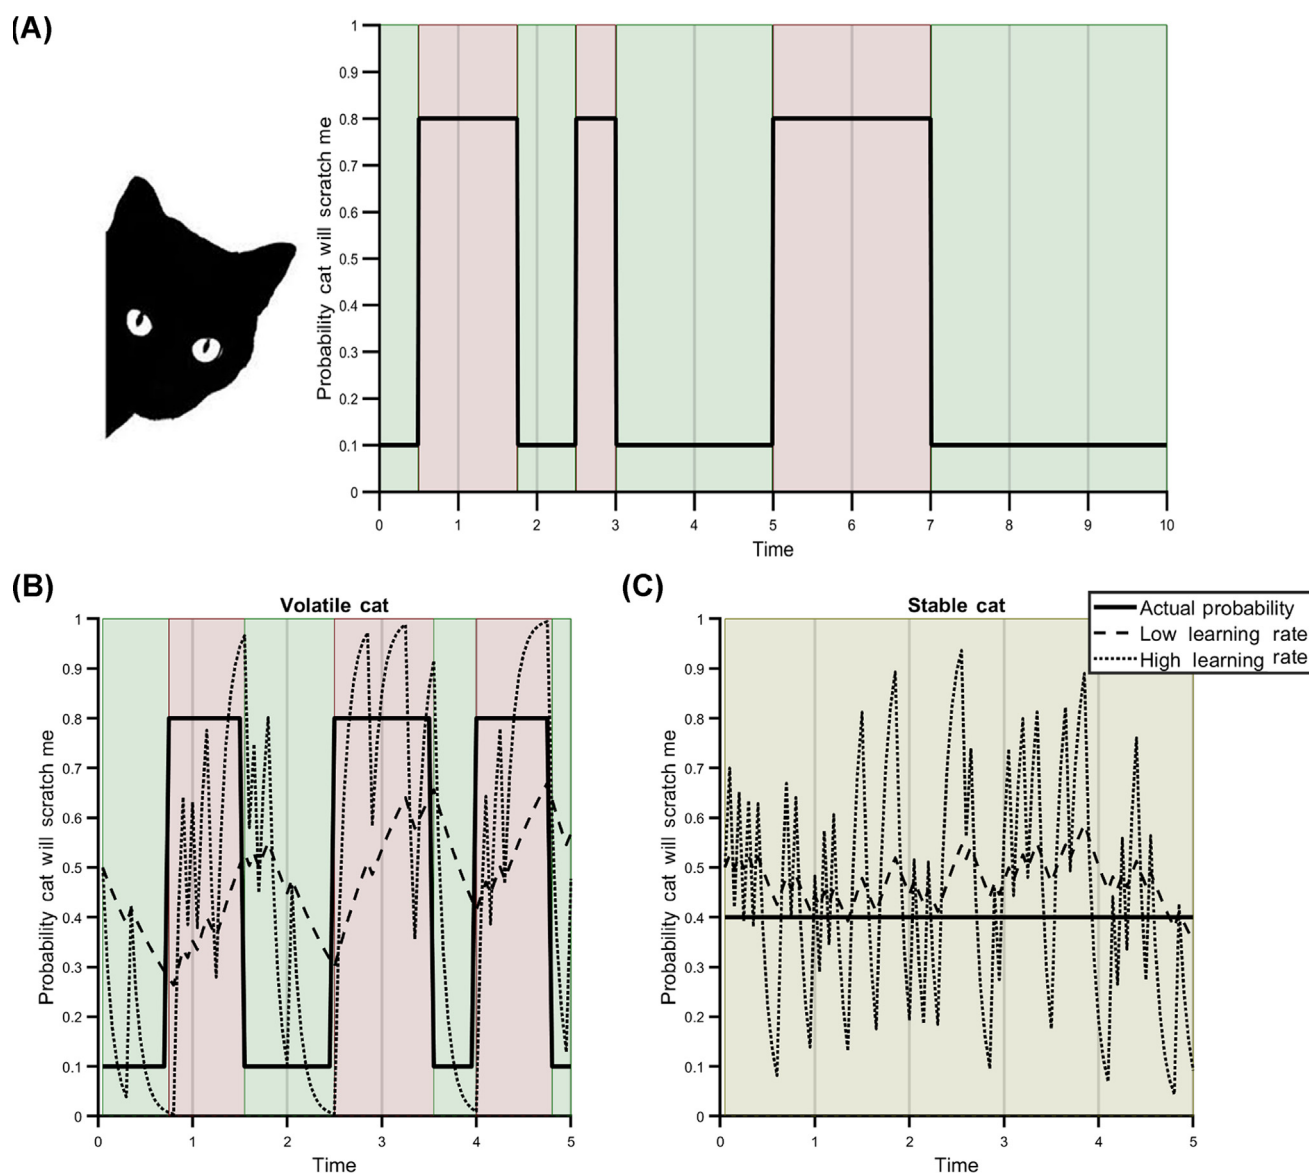
²Oxford Health Foundation Trust, Oxford, UK

*Correspondence:
Michael.browning@psych.ox.ac.uk
(M. Browning).



Key Figure

Uncertainty When Learning about Dynamic, Hidden Processes



Trends in Cognitive Sciences

Figure 1. (A) Your cat scratches you on 10% of the times you stroke it when it is in a good mood (green areas) and 80% of the time you stroke it when it is in a bad mood (red areas). You cannot observe the mood of the cat; all you can observe is whether it has scratched you when you previously stroked it. To learn what the mood of the cat is and, therefore, how likely it is to scratch you the next time you stroke it, you need to account for two important sources of uncertainty: (i) expected uncertainty: even if you know exactly what the mood of the cat is, you cannot be certain what its behaviour will be when you stroke it (e.g., even when it is in a good mood, it will scratch you 10% of the time); and (ii) unexpected uncertainty: the mood of the cat changes over time; thus, even if it was in a good mood the last time your stroked it, it might now be in a bad mood. If you stroke the cat and it does something surprising (e.g., scratches you when you thought it was in a good mood), how you should update your belief about its mood depends on what caused the surprising event. If you think it was caused by chance (expected uncertainty, it just happened to scratch you) then you should not change your belief about its mood much (use a low learning rate), whereas if it was caused by a change in the mood of the cat (unexpected uncertainty), you need to update your belief quickly (use a high

(Figure legend continued at the bottom of the next page.)

In the first of these cases, the surprising event is caused by a probabilistic relationship between the cat scratching and its mood. In other words, the cat may occasionally scratch even if it is feeling friendly and/or not scratch if it is feeling angry. This means that, even after experiencing whether the cat has scratched us or not, we cannot be sure of its mood. This form of uncertainty is called 'irreducible' or 'expected' uncertainty [1] and increases as the association between an observation and a hidden process becomes less deterministic. For example, expected uncertainty would be relatively higher if the cat scratches on 30%:75% of the times that it is in a good:bad mood and relatively lower if it scratches on 10%:88% of occasions. When considering how to respond to this source of surprising events, the higher the expected uncertainty, the less a specific observation tells us about the hidden process we are interested in, and the more observations we need to collect to confidently estimate the state of the hidden process (see Box 2 for a description of the effect of different sources of uncertainty on the information content of events). More concretely, having high expected uncertainty about the behaviour of the cat means that we should consider its average behaviour across multiple previous time points to be confident of whether it is in a good or bad mood now, whereas a low expected uncertainty would allow us to make a similar judgement based on fewer events.

A second reason why the cat may scratch is because its mood has changed and, even though it was previously feeling friendly, it is now angry. This form of uncertainty, called 'unexpected uncertainty', is produced when the hidden process we are estimating changes [1]. It is high if the mood of the cat changes frequently (is **volatile**) relative to if it is stable. Unexpected uncertainty requires the opposite response to expected uncertainty. Specifically, when learning about a volatile process, we cannot rely so much on previous observations because it is more likely that the hidden process will have changed since they occurred; rather, we have to base our current beliefs on more recent observations. Thus, if the behaviour of the cat has high unexpected uncertainty, we should base our belief about its mood mainly on how it responded the last few times we stroked it, whereas, if it has low unexpected uncertainty, we can incorporate a greater number of prior events in our estimate.

Reinforcement learning models [3] offer a relatively straightforward framework for describing how agents update their beliefs in response to experience and, thus, provide a useful tool for thinking about and measuring adaptation to uncertainty. One of the simplest models is the Rescorla–Wagner learning rule [4] (Equation 1), in which a belief at time t , $r_{(t)}$ is updated by a prediction error, the difference between an event which has just occurred $o_{(t)}$ and what the agent expected to happen (its current belief about the event), multiplied by a learning rate, α :

$$r_{(t+1)} = r_{(t)} + \alpha(o_{(t)} - r_{(t)}) \quad [1]$$

The learning rate in Equation 1 is simply some number between 0 and 1. The role of the learning rate becomes more obvious if we rearrange Equation 1 to collect the terms on the right (Equation 2):

$$r_{(t+1)} = (1-\alpha)r_{(t)} + \alpha(o_{(t)}) \quad [2]$$

Glossary

Expected uncertainty: uncertainty that arises because of a nondeterministic relationship between observable events and some hidden process. For example, the cat still scratches 10% of the time it is stroked, even when it is in a good mood (see Figure 1 in the main text).

Hidden process: any process that we cannot directly observe but that influences observations is a hidden process (sometimes called a 'partially observable' process). Examples include other people's mood, the weather, or how enjoyable a party will be.

Surprise: a measure of the absolute difference between what occurred and what was expected. In the reinforcement learning literature, surprise is often quantified as the absolute value of the prediction error. In the Bayesian literature, it is formalised as the negative logarithm of the conditional probability of the event, given the belief of the learner [2].

Unexpected uncertainty: uncertainty that arises because of a change in the thing that is being learned. For example, the mood of the cat changes (see Figure 1 in the main text).

Volatility: the changeability of a process, source of unexpected uncertainty. A related concept is a 'change point', which is when a process abruptly changes (see Box 1 in the main text).

learning rate). In other words, to learn about a hidden, changeable process such as the mood of the cat, you need to identify the potential causes of events you were not expecting and use estimates of the extent of these causes to guide how you learn. Misestimating uncertainty leads to suboptimal learning. (B) A volatile cat (mood changes frequently) and (C) a stable cat (mood remains constant). The dashed lines represent the belief of an agent learning about the cat who is using a low learning rate (i.e., assumes surprising events are caused by chance) and the dotted lines the belief of an agent using a high learning rate (assumes surprising events are caused by changes in the cat's mood). In (B), the low learning-rate learner does not adapt to changes in the mood of the cat fast enough and, thus, never accurately estimates its current mood. In (C), the high learning rate learner overreacts to chance events and so never settles on an accurate estimate of the mood of the cat.

Box 1. Forms of Uncertainty

In this paper, we focus specifically on expected and unexpected uncertainty. Here, we describe these and other forms of uncertainty [36,38,39].

Estimation Uncertainty

The first time we meet a new cat, we do not know whether it is likely to scratch us. This lack of knowledge is called 'estimation uncertainty' and, in some respects is similar to expected uncertainty. The important distinction is that estimation uncertainty occurs because of lack of experience; thus, we can reduce it by gathering more evidence (e.g., stroking the cat and seeing what happens), whereas, with expected uncertainty, the lack of knowledge is a fundamental aspect of the system that is not reduced by gathering more evidence.

Expected Uncertainty

Expected uncertainty arises when events occur probabilistically, rather than deterministically (Figure 1 in the main text) and is linked to 'risk'. Generally, the term 'risk' is used when a probability is explicit and expected uncertainty when it is estimated. Gathering more experience does not reduce expected uncertainty, which has resulted in it sometimes being called 'irreducible uncertainty'.

Unexpected Uncertainty

Unexpected uncertainty arises when the thing we are learning about changes (Figure 1 in the main text). Unexpected uncertainty may be conceptualised in several ways. In the example illustrated in Figure 1 in the main text, being scratched by the cat is ambiguous, so we can never be sure whether the mood of the cat has changed. In these cases, unexpected uncertainty is often described as 'volatility', roughly an estimate of the average amount that an association will change between observations. However, if we are estimating how much food our cat eats and, after eating 100–120 g a day for 3 months, it suddenly eats 300 g, we can clearly say that a change has just occurred (sometimes called a 'change point'). In fact, the two descriptions are complementary: the volatility of an association is the probability that a change point has occurred over a specific time period.

Higher Order Uncertainty

Uncertainty may change over time. This allows the possibility of higher order uncertainties, uncertainty about the current level of uncertainty. While the number of higher order uncertainties may theoretically be infinite, their calculation becomes progressively more demanding, and their impact on estimates of the environment more subtle, meaning there is likely to be a limit beyond which they are not accounted for in human learning. This highlights a more general question when investigating how humans learn: what are the hidden processes that are (or should be) estimated?

As can be seen, the Rescorla–Wagner learning rule defines the updated belief as a weighted mean of new information from the recently observed outcome, $o_{(t)}$, and old information in the form of the previous belief, $r_{(t)}$, with the learning rate acting as the weight. When the learning rate is high, the learner places more weight on the recent observation; when it is low, the learner places more weight on prior belief. As described earlier, if an agent believes that surprising events are being generated by noise in the relationship between observations and the hidden process (expected uncertainty), then it should be more influenced by prior events, and the agent should use a lower learning rate (Figure 1C). By contrast, if an agent thinks the surprising events are caused by changes in the hidden process being estimated (unexpected uncertainty) then it should be more influenced by recent events and use a higher learning rate (Figure 1B). Here, it is important to note that the Rescorla–Wagner learning rule itself does not estimate uncertainty: it assumes that learning rate does not change. It is presented here as a simple framework to illustrate how learning rates should adapt to changes in expected and unexpected uncertainty. A range of more complex models, which are able to estimate uncertainty, are described in Box 3. These more complex models use their estimates of uncertainty to adapt their learning rate in the manner described earlier.

In summary, surprising events can arise from a variety of causes. Different causes of surprising events may motivate diametrically opposing inferential responses, with an increased learning

Box 2. The Effect of Uncertainty on the Information Content of Events

As described in the main text, different forms of uncertainty influence how useful recent relative to more distant previous events are when learning about a hidden process. Expected uncertainty, arising from a probabilistic relationship between the hidden process being learned and observations, reduces the amount of information individual events provide, requiring learners to rely on more distant previous events during learning (i.e., use a lower learning rate, see Equations 1 and 2 in main text). By contrast, unexpected uncertainty erodes the information content of past observations by making it more likely that the underlying generative process has changed since they were observed (i.e., prompting the use of a higher learning rate). Here, the term 'information content' can be understood as the amount of information observing an event provides to a learner. The information content is challenging to quantify because it depends on the nature of the event as well as the learner's prior belief(s) about the state and structure of the world. One approach that has been used to measure the information content of events is to construct a Bayesian Ideal Observer model of a task and use the Kullback–Leibler divergence (KLDiv) of the belief of the model from before to after a particular event as a measure of the information content of the event [5].

Here, the Bayesian Observer (see [7] for a description of an ideal observer model) maintains its belief about the state of relevant aspects of the task as dimensions of a multivariate probability distribution, with updates applied to this distribution that reflect the assumptions of the model about the structure of the task (e.g., that the process it is learning about may change over time; Figure 1). The KLDiv is a distance measure, in this case of the change in the ideal observer's belief from before to after the event, and can be conceptualised as the degree to which the event changes the learner's belief about the generative process that it uses to predict future outcomes. An example of this approach is shown in Figure 1. Data from this model allow quantification of the effects of unexpected and expected uncertainty ('volatility' and 'noise' in Figure 1) on the information content ('KLDiv' in Figure 1) of events: KLDiv is significantly higher when volatility is higher [$F(1,831)=4.3$, $P=0.038$] and when noise is lower [$F(1,831)=142$, $P<0.001$]. In other words, the learning rate adaptation in response to uncertainty described in the main text reflects the relative information content of events: a higher learning rate is used when events have relatively higher information content.

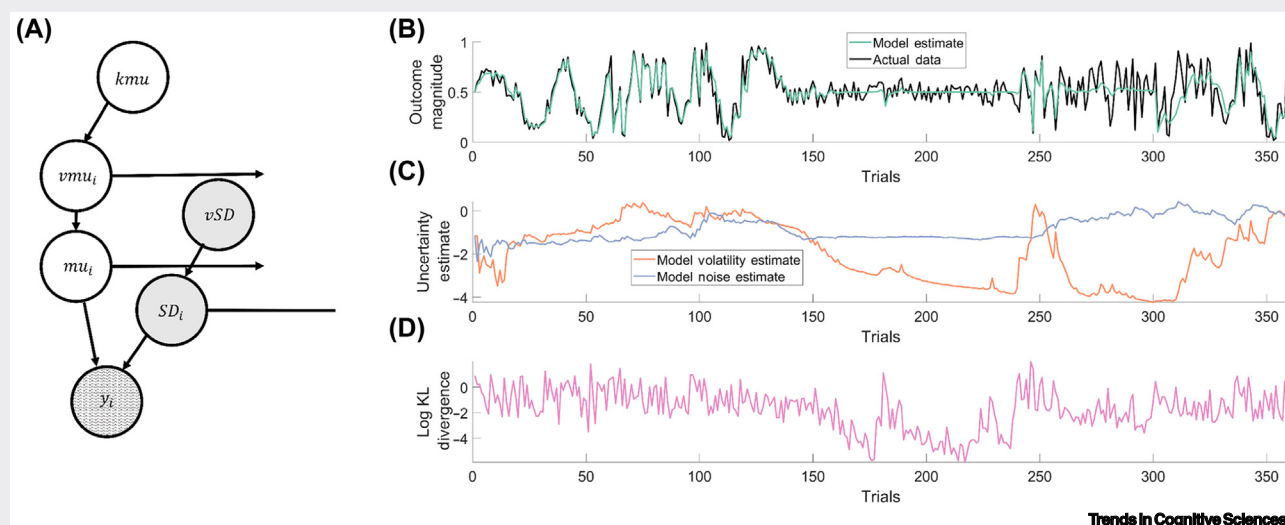


Figure 1. An Ideal Observer Model Designed to Estimate Both Expected and Unexpected Uncertainty. (A) A graphical model of a Bayesian Ideal Observer Model, which estimates the magnitude of an outcome (y_i) from previous observations. The model believes that the outcomes are generated from a Gaussian distribution with mean mu_i and a standard deviation SD_i , which produces expected uncertainty. The mean of the process is volatile with a change rate vmu_i , producing unexpected uncertainty. km_u and vSD are higher level change parameters that allow the model to conceptualise changes in vmu_i and SD_i . (B) Data (black line) and model-estimated mu_i (green line) when learning about the data. Note that the data include periods of high (trials 1–120 and 300–360) and low (trials 121–299) volatility as well as high (trials 240–360) and low (trials 1–239) noise. From trial 1–120, the model thinks (correctly) that volatility is high and so changes its belief about the mean to match its observations. Between trials 260 and 300, the model thinks volatility is low and noise is high, so it (correctly) ignores variation in the observations when estimating the mean of the process. Between trials 240 and 260, the model mistakes noise for volatility and so rapidly changes its belief in response to random fluctuations in the outcome (i.e., uses a high learning rate when it should be low). (C) Model-derived estimates of unexpected uncertainty (volatility, vmu_i) and expected uncertainty (noise, SD_i). (D) Model-derived Kullback–Leibler divergence (KLDiv) which provides a measure of the information content of each trial.

rate being appropriate for higher unexpected uncertainty and a decreased learning rate for higher expected uncertainty. An alternative way to frame this problem is to consider whether the information contained in a surprising event is useful at predicting the future (i.e., like unexpected uncertainty) or not (i.e., like expected uncertainty) [5,6]. This illustrates that one way of mitigating the impact of uncertainty on decision-making is to identify the potential causes of the uncertainty,

Box 3. Models of Uncertainty Estimation

In this paper, we have framed the logic behind using estimates of uncertainty to guide learning. We have also described studies in which the simple Rescorla–Wagner learning rule was used to obtain separate estimates of learning rate when uncertainty was high versus low [7,8,10,11]. In these studies, the Rescorla–Wagner rule was used as a method of measuring learning rate; it does not describe how learners might calculate whether learning rate should be changed (the model is simply fit separately to blocks in which uncertainty is high or low). Several more complex models, which are able to adapt their effective learning rate in response to uncertainty, have been described, which we briefly summarise here.

A relatively simple adaptation of the Rescorla–Wagner learning rule was suggested by Pearce and Hall [40] and developed by others [41]. The broad class of ‘Pearce–Hall type’ models are able to adapt to changes in unexpected uncertainty by estimating the surprise associated with a stimulus (i.e., learning the absolute magnitude of the prediction error) and using this estimate to adjust the main learning rate of the model. More elaborate than this are models in which various sources of uncertainty are estimated separately before impacting learning [9,42]. Finally, Bayesian learners, which invert generative models of the environment, including estimates of various sources of uncertainty, have been described. The most prominent of these is the Kalman Filter [43], which can efficiently estimate and adapt to changes in unexpected uncertainty. More recent Bayesian ‘filters’ have also been described which provide estimates of higher order and/or expected uncertainty [7,38] (see also Box 2 in the main text). These different models solve similar problems in different ways, with the complexity of the model increasing from the Pearce–Hall to the Bayesian examples. To date, it is not clear which model provides a more veridical description of how humans solve the problem of estimation of uncertainty.

estimate the extent of each cause, and then to use these estimates to sculpt how we learn. In the next section, we consider the evidence suggesting that human learners maintain and use such estimates.

The Impact of Estimates of Uncertainty on Human Learning

As explained in the previous section, learning is more accurate if a higher learning rate is used for volatile than stable processes (i.e., when unexpected uncertainty is high relative to when it is low). Several studies using a range of reward and/or punishment-based learning tasks [7–13] have examined whether humans increase their learning rate when the volatility of the process being learned increases. The consistent finding across studies is that, as predicted by the normative account of learning, participants use a higher learning rate in volatile relative to stable contexts (i.e., when unexpected uncertainty is high compared with low). This effect is apparent in tasks using both rewarding [7–9,12,13] and punishing [10,11] outcomes.

Investigation of physiological markers of this volatility adaptation process have reported activity in the dorsomedial prefrontal cortex (dmPFC) that correlates with estimated volatility [7,8,13–15], with a consistent effect found in single-neuron recordings from primates [16]. A second line of enquiry, arising from an early synthesis of animal work [1], suggests that phasic activity of the central norepinepheric (NE) system also contains an estimate of volatility. This proposition is consistent with theories [17,18] on the broader role of NE, which argue that it increases the gain of sensory representations and, thus, increases their impact on behaviour (i.e., analogous to an increased learning rate). Phasic activity of the central NE system is correlated with pupil dilation in primates [19], suggesting that it would be possible to estimate activity of this system using pupillometry. Consistent with the proposed role of central NE activity in representing estimated volatility, several previous studies have reported greater pupillary dilation following outcomes when estimated volatility was high relative to when it was low [9–12,14].

In contrast to unexpected uncertainty, increased expected uncertainty should prompt reduced learning rates. Most studies that have explicitly manipulated expected uncertainty have used tasks in which participants have to learn the magnitude, rather than the probability, of an outcome. In such tasks, outcomes are generated using a random process, with the standard deviation (SD) of this generative process controlling the expected uncertainty. An advantage of this type of task is that expected uncertainty (the SD of the process) and unexpected uncertainty

(shifts in the mean of the process) can be independently manipulated. The predicted reduction in learning rate as expected uncertainty increases has been reported using two reward-based versions of this magnitude learning task [9,12,20–22]. Subsequent work showed that the prediction errors encoded in the activity of midbrain dopaminergic areas (substantia nigra and/or ventral tegmental area) displayed a similar scaling as that seen for the learning rate [21]. Given that learning rate multiplies prediction error to produce the belief update (Equation 1), the authors interpreted this result as evidence of the mechanism by which estimates of expected uncertainty modify belief updating.

In summary, there is robust evidence that human learners adapt to changes in unexpected uncertainty in the manner predicted by normative accounts of learning and initial evidence of a similar process for expected uncertainty. These observations suggest that humans maintain estimates of the uncertainty of the associations they are learning and use these estimates to tune their learning.

The Impact of the Misestimation of Uncertainty on Human Learning

The above-described work indicates that, overall, humans are able to adapt to changes in the unexpected and expected uncertainty of rewarding outcomes, at least in some situations. When investigating the role of uncertainty estimation in affective disorders, as we do later, it is useful to consider what the impact on behaviour might be if an individual was generally less able to estimate the degree or nature of uncertainty during learning.

The effects of misestimating the level of uncertainty can be derived from the normative account of learning described earlier. Overestimating expected uncertainty and/or underestimating unexpected uncertainty would lead an individual to be relatively less sensitive to recent events and slower to adapt to important changes, particularly in volatile, less noisy environments (Figure 1B). By contrast, underestimating expected uncertainty and/or overestimating unexpected uncertainty would lead to an increased influence of recent events and unstable beliefs, particularly in stable, noisy environments (Figure 1C). Therefore, a general consequence of uncertainty misestimation is a reduced ability to accurately track the state of dynamic environments and, thus, to select the actions most likely to obtain positive outcomes or avoid negative outcomes. As a result, the misestimation of uncertainty may provide a mechanistic link between the internal cognitive facets of depression and anxiety and the external risk factors for the disorders, such as exposure to adversity. More targeted consequences of uncertainty misestimation would be expected if the uncertainty of specific associations were misjudged. For example, as discussed later, a specific tendency to overestimate the unexpected uncertainty of negative relative to positive outcomes would lead an individual to be more influenced by recent negative than positive events, a cognitive profile that is characteristic of depression and anxiety [23,24].

As well as misestimating the degree of uncertainty, its nature may also be misjudged. This may occur, for example, by mistaking the random variation that produces expected uncertainty for the meaningful change to the environment that produces unexpected uncertainty (e.g., see trials 245–260 of Figure 1 in Box 2). The opposing responses required by expected and unexpected uncertainty on learning rates mean that this form of misattribution would lead to particularly inefficient estimation of the state of the environment and exaggerate the effects of uncertainty misestimation described earlier. It is notable that the research examining the effects of uncertainty manipulations, reviewed earlier, has either manipulated one form of uncertainty [7,10], has explicitly instructed participants about which form of uncertainty has changed [5,20,21], or has made the effects of different forms of uncertainty unambiguously different (e.g., by making changes caused by unexpected uncertainty much bigger than changes caused by expected uncertainty

[9,12,22]). As a result, the degree to which humans can accurately attribute uncertainty to a specific source when it is not unambiguously specified has yet to be investigated (see Outstanding Questions).

Abnormal Learning in Depression and Anxiety

Adverse experience is a risk factor for both depression and anxiety [25], suggesting that learning processes are involved in the aetiology of these illnesses. Several recent reviews [24,26–29] have described in detail the evidence that altered learning processes are associated with these illnesses. Here, we briefly summarise the pertinent aspects of this literature before considering the possible role of the misestimation of uncertainty in anxiety and depression.

The case for abnormal learning in anxiety is supported by a large case-control literature examining fear conditioning and aversive learning in anxious relative to control groups, with evidence that anxious participants are slower to extinguish fear associations and show greater generalisation of such associations compared with non-anxious controls [26,29,30]. A complementary mechanistic literature has suggested that anxious individuals show increased limbic (including amygdala) and reduced inhibitory frontal activity in response to aversive stimuli [28,31,32]. The learning literature in depression is somewhat less developed and, while there exists a degree of inconsistency in the findings, current evidence suggests that depressed patients are generally less behaviourally sensitive to rewarding outcomes than are control participants [24,27]. From a mechanistic perspective, a similar network of brain regions to those described for anxiety have been reported in depressed cohorts [33], with more focussed work looking specifically at anhedonia also describing aberrant activity in the substantia nigra and ventral tegmental area [34].

A common feature of the clinical literatures on both anxiety and depression has been that they have generally reported ‘first-order’ measures of learning and decision-making that assess gross effects, such as how quickly patients learn associations or their behavioural sensitivity to a specific valence of outcome. To date, little work has assessed questions related to the uncertainty of outcomes that require ‘second-order’ measures of learning, such as whether patients are able to adapt their learning rate to changes in unexpected and/or expected uncertainty. As described in earlier sections, tuning learning to the estimated levels of uncertainty is important when learning about hidden, dynamic processes. This suggests that some aspects of the aberrant learning associated with depression and anxiety are related to the misestimation of uncertainty.

In the next section, we review early work that provides evidence for this proposal and suggests how such misestimation influences cognitive processing generally.

The Misestimation of Uncertainty in Depression and Anxiety

One approach to assessing whether symptoms of anxiety and depression may be associated with misestimation of the level of uncertainty is to measure how individuals with high versus low symptom scores adapt their learning in response to changes in the levels of uncertainty. A single study [10] has examined the degree to which trait anxiety is associated with learning rate adaptation to the changing volatility of an aversive outcome. In this study, a group of nonclinical participants, selected to have a range of trait anxiety scores, completed a learning task that comprised two blocks of trials during which participants had to learn the association between shape stimuli and electric shocks. In one block of the task, the stimulus–outcome associations were stable, whereas they were volatile in the other block. The authors estimated the learning rates used by participants in the two blocks and showed that, across all participants, a higher learning rate was used for volatile than for stable blocks. The key result of the study was that

this difference in learning rate decreased as trait anxiety increased. In other words, individuals with higher anxiety adjusted their learning rate less in response to changes in unexpected uncertainty. Consistent with the putative role of central NE in representing estimated unexpected uncertainty [1,17,18], pupil dilation was found to be greater when outcomes were presented in the volatile than in the stable block, with the size of this volatility-related pupillary effect also being significantly smaller in anxious participants. While this study provides initial evidence that symptoms of anxiety are associated with a relative insensitivity to changes in unexpected uncertainty, several crucial issues regarding the specificity of the findings and the causal (or otherwise) relationship between uncertainty estimation and symptoms remain (see Outstanding Questions).

As described earlier, another route by which the misestimation of uncertainty may impact symptoms of anxiety and depression is by skewing the processing of affective stimuli in a manner that favours negative over positive events. The preferential processing of negative events, known as 'negative cognitive bias', is believed to be causally linked to both anxiety and depression [23,24,35]. As described in the introductory sections of this opinion article, the optimal response to changes in estimated uncertainty is to adjust the degree to which beliefs are influenced by current relative to previous events (i.e., to use a higher or lower learning rate), to reflect the relative information content of the events (Box 2). Patients with anxiety and/or depression are more influenced by negative than by positive events [21,31,32], suggesting that they misestimate the uncertainty associated with the two class of events such that negative events are judged to be more informative. This mechanism would require individuals to maintain independent estimates of the uncertainties associated with positive and negative events. An initial study [11] tested whether such separable estimates exist by using a task in which choices led to independent positive and negative outcomes (wins and losses of money). In the task, the volatility and, therefore, the unexpected uncertainties, of the two outcomes were manipulated such that, in one block, wins were more volatile than losses, whereas, in another, the losses were more volatile. This manipulation serves to make the relatively more volatile outcome more informative (Box 2). The degree to which participant belief was updated in response to the separate outcomes was measured by fitting separate learning rates for win and loss outcomes. A clear effect was found such that participants independently adjusted their learning rates for the positive and negative outcomes with a higher rate being used for whichever outcome was more volatile and pupil dilation during outcome receipt reflecting this learning rate effect. These results provide evidence that human learners maintain separable estimates of the unexpected uncertainties of positive and negative outcomes and use these estimates to modify the degree to which they update their beliefs in response to a particular valence of outcome. The existence of independent estimates of the uncertainties associated with positive and negative events suggests that the estimates may motivate the negative cognitive biases underlying anxiety and depression.

In summary, a few studies have begun to examine the potential role of misestimation of the level of uncertainty in anxiety and depression. The results to date are promissory rather than conclusive, and have focussed on unexpected rather than on expected uncertainty, although they suggest that considering questions of uncertainty estimation would allow us to identify novel cognitive mechanisms of the illnesses.

While the focus of this paper has been on the affective disorders, misestimation of uncertainty can produce profound changes in how an individual interacts with their environment (see section on the impact of the misestimation of uncertainty). This suggests that difficulties in estimating uncertainty may underlie other psychiatric phenomena, such as response to stress [36] or autism [37]. Thus, it will be essential to determine the specificity of these findings by assessing uncertainty estimation processes in populations of participants with a range of presentations. This and other issues are summarised in the Outstanding Questions.

Concluding Remarks

Humans are sophisticated learners, which is fortunate given that the complexity of the world we need to learn about, and our inability to directly observe many important aspects of it, make accurate learning difficult. There is strong preclinical evidence that one strategy used to improve inference involves identifying the causes of surprising events, estimating the magnitude of these causes and using these estimates to sculpt how much we learn from the events. Anxiety and depression are associated with difficulties in learning, particularly about valenced events, where an enhanced impact of negative over positive outcomes is seen. Misestimation of the level, and potentially the source, of uncertainty may underlie some of these difficulties. Recent advances in the computational analysis of behaviour and physiology mean that understanding how this misestimation arises is now within our reach and that testing the impact of interventions that modify it should be a translational priority.

Acknowledgements

The authors would like to thank Phil Cowen for interesting discussions on the topics presented in this manuscript. This study was supported by a MRC Clinician Scientist Fellowship awarded to M.B. (MR/N008103/1) and by the NIHR Oxford Health Biomedical Research Centre.

Disclaimer Statement

M.B. has received travel expenses from Lundbeck for attending conferences, and M.B. and E.P. act as consultants for J&J.

References

1. Yu, A.J. and Dayan, P. (2005) Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692
2. MacKay, D.J.C. (2003) *Information Theory Inference and Learning Algorithms*, Cambridge University Press
3. Sutton, R. and Barto, A.G. (1998) *Reinforcement Learning*, MIT Press
4. Rescorla, R.A. and Wagner, A.R. (1972) A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In *Classical Conditioning II: Current Research and Theory* (Black, A.H. and Prokasy, W.F., eds), pp. 64–99, Appleton-Century-Crofts
5. O'Reilly, J.X. et al. (2013) Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proc. Natl. Acad. Sci. U. S. A.* 110, E3660–E3669
6. Nour, M.M. et al. (2018) Dopaminergic basis for signaling belief updates, but not surprise, and the link to paranoia. *Proc. Natl. Acad. Sci. U. S. A.* 115, E10167–E10176
7. Behrens, T.E.J. et al. (2007) Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221
8. Behrens, T.E.J. et al. (2008) Associative learning of social value. *Nature* 456, 245–249
9. Nassar, M.R. et al. (2012) Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* 15, 1040–1046
10. Browning, M. et al. (2015) Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nat. Neurosci.* 18, 590–596
11. Pulcu, E. and Browning, M. (2017) Affective bias as a rational response to the statistics of rewards and punishments. *Elife* 6, e27879
12. Krishnamurthy, K. et al. (2017) Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nat. Hum. Behav.* 1, 0107
13. Nassar, M.R. et al. (2019) Dissociable forms of uncertainty-driven representational change across the human brain. *J. Neurosci.* 39, 1688–1698
14. Muller, T.H. et al. (2019) Control of entropy in neural models of environmental state. *Elife* 8, e39404
15. McGuire, J.T. et al. (2014) Functionally dissociable influences on learning rate in a dynamic environment. *Neuron* 84, 870–881
16. Massi, B. et al. (2018) Volatility facilitates value updating in the prefrontal cortex. *Neuron* 99, 598–608
17. Aston-Jones, G. and Cohen, J.D. (2005) An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* 28, 403–450
18. Sales, A.C. et al. (2019) Locus Coeruleus tracking of prediction errors optimises cognitive flexibility: an active inference model. *PLoS Comput. Biol.* 15, e1006267
19. Joshi, S. et al. (2016) Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron* 89, 221–234
20. Diederer, K.M.J. and Schultz, W. (2015) Scaling prediction errors to reward variability benefits error-driven learning in humans. *J. Neurophysiol.* 114, 1628–1640
21. Diederer, K.M.J. et al. (2016) Adaptive prediction error coding in the human midbrain and striatum facilitates behavioral adaptation and learning efficiency. *Neuron* 90, 1127–1138
22. Nassar, M.R. et al. (2016) Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nat. Commun.* 7, 11609
23. Mathews, A. and MacLeod, C. (2005) Cognitive vulnerability to emotional disorders. *Annu. Rev. Clin. Psychol.* 1, 167–195
24. Eshel, N. and Roiser, J.P. (2010) Reward and punishment processing in depression. *Biol. Psychiatry* 68, 118–124
25. Heim, C. and Nemeroff, C.B. (2001) The role of childhood trauma in the neurobiology of mood and anxiety disorders: pre-clinical and clinical studies. *Biol. Psychiatry* 49, 1023–1039
26. Duits, P. et al. (2015) Updated meta-analysis of classical fear conditioning in the anxiety disorders. *Depress. Anxiety* 32, 239–253
27. Robinson, O.J. and Chase, H.W. (2017) Learning and choice in mood disorders: searching for the computational parameters of anhedonia. *Comput. Psychiatr.* 1, 208–233
28. Craske, M.G. et al. (2017) Anxiety disorders. *Nat. Rev. Dis. Primers* 3, 17024
29. Pittig, A. et al. (2018) The role of associative fear and avoidance learning in anxiety disorders: gaps and directions for future research. *Neurosci. Biobehav. Rev.* 88, 117–140
30. Lissek, S. et al. (2005) Classical fear conditioning in the anxiety disorders: a meta-analysis. *Behav. Res. Ther.* 43, 1391–1424

Outstanding Questions

How accurately can human learners distinguish between expected and unexpected uncertainty? The published literature has used tasks in which expected and unexpected uncertainty are unambiguously differentiated, leaving open the question of how humans adjust their learning when this is not the case.

How is the misestimation of uncertainty linked to prior experience? The process by which levels of uncertainty are estimated may be influenced by prior experiences, suggesting a potential mechanism by which adverse experiences could impact risk for emotional disorders. It would be interesting to test, for example, whether experience of adversity was associated with overestimates of the unexpected uncertainty of negative events.

How is uncertainty represented in the brain? While animal and human work has identified a range of neural areas in which activity correlates with estimates of uncertainty, it is not clear how the separable estimate of uncertainty related to positive and negative outcomes are represented.

Is insensitivity to changes in volatility specific to anxiety disorders? The finding that symptoms of anxiety correlate with insensitivity to changes in volatility does not provide strong evidence that this relationship is specific to anxiety. Misestimation of unexpected uncertainty has also been reported in individuals with autism and, thus, may be characteristic of a range of psychological difficulties rather than anxiety specifically.

What is the causal relationship between misestimation of uncertainty and affective disorders? The association between a cognitive process, such as the (mis)estimation of uncertainty, and symptoms is interesting if it tells us something about the aetiology of, or a potential treatment for, the symptoms. Testing this requires experimental studies in which the cognitive process is manipulated, for example by using targeted pharmacological or cognitive interventions.

31. Indovina, I. *et al.* (2011) Fear-conditioning mechanisms associated with trait vulnerability to anxiety in humans. *Neuron* 69, 563–571
32. LeDoux, J.E. and Pine, D.S. (2016) Using neuroscience to help understand fear and anxiety: a two-system framework. *Am. J. Psychiatry* 173, 1083–1093
33. Kaltenboeck, A. and Harmer, C.J. (2018) The neuroscience of depressive disorders: a brief review of the past and some considerations about the future. *Brain Neurosci. Adv.* Published online October 8, 2018. <https://doi.org/10.1177/2398212818799269>
34. Husain, M. and Roiser, J.P. (2018) Neuroscience of apathy and anhedonia: a transdiagnostic approach. *Nat. Rev. Neurosci.* 19, 470–484
35. Browning, M. *et al.* (2010) The modification of attentional bias to emotional information: a review of the techniques, mechanisms, and relevance to emotional disorders. *Cogn. Affect. Behav. Neurosci.* 10, 8–20
36. de Berker, A.O. *et al.* (2016) Computations of uncertainty mediate acute stress responses in humans. *Nat. Commun.* 7, 10996
37. Lawson, R.P. *et al.* (2017) Adults with autism overestimate the volatility of the sensory environment. *Nat. Neurosci.* 20, 1293–1299
38. Mathys, C.D. *et al.* (2014) Uncertainty in perception and the Hierarchical Gaussian Filter. *Front. Hum. Neurosci.* 8, 825
39. Marshall, L. *et al.* (2016) Pharmacological fingerprints of contextual uncertainty. *PLoS Biol.* 14, e1002575
40. Pearce, J.M. and Hall, G. (1980) A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* 87, 532–552
41. Li, J. *et al.* (2011) Differential roles of human striatum and amygdala in associative learning. *Nat. Neurosci.* 14, 1250–1252
42. Nassar, M.R. *et al.* (2010) An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* 30, 12366–12378
43. Kalman, R.E. (1960) A new approach to linear filtering and prediction problem. *J. Basic Eng.* 82, 35–45