Emotional Valence and the Free-Energy Principle

Mateus Joffily[1,2]*, Giorgio Coricelli[1,2,3]

1 Center for Mind/Brain Sciences, University of Trento, Trento, Italy, **2** Groupe d'Analyse et de Theorie Economique, Centre National de la Recherche Scientifique, Lyon, France, **3** Department of Economics, University of Southern California, Los Angeles, California, United States of America

Abstract

The free-energy principle has recently been proposed as a unified Bayesian account of perception, learning and action. Despite the inextricable link between emotion and cognition, emotion has not yet been formulated under this framework. A core concept that permeates many perspectives on emotion is valence, which broadly refers to the positive and negative character of emotion or some of its aspects. In the present paper, we propose a definition of emotional valence in terms of the negative rate of change of free-energy over time. If the second time-derivative of free-energy is taken into account, the dynamics of basic forms of emotion such as happiness, unhappiness, hope, fear, disappointment and relief can be explained. In this formulation, an important function of emotional valence turns out to regulate the learning rate of the causes of sensory inputs. When sensations increasingly violate the agent's expectations, valence is negative and increases the learning rate. Conversely, when sensations increasingly fulfil the agent's expectations, valence is positive and decreases the learning rate. This dynamic interaction between emotional valence and learning rate highlights the crucial role played by emotions in biological agents' adaptation to unexpected changes in their world.

Introduction

Free-energy is an information theoretic quantity that upper bounds the surprise on sampling some data, given a generative model of how data were caused. The free-energy principle assumes that biological agents encode a probabilistic model of the causes of their sensations, and postulates that any adaptive agent that resists a tendency to disorder must minimize its free-energy [1,2]. Under simplifying (Gaussian) assumptions, free-energy minimization can be understood as the minimization of the prediction error between the actual and predicted sensory inputs. In order to comply with the free-energy principle, adaptive agents have two tactics at their disposal: (1) adjusting their internal states to generate more accurate predictions and (2) acting on the environment to sample sensations that fulfil their predictions. The principle is based upon the realization that perceptual inference and learning [3,4] and active inference [5,6] rest on the same Bayesian scheme. Perceptual inference refers to inferring the states of the world causing sensory inputs. Perceptual learning relates to learning the relationship between inputs and causes. Active inference corresponds to acting on the world to fulfil prior expectations about sensory inputs. The computational implementation of the free-energy principle has been shown to replicate in many aspects neural mechanisms and the cortical organization of the brain [4,7].

Crucially, when inferring and learning the causes of their sensations in a changing world, adaptive agents need to deal with different forms of uncertainty, namely estimation uncertainty [8], volatility [9,10] and unexpected uncertainty [11,12]. Estimation uncertainty refers to the known estimation variance of states of the world causing sensory inputs and can be reduced through

learning. Volatility refers to slow and continuous changes in states of the world, and has usually been modelled by making estimation uncertainty follow some latent stochastic process [13]. Finally, unexpected uncertainty arises from surprising sensory inputs caused by discrete and fast changes in states of the world, and calls for forgetting the past and restarting learning from new sensory data. Dealing with different forms of uncertainty is fundamental to Bayesian models of learning in a non-stationary environment [14]. In fact, a critical challenge faced by many recent Bayesian schemes of human learning is how to dynamically update beliefs about states of the world in order to optimize predictions in a changing environment [9–12,15–18].

Despite the major role attributed to emotions in influencing the content and the strength of the agent's beliefs and the resistance of these beliefs to modification [19], emotion has not been considered in - much less integrated into - these computational models. The concept of emotional valence, or simply valence, has been used among emotion researchers to refer to the positive and negative character of emotion or some of its aspects, including elicitors (events, objects), subjective experiences (feeling, affect) and expressive behaviours (facial, bodily, verbal) [20]. The valence of feelings has been argued to be a pivotal criterion for demarcating emotion [21] and a core dimension of the subjective experience of moods and emotions [22]. Traditionally, mood has been defined, in contrast to emotion, as an affective state that lacks a clear referent, changes slowly and lasts for an extended period of time [23].

In the present paper, we propose a mathematical definition of emotional valence in terms of the negative rate of change of free-energy over time. As we shall see later, this formalism entails the dynamic attribution of emotional valence to every state of the

Author Summary

Emotion plays a crucial role in the adaptation of humans and other animals to changes in their world. Nevertheless, emotion has been neglected in Bayesian models of learning in non-stationary environments. The free-energy principle has recently been proposed as a unified account of learning, perception and action in biological agents. In this paper, we propose a formal definition of emotional valence (i.e., the positive and negative character of emotion) in terms of the rate of change of free-energy or, under some simplifying assumptions, of prediction error over time. This formalization leads to a straightforward and simple meta-learning scheme that accounts for the complex and reciprocal interaction between cognition and emotion. We instantiate this scheme with an emotional agent who is able to dynamically assign emotional valence to every new state of the world that is visited and to experience basic forms of emotion. Crucially, our hypothetical agent uses emotional valence to dynamically adapt to unexpected changes in the world. The proposed scheme is very general in the sense that it is not tied to any particular generative model of sensory inputs.

world that an adaptive agent might visit and prescribes the dynamics of basic forms of emotion such as *happiness*, *unhappiness*, *hope*, *fear*, *disappointment* and *relief*.

We will first introduce the free-energy principle and present our computational model of emotional valence. We then demonstrate this scheme by simulating and comparing two artificial agents. One agent explicitly optimises posterior beliefs about volatility and does not use its internally generated emotional valence signal. The other does not estimate volatility but instead implements the emotional regulation of estimation uncertainty. The contribution of this work is two-fold. First, we provide a phenomenological account of emotion in terms of changes in free energy - as a proxy for changes in the quality of how the world is being modelled during inference and learning. Second, emotion is coupled back to inference using a form of regularization or meta-learning. In other words, changes in the quality of inference are used to regularize the rate of evidence accumulation to provide adaptive learning rates. These learning rates correspond to expected uncertainty about inferences, under hierarchical models of the world.

Models

In this section, we introduce the free-energy principle as it was originally proposed by Friston, Kilner and Harrison [1], and then we propose a new mathematical definition of emotional valence and some basic forms of emotion in terms of free-energy. Next, we put forward a meta-learning rule by means of which emotional valence regulates estimation uncertainty, and outline the relationship between the dynamics of free-energy and some basic forms of emotion.

The free-energy principle

In statistical physics, variational free-energy minimization is a method for approximating a complex probability density $p(\vartheta)$ by a simpler ensemble density $q(\vartheta; \mu)$ that is parametrized by adjustable parameters μ [24]. In neuroscience, the free-energy principle assumes that biological agents encode the parameters μ of an arbitrary recognition (ensemble) density $q(\vartheta; \mu)$ of

environmental quantities ϑ that are the presumed causes of their sensations [2]. The recognition density $q(\vartheta; \mu)$ is an approximation to the true posterior density $p(\vartheta|\tilde{s},m)$ of ϑ, given the sampling of some sensory data $\tilde{s}$ and the generative model m entailed by the agent.

The environmental quantities ϑ may be any forces or fields that act upon the agent, such as heat or light-stimulating sensory receptors. In more complex agents, ϑ may also refer to very abstract quantities such as 'social rank' or 'moral norms'. The learning of the environmental quantities ϑ and inferences about their states rest on empirical Bayes and hierarchical generative models [3,4]. In this framework, perceptual learning corresponds to estimating the parameters μ of the recognition density $q(\vartheta; \mu)$ after many sensations, whereas perceptual inference corresponds to inferring the state of ϑ after a single sensation. In a hypothetical environment, learning could refer to the estimation of the categories associated with sensations while inference would be the classification of a particular sensation into one of these categories. In what follows, we shall see how free-energy minimization can account in a unified way for perception, learning and action.

The divergence from the recognition density $q(\vartheta; \mu)$ to the true posterior density $p(\vartheta|\tilde{s},m)$ is measured by the Kullback-Leibler (KL) divergence, which can be decomposed into two quantities known as free-energy and surprise:

$$\begin{aligned} \mathrm{KL}(q(\vartheta;\mu)\|p(\vartheta|\tilde{s},m)) &= \textit{Free-Energy} - \textit{Surprise} \\ &= \int_{\vartheta} q(\vartheta;\mu) \ln \frac{q(\vartheta;\mu)}{p(\tilde{s},\vartheta|m)} d\vartheta + \ln p(\tilde{s}|m) \end{aligned} \quad (1)$$

The first term on the right side of the equation is the free-energy that may be efficiently treated by adjusting the parameters μ in order to minimize the divergence. The second term is surprise, which informs about the probability that some data $\tilde{s}$ has been generated under the assumptions of the model m. In Bayesian model selection, the marginal likelihood $p(\tilde{s}|m)$ is also known as the evidence for the model m. Rearranging (1), one obtains a formulation of free-energy $F(\mu,\tilde{s})$ in terms of divergence plus surprise:

$$\begin{aligned} F(\mu,\tilde{s}) &= \textit{Divergence} + \textit{Surprise} = \\ &\quad \mathrm{KL}(q(\vartheta;\mu)\|p(\vartheta|\tilde{s},m)) - \ln p(\tilde{s}|m) \end{aligned} \quad (2)$$

The free-energy principle states that any adaptive agent that is at equilibrium with its environment must minimize its free-energy [1,2]. Minimizing free-energy with respect to μ reduces divergence, thereby making the recognition density $q(\vartheta; \mu)$ an approximate posterior density $p(\vartheta|\tilde{s},m)$. Notice that divergence depends on μ while surprise does not. Because the divergence is always non-negative (Gibb's inequality), free-energy is said to be an upper bound on surprise.

Crucially, biological agents can minimize free-energy not only by changing their beliefs but also by changing their sensations through acting on the environment. The dependency of sensation $\tilde{s}$ on action α is expressed by $\tilde{s}(\alpha)$. A new rearrangement of (1) shows more clearly what acting on the environment to minimize free-energy $F(\mu,\alpha)$ implies (here, we replace the dependency of free-energy on sensation $\tilde{s}$ by expressing it directly as a function of α):

$$F(\mu,\alpha) = Complexity - Accuracy =$$
$$\mathrm{KL}(q(\vartheta;\mu)\|p(\vartheta|m)) - \langle \ln p(\tilde{s}(\alpha)|\vartheta,m)\rangle_q \quad (3)$$

where $\langle \cdot \rangle_q$ means the expectation under the density q.

In this second formulation, free-energy is expressed as complexity minus accuracy. Complexity is the divergence from the recognition density $q(\vartheta;\mu)$ to the true prior density $p(\vartheta|m)$ of the causes ϑ. Accuracy is the surprise about sensations that are expected under the recognition density; note that accuracy depends on action α whereas complexity does not. This means biological agents will act to minimise free-energy through maximising accuracy. That is, biological agents will act in the environment to sample sensations that are expected by their recognition density.

This perspective on behaviour contrasts with the traditional one in Pavlovian and instrumental conditioning, where behaviour is chiefly understood in terms of maximizing expected reward or pleasure (or conversely minimizing expected loss or pain) [25,26]. In active inference, behaviour is driven by an attempt to fulfil the sensory expectations of posterior beliefs (recognition density). This prevents the dichotomization of the states of the world in terms of pairs of opposites, such as 'reward-loss' or 'pleasure-pain', and implies that the notion of desired states is replaced with that of expected states. States with high probability under the recognition density (low surprise) are more frequently approached whereas states with low probability (high surprise) are avoided by the agent. Agents that expect to visit states that are noxious to their structure will compromise their chances of survival and transmitting their phenotype to future generations (e.g., a rabbit that expects to visit foxes). The adaptive fitness of a phenotype is thus the negative surprise averaged over all the states the agent visits [2].

Emotional valence

In order to harmonize the principled assumption that any biological agent that is at equilibrium with its environment must minimize its free-energy [2] and the traditional notion that humans approach pleasure and avoid pain [27], we related positive and negative valence to the decrease and increase of free-energy over time, respectively. In a continuous time domain, the rate of change of free-energy $F(t)$ is the first time-derivative of free-energy $F'(t)$ at time t. We thus formally define the valence of a state visited by an agent at time t as the negative first time-derivative of free-energy at that state or, simply, $-F'(t)$.

Here, we recall that adaptive agents encode a hierarchical generative model of the causes of their sensations [3,4]. States of the world of increasing complexity and abstraction are encoded in higher levels of the hierarchy, whereas sensory data *per se* are encoded at the lowest level. Free-energy is minimized for each level of the hierarchy separately, and the quantity $F_i(t)$ corresponds to the free-energy associated with the hidden state at the i-th level of the hierarchical model.

According to our definition of emotional valence, when $F_i'(t)$ is positive (i.e., free-energy is increasing over time at level i of the hierarchy) the valence of the state at this level i is negative at time t. When $F_i'(t)$ is negative (i.e., free-energy is decreasing over time at level i) the valence of the state at this level i is positive at time t. When $F_i'(t)$ is zero (i.e., free-energy is constant at level i) the valence of the state at this level i is neutral at time t. Importantly, free-energy is an upper bound on surprise, and neutral valenced states may also be characterized by low or high levels of surprise.

The factorization of emotional valence across levels of the hierarchical model means that positive and negative valence can

be independently attributed to each state in the model, and thus positive and negative valences can be concurrently elicited for the same sensation. Note that free-energy and the rate of change of free-energy are functions not just of current sensations but the posterior beliefs about the causes of those sensations. This means that the free-energy can change in a way that is context-sensitive, depending upon (different) current beliefs about (exactly the same) sensations.

Basic forms of emotion

Cognitive theories of emotion have widely relied on degrees of belief about states of affairs (environmental states) for their analyses of some basic forms of emotion. It has been suggested that a large group of emotions, which includes *happiness, unhappiness, relief* and *disappointment*, is related to certain (firm) beliefs that states of affairs obtain, while a second smaller group of emotions, mainly represented by *hope* and *fear*, is related to uncertain beliefs [28–31]. These two classes of emotions have been referred to as *factive* and *epistemic*, respectively [29]. In philosophy, states of affairs are formally said to either obtain or not whereas beliefs can be true or false (see [32]). Henceforth, we will adopt this terminology.

To illustrate the difference between *factive* and *epistemic* emotions, imagine the case of Lucia who is waiting for a train at the station. Lucia is *happy* that the train is on time (state of affairs p), if she desires p and is certain (i.e., firmly believes) that p obtains. Conversely, Lucia is *unhappy* that p, if she does not desire p and is certain that p obtains. However, Lucia *hopes* that p, if she desires p but is uncertain that p obtains; and, alternatively, Lucia *fears* that p, if she does not desire p but is uncertain that p obtains. On the other hand, *relief* and *disappointment* are better related to the transition from uncertain to certain beliefs [31]. For instance, Lucia is *relieved* that *not-p* if she does not desire p and up to now was uncertain about p, but now is certain that *not-p* obtains; conversely, Lucia is *disappointed* that *not-p* if she desires p and up to now was uncertain about p, but now is certain that *not-p* obtains.

Beliefs and desires can be intuitively related to bottom-up conditional expectations and top-down predictions, respectively, in a predictive coding scheme of free-energy minimization [2]. In this formulation, states of affairs cannot be directly assessed but must be inferred from sensory inputs. Assigning absolute certainty (or zero uncertainty) to any belief impairs the learning of new relationships between sensory inputs and their causes. Here, we consider it more appropriate to circumvent the assumption of certain beliefs proposed in cognitive theories of factive and epistemic emotions, and present a new formulation that relies only on the dynamics of free-energy without any explicit reference to uncertainty. Later, we shall see that factive and epistemic emotions are indeed associated with low and high levels of uncertainty, respectively, but this comes as a consequence and not as a necessary condition of their definition (see Results).

In a continuous time domain, the rate of change of the first time-derivative of free-energy $F_i'(t)$ at the i-th level of the hierarchical model is the second time-derivative of free-energy $F_i''(t)$. By analogy with mechanical physics, $F_i'(t)$ and $F_i''(t)$ can be understood as the velocity and acceleration of free-energy $F_i(t)$ at time t, respectively. Our proposal stands on the assumption that, when both $F_i'(t)$ and $F_i''(t)$ are negative (i.e., free-energy $F_i(t)$ is decreasing 'faster and faster' over time) the agent *hopes* to be visiting a state of lower free-energy in the near future at this level i. However, when $F_i'(t)$ is negative and $F_i''(t)$ is positive (i.e., free-energy is decreasing 'slower and slower' over time) the agent is *happy* to be currently visiting a state of lower free-energy than the previous one at this level i. Equivalently, when $F_i'(t)$ and $F_i''(t)$ are

positive (i.e., free-energy is increasing 'faster and faster' over time) the agent *fears* to be visiting a state of greater free-energy in the near future at this level *i*. However, when $F_i'(t)$ is positive and $F_i''(t)$ is negative (i.e., free-energy is increasing 'slower and slower' over time) the agent is *unhappy* to be currently visiting a state of higher free-energy than the previous one at this level *i*. Additionally, when the rate of change of free-energy $F_i'(t)$ changes sign from negative to positive, the agent is *disappointed* not to be visiting a state of lower free-energy than the current one at this level *i*. Conversely, when $F_i'(t)$ changes sign from positive to negative, the agent is *relieved* not to be visiting a state of higher free-energy than the current one at this level *i*. Finally, when $F_i'(t)$ and $F_i''(t)$ are zero (i.e., free-energy is constant over time) the agent may be low or high *neutrally surprised* in this level *i*. This analysis is summarized in Table 1. Note that since free-energy is minimized for each level of the hierarchical model separately, our formulation also predicts that different emotions can occur concurrently.

The dynamics of free-energy reveal an interesting temporal dependency among the basic forms of emotion. Figure 1 illustrates two hypothetical dynamics of free-energy (top and bottom rows) that elicit distinct patterns of emotion over time (left column). From the two-dimensional space defined by the first and second time-derivatives of free-energy (right column), it becomes clear that transitions from negative to positive emotions can only occur by passing through *relief*, and transitions from positive to negative emotions can only occur by passing through *disappointment*, but transitions between negative (e.g., *fear* and *unhappiness*) or positive (e.g., *hope* and *happiness*) emotions can occur bidirectionally. More importantly, each basic form of emotion is mapped onto a particular region of this two-dimensional space.

Emotional regulation of estimation uncertainty

So far, we have described how emotional valence and some basic forms of emotion can be elicited by the dynamics of free-energy. What, however, is the function of these quantities in a scheme originally developed to explain perception, learning and action? We propose that valence, computed as the negative rate of change of free-energy, is crucial because it informs biological agents about unexpected changes in their world. When valence is positive, sensory inputs fulfil the agent's expectations and the probability of unexpected changes is low. However, when valence is negative, the agent's expectations are violated and unexpected changes in the world are likely to have taken place. In settings

where recent information is a better predictor of states of the world than past information, that is, in a changing world, recent information must be more heavily weighted and, therefore, the learning rate should be high [14]. Conversely, in a stationary world, in which past and recent information are equally informative, the learning rate should be low in order to take into account both past and recent information.

We formalise this notion in terms of emotional meta-learning in which estimation uncertainty is determined not just by free-energy but by the rate of change of free-energy. More specifically, when the free-energy associated with posterior beliefs about states at a particular level in the agent's hierarchical model is increasing, the posterior certainty about these states decreases. In other words, the agent interprets decreasing evidence for its estimates of states of the world as evidence that it is too confident about those states. This can be implemented fairly simply with the augmented Bayesian update:

$$\ln \hat{\sigma}_i(t) = \ln \sigma_i(t) + \lambda F'_i(t) - \tau \qquad (4)$$

$$\sigma_i(t) = \arg \min_{\sigma_i} F_i \qquad (5)$$

$$F_i(t) = -\int q(x_{i-1}, x_i) \ln p(x_{i-1}, x_i) - S(q(x_i)) \qquad (6)$$

Here, the variances $\hat{\sigma}_i(t)$ and $\sigma_i(t)$ correspond to the posterior estimation uncertainty with and without emotional regulation, respectively. The variance $\sigma_i(t)$ is the one that changes to minimize the free-energy $F_i(t)$ at the *i*-th level of the generative model. The quantity $S(\cdot)$ denotes the Shannon entropy, which in this case is a measure of the uncertainty associated with the states at level *i* in the recognition density. The parameter λ can be interpreted as the sensitiveness or 'awareness' of the agent to its emotional valence signals, which informs the agent about changes in the world. The parameter τ represents a long-lasting valenced level that lacks a clear referent, which we thus interpret as mood [23]. The parameters λ and τ are both state and agent dependent. They can also be interpreted as the agent's meta-cognition about the extent to which the agent knows that it does not know the structure of the world.

We have framed the emotional regulation of uncertainty as meta-learning to emphasise that learning (the update) is informed by the consequences of learning, here, the rate of change of variational free-energy. Note that this is a very general scheme that is not tied to any particular generative model. Crucially, expectations about various states, which define them as surprising or not, rest upon prior beliefs that are themselves optimised with respect to variational free-energy; either at an evolutionary timescale or during experience dependent learning.

From equation 4, one can see that positive and negative valence exponentially decreases and increases, respectively, the estimation uncertainty about states of the world. The mood τ induces a constant level of over or under-confidence in the estimates of states irrespective of how surprising the sensory inputs may be. In a negative mood ($\tau < 0$), the agent overweights recent inputs, tracking more easily any volatility in the environment. In a positive mood ($\tau > 0$), the agent overweight past inputs, becoming more attached to past information and less susceptible to tracking environmental changes.

Table 1. Basic forms of emotion and the dynamics of free-energy.

Emotion at time t	Valence	Factive/Epistemic	$F'(t)$[a]	$F''(t)$[b]
happy(t)	positive	factive	<0	>0
unhappy(t)	negative	factive	>0	<0
hopes(t)	positive	epistemic	<0	<0
fears(t)	negative	epistemic	>0	>0
surprised(t)	neutral	factive	0	0
relieved(t)	positive	factive	-0[c]	<0
disappointed(t)	negative	factive	$+0$[d]	>0

[a]First time-derivative of free-energy at time t.
[b]Second time-derivative of free-energy at time t.
[c]Negative value very close to zero.
[d]Positive value very close to zero.
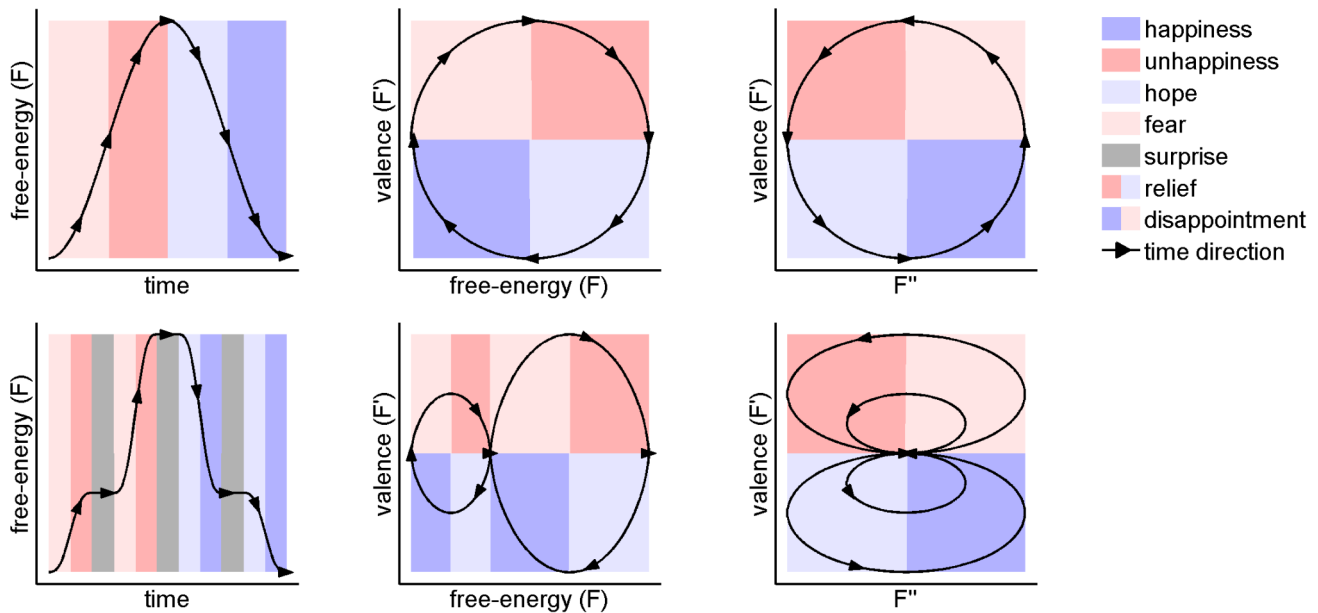doi:10.1371/journal.pcbi.1003094.t001

Figure 1. Basic forms of emotion and the dynamics of free-energy. (top and bottom rows) Two hypothetical dynamics of free-energy and their corresponding basic forms of emotion. (left column) Free-energy F plotted as a function of time. (middle column) The same free-energy F and its first time-derivative F' (valence) as a function of time. (right column) The first F' and second F'' time-derivatives of the same trajectory of free-energy F as a function of time. Notice that the basic forms of emotion are mapped to specific quadrants in the first and second time-derivative spaces independently of the free-energy trajectory. The black arrows indicate the direction of increasing time. The background colours identify the basic forms of emotion elicited at each time point: *happiness* (dark blue), *unhappiness* (dark red), *hope* (light blue), *fear* (light red), *relief* (transition from dark red to light blue), *disappointment* (transition from dark blue to light red) and *surprise* (grey).
doi:10.1371/journal.pcbi.1003094.g001

This emotionally regularized update scheme may appear a bit ad hoc. However, there are several important heuristics in the optimisation literature that are closely related to Equation 4. These are generally described as regularization schemes - for example Levenberg Marquardt Regularization - in which the gradient descent or learning rate is generally decreased when the objective function being optimized does not change as expected. Usually, this regularization can be cast as changing the relative precision of the data at hand. In short, like our scheme, regularization schemes detect a failure in optimization in terms of adverse changes in the objective function (here the free energy) and respond by making more cautious updates - through changing the expected uncertainty about data or prior beliefs. We will see later that, in a hierarchical setting, this can lead to an adaptive change in the rate of optimization or learning at various levels of a hierarchical model.

Results

In this section, we apply equation 4 to demonstrate how one can simulate an emotional agent. In brief, we will compare and contrast two schemes that are exposed to exactly the same sensory inputs and do or do not have emotional updates. The first is based on a hierarchical Bayesian treatment of volatility models that explicitly optimises posterior beliefs about estimation uncertainty. The second uses a simpler generative model that does not optimise the estimates of uncertainty explicitly but implements valence. Using this simpler scheme we show that the same adaptive behaviour can be reproduced using the emotional updates above.

The dynamic perceptual model

Mathys et al. [10] have proposed a generic hierarchical Bayesian scheme that accounts for learning under multiple forms of uncertainty and environmental states. The environmental states can be either discrete or continuous, and the uncertainty can range from probabilistic relations between environmental and perceptual states (perceptual ambiguity) to environmental volatility. Here, we focus only on the simplest discrete and deterministic (i.e., without perceptual ambiguity) environment which nevertheless includes volatility.

In our example of a discrete and deterministic environment, we simulate an agent that learns the probability of a slot machine (one-armed bandit) to generate outcomes (x_1) equal to either \$1 ($x_1 = 1$) or \$0 ($x_1 = 0$). The agent's sensations (u) of the outcomes (x_1) are unambiguous, meaning that $u = x_1$ for both $x_1 = 1$ and $x_1 = 0$. The reward probability of the slot machine is governed by the tendency (x_2) of the machine to generate \$1. In the dynamic perceptual model, the agent knows that the reward tendency may change over time and therefore they also estimate its volatility (x_3).

This discrete and deterministic environment can be formalized with the statement that the sensory input $u^{(k)} \in \{0,1\}$ is binary and the environmental state $x_1^{(k)} \in \{0,1\}$ is the deterministic cause of input $u^{(k)}$ at trial k. The likelihood of state x_1 given sensory input u has the following form (for simplicity, we omit the trial reference k):

$$p(u|x_1) = (u)^{x_1}(1-u)^{1-x_1} \tag{7}$$

Therefore, $u = x_1$ for both $x_1 = 1$ and $x_1 = 0$. At the next level of the hierarchy, the tendency of $x_1 = 1$ (i.e., outcome equal to \$1) is defined by the state $x_2^{(k)} \in \Re$. The probability of $x_1 = 1$ approaches zero when $x_2 \to -\infty$ and approaches one when $x_2 \to \infty$. The mapping from the tendency x_2 to the probability of x_1 is defined

by the following empirical (conditional) prior density:

$$p(x_1|x_2) = s(x_2)^{x_1}(1-s(x_2))^{1-x_1} = \text{Bernoulli}(x_1; s(x_2)) \quad (8)$$

where $s(\cdot)$ is the sigmoid function:

$$s(x) \stackrel{\text{def}}{=} \frac{1}{1+\exp(-x)} \quad (9)$$

It is also assumed that the state $x_2^{(k)}$ at trial k is normally distributed around its value at the previous trial $x_2^{(k-1)}$ with variance $\exp(\kappa x_3^{(k)} + \omega)$. In other words, x_2 evolves in time as a Gaussian random walk:

$$p\left(x_2^{(k)}|x_2^{(k-1)}, x_3^{(k)}\right) = \mathcal{N}\left(x_2^{(k)}; x_2^{(k-1)}, \exp(\kappa x_3^{(k)} + \omega)\right) \quad (10)$$

where the parameters κ and ω are agent dependent.

The state $x_3^{(k)}$ determines the log-volatility of the environment and is represented at the third level of the model. Again, x_3 also evolves in time as a Gaussian random walk but with a step size defined by the constant ϑ that may also differ among agents:

$$p\left(x_3^{(k)}|x_3^{(k-1)}, \vartheta\right) = \mathcal{N}\left(x_3^{(k)}; x_3^{(k-1)}, \vartheta\right) \quad (11)$$

This structure defines a four-level generative model, where ϑ is represented at the last level, and its inversion corresponds to optimizing the posterior densities over unknown hidden states $x = \{x_1, x_2, x_3\}$ and parameters $\chi = \{\kappa, \omega, \vartheta\}$. Here, states and parameters are distinguished in terms of the timescale at which they change. More specifically, states change quickly and parameters change either slowly or not at all for the duration of the observations.

The static perceptual model with emotional valence

Alternatively, we propose a generative model that does not explicitly estimate the volatility (e.g., x_3) of some environmental states (e.g., x_2) but instead makes use of emotional valence (i.e., the negative rate of change of free-energy over time) to assess unexpected changes in the environment. For that purpose, we implement the static perceptual model proposed by Daunizeau et al. [15] with two modifications. First, we consider unambiguous sensory inputs as in Mathys et al. [10] and, second, we use valence to update the posterior variance (estimation uncertainty) of states according to equation 4.

At the first level of the hierarchy, the dynamic model and static perceptual model with valence are exactly the same. At the second level, the static model assumes that the tendency x_2 of outcome x_1 to be equal to \$1 is constant across trials:

$$x_2^{(k)} = x_2^{(0)} : \forall k \quad (12)$$

After inverting this generative model using variational free-energy minimization as described in [10,15], we obtain the updated equations of the posterior distribution of $x_i^{(k)}$, which can be used to investigate the behaviour of the agent on a trial-by-trial basis:

$$\mu_1^{(k)} = u^{(k)} \quad (13)$$

$$\sigma_2^{(k)} = \frac{1}{1/\hat{\sigma}_2^{(k)} + \hat{\sigma}_1^{(k)}} \quad (14)$$

$$\mu_2^{(k)} = \mu_2^{(k-1)} + \sigma_2^{(k)}\delta_1^{(k)} \quad (15)$$

where the following definitions have been used:

$$\hat{\mu}_1^{(k)} \stackrel{\text{def}}{=} s(\mu_2^{(k-1)}) \quad (16)$$

$$\delta_1^{(k)} \stackrel{\text{def}}{=} \mu_1^{(k)} - \hat{\mu}_1^{(k)} \quad (17)$$

$$\hat{\sigma}_1^{(k)} \stackrel{\text{def}}{=} \hat{\mu}_1^{(k-1)}\left(1 - \hat{\mu}_1^{(k-1)}\right) \quad (18)$$

$$\hat{\sigma}_2^{(k)} \stackrel{\text{def}}{=} \sigma_2^{(k-1)}e^{\lambda \nabla F_2^{(k-1)} - \tau} \quad (19)$$

Here, $\mu_1^{(k)}$ and $\mu_2^{(k)}$ are the posterior expectations of x_1 and x_2 after sensory input $u^{(k)}$, which can be interpreted as the expected probability and the expected tendency of reward, respectively. Accordingly, the uncertainty $\sigma_2^{(k)}$ is the posterior variance of x_2. The prediction error at the first level $\delta_1^{(k)}$ is the difference between the expectation $\mu_1^{(k)}$ and the prediction $\hat{\mu}_1^{(k)}$ before seeing the input $u^{(k)}$. Equivalently, $\hat{\sigma}_1^{(k)}$ is the variance of the prediction $\hat{\mu}_1^{(k)}$ before seeing the input $u^{(k)}$.

In order to adapt to unexpected changes in the environment, the agent needs to update the posterior variance $\sigma_2^{(k)}$ proportionally to the valence of the state x_2 at time $k-1$. In discrete time, the valence of the state $x_i^{(k)}$ is, by definition, the negative first backward difference of free-energy $\nabla F_i^{(k)}$ at time k:

$$\nabla F_i^{(k)} \stackrel{\text{def}}{=} F_i^{(k)} - F_i^{(k-1)} \quad (20)$$

Specific to the proposed generative model, the free-energy $F_2^{(k)}$ of state $x_2^{(k)}$ is:

$$
\begin{aligned}
F_2^{(k)} = & -E\left[\ln p\left(x_1^{(k)}|x_2^{(k)}\right) + \ln p\left(x_2^{(k)}\right)\right] - S\left(q\left(x_2^{(k)}\right)\right) \\
= & -\mu_2^{(k)}\left(\mu_1^{(k)} - 1\right) - \ln s\left(\mu_2^{(k)}\right) - \frac{1}{2}\left(s\left(\mu_2^{(k)}\right)^2 - s\left(\mu_2^{(k)}\right)\right)\sigma_2^{(k)} \\
& + \frac{1}{2\hat{\sigma}_2^{(k)}}\left(\left(\mu_2^{(k)} - \mu_2^{(k-1)}\right) + \sigma_2^{(k)}\right) + \frac{1}{2}\ln\hat{\sigma}_2^{(k)} + \frac{1}{2}\ln 2\pi \\
& - \frac{1}{2}\ln\sigma_2^{(k)} - \frac{1}{2}\ln 2\pi e
\end{aligned}
\quad (21)
$$

where the expectation is taken under the approximate posterior densities $q\left(x_1^{(k)}\right)$ and $q\left(x_2^{(k)}\right)$.

The parameters λ and τ are constant and dependent on the agent. They represent the sensitiveness to emotional valence and the mood of the agent, respectively. According to our assumptions, the uncertainty of a hidden state x_i should increase or decrease when its valence $-\nabla F_i^{(k)}$ is negative or positive, respectively. Therefore, λ is constrained within the interval $[0,\infty]$. Notice that, when λ and τ are equal to zero, the static perceptual model with valence becomes the same as the standard static perceptual model described in [15].

The reference scenario

Having defined the two competing schemes, we implemented two agents under the dynamic perceptual model (DP) and the static perceptual model with valence (SPV), hereafter referred to as the DP agent and the SPV agent. These agents were exposed to 320 sensory inputs (outcomes) sampled from a three-stage reference scenario as proposed in [10]. In the first stage (low volatility) of the scenario, the agents were exposed to a sequence of 100 outcomes where the probability of $x_1=1$ (outcome equal to \$1) was 0.5. In the second stage (high volatility), the probability that $x_1=1$ alternated between 0.9 and 0.1 every 20 inputs. Finally, in the third stage (low volatility again), the first 100 outcomes were repeated in exactly the same order. The initial values of the hidden states x_2 and x_1 were $\mu_2^{(0)}=0$, $\sigma_2^{(0)}=1$ and $\mu_1^{(0)}=s(\mu_2^{(0)})$ for both the DP and SPV models. In the DP model, the initial values of the hidden state x_3 were $\mu_3^{(0)}=-0.4$ and $\sigma_3^{(0)}=1$.

We replicated the results reported by Mathys et al. [10] for the DP model with the same parameters $\vartheta=0.5$, $\omega=-2.2$ and $\kappa=1.4$ (see Figure 2). Overall, the posterior expectation of x_1, which is the reward probability, fluctuated around the true probability of $x_1=1$ both in the low and high volatility stages. Nevertheless, one can observe increasing instability during the third stage relative to the first, even though the inputs were presented exactly in the same order in both of them. Mathys et al. [10] explained this in terms of a strong tendency for the agent to increase its posterior expectation of log-volatility μ_3 in response to surprising stimuli (given the parameters used in the reference scenario). The increase of μ_3 was followed by an increase in the posterior variance σ_2 of state x_2, which regulates the learning rate at the second level. Despite the different levels of volatility in each stage, the posterior variance σ_2 smoothly increased with a constant rate during the whole scenario.

We first evaluated the SPV model setting both the sensitiveness λ and mood τ equal to zero. In this case, the agent learns according to a standard static perceptual model and is completely insensitive to any volatility or unexpected change in the environment. As illustrated in Figure 3, the posterior expectation of $x_1=1$ converges to 0.5, which is the true probability of $x_1=1$ across the three (low and high volatility) stages. Concomitantly, the posterior variance (estimation uncertainty) σ_2 asymptotically decreases toward zero, reflecting the decreasing uncertainty of the estimates across sensory inputs.

When setting the parameters λ and τ to values different than zero, the agent becomes sensitive to changes in its environment. In Figure 4, one can observe the effect of mood τ alone. When τ is set to -0.13 and λ is kept equal to 0, a negative mood is sufficient to make the SPV model reactive to the volatility of the environment similar to the DP model. Importantly, the dynamic model also has a constant parameter ω that is agent dependent, which has a similar function to τ in our model. Nevertheless, the SPV does not

show the increasing instability in the last (low volatility) stage observed in the DP model. In fact, the posterior variance σ_2 returns to a stable baseline even after the increased fluctuation during the high volatility stage.

With the addition of emotional valence to the model, the agent becomes even more reactive and is able to track fast changes in the environment. In Figure 5, the sensitiveness λ is set to 0.8. The posterior variance σ_2 now changes more quickly in response to surprising sensory inputs and there is a clear distinction between the low and high volatility stages. More specifically, the elicitation of negative valence is the main cause of increases in σ_2, whereas positive valence causes σ_2 to decrease. Despite the phasic reaction to unexpected changes during the high volatility stage, the agent returns again to a fairly stable baseline similar to the first low volatility stage in the last low volatility stage.

Critically, an optimal tracking of environmental volatility requires mood to be set to some appropriate negative value. An extremely low mood, characterized by a large negative tau, would cause a very large increase in estimation uncertainty, consequently impairing discrimination between high and low volatility stages.

Uncertainty associated with factive and epistemic emotions

We also investigated the estimation uncertainty associated with the *factive* (*happiness* or *unhappiness*) and *epistemic* (*fear* or *hope*) emotions in the reference scenario. It is noteworthy that we defined these emotions simply in terms of the dynamics of free-energy without any assumptions about uncertainty, contrary to the traditional analysis of these emotions in psychology and philosophy (see [28–31]). For this purpose, we performed 100 realizations of the reference scenario (i.e., we repeated the simulation with the reference scenario 100 times, sampling new sensory inputs at each time) and we computed the mean of the posterior variance (estimation uncertainty) $\hat{\sigma}_2$ of state x_2 immediately after the onset of *factive* and *epistemic* emotions. The posterior variance $\hat{\sigma}_2$ represents the change in estimation uncertainty after the elicitation of the emotion and before the observation of the next sensory input (see equation 19). For this analysis, we set the sensitiveness λ to an intermediate value equal to 0.4 and we kept the mood τ equal to -0.13.

The distribution of the mean $\hat{\sigma}_2$ across simulations grouped within the low and high volatility stages of the reference scenario is shown in Figure 6. In both the low and high volatility stages, the mean $\hat{\sigma}_2$ was higher on average for the *epistemic* (low volatility: M = 0.68, SD = 0.03; high volatility: M = 1.07, SD = 0.19) than the *factive* (low volatility: M = 0.58, SD = 0.02; high volatility: M = 0.69, SD = 0.06) emotions. Furthermore, the mean $\hat{\sigma}_2$ was also higher on average during the high (M = 0.88, SD = 0.24) than the low volatility (M = 0.63, SD = 0.06) stages.

Discussion

In this paper, we have proposed a biologically plausible computational model of emotional valence inspired by the free-energy principle. The mathematical definition of emotional valence in terms of the negative rate of change of free-energy not only accounts for how beliefs determine emotions but also provides a formal account of how emotions determine the content and the degree of posterior beliefs (see [19]). In our framework emotional valence regulates estimation uncertainty signalling unexpected changes in the world, thereby performing an important meta-learning function.
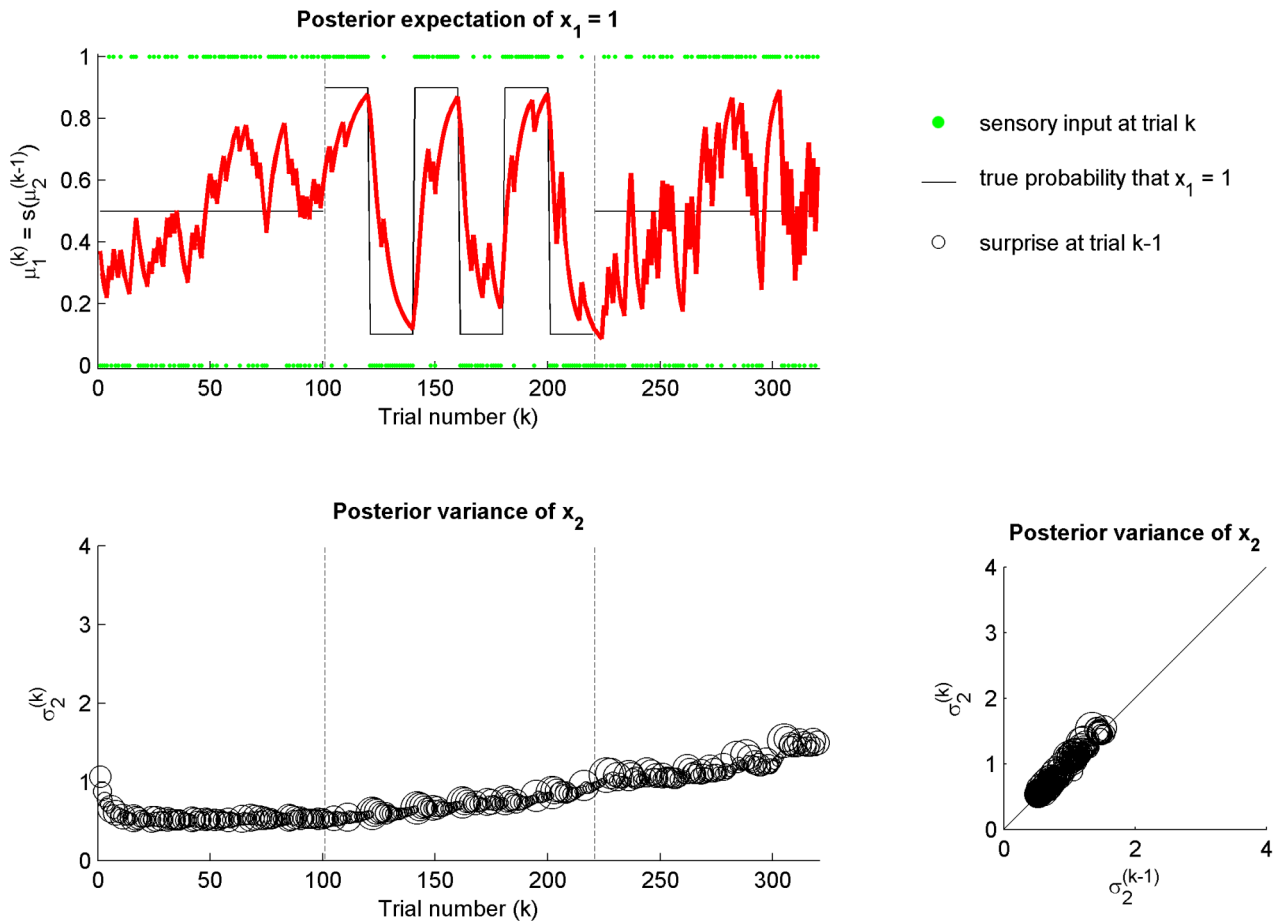
Figure 2. Dynamic perceptual model: $\vartheta = 0.5$, $\omega = -2.2$ and $\kappa = 1.4$. A simulation of 320 trials. The first (low volatility), second (high volatility) and third (low volatility) stages are separated by vertical dashed lines. (top) The agent's posterior expectation $s(\mu_2)$ that $x_1 = 1$ (red line) after sensory input u (green dots), is plotted over the true probability that $x_1 = 1$ (black line), which is unknown to the agent. (bottom left) The time course of the posterior variance σ_2 of x_2 over trials. The size of the black circles is proportional to the surprise of sensory input u at trial k. (bottom right) The change in the posterior variance of x_2 from trial $k-1$ to trial k as a function of the surprise of sensory input u at trial k.
doi:10.1371/journal.pcbi.1003094.g002

The relationship between emotional valence and state transition also finds support in previous studies of emotion (see [33–36]). Batson et al. [35] have argued that the shift from a less valued state (i.e., high free-energy) to a more valued state (i.e., low free-energy) is accompanied by positive affect, while a shift in the opposite direction is accompanied by negative affect. Likewise, Ben-Zeev [36] has suggested that emotions are generated when the level of stimulation we have experienced for long enough to get accustomed to it changes, and the change, rather than the general level, is of emotional significance. Accordingly, in the words of the same author, "loss of satisfaction does not produce a neutral state, but misery, and loss of misery does not produce a neutral state either, but happiness" [36].

Similar situations can also be found when people are entertained by magicians or humorists. In both cases, following the surprise elicited by the apparent violation of the physical laws in magic [37] or the incongruity of the situation in humour [38], greatest pleasure is experienced when the trick or the joke is understood. Our suggestion is that pleasure is elicited in the transition from a state of high to low surprise. Critically, magic tricks are performed on a stage where people know that there is no real violation of the physical laws; if such surprising events would happen in everyday life, they would probably be experienced as quite disturbing and unpleasant.

According to our scheme, emotional valence is not estimated itself by the agent but emerges naturally from the process of estimating hidden states by means of free-energy minimization. One could eventually hypothesize that some living organisms, such as humans, explicitly represent valence as one of the causes of their sensations. This means that these agents should also estimate valence (and its uncertainty) like any other hidden state in their generative model. Nevertheless, the explicit representation of valence is not a requirement for emotional valence to exist in our scheme and to play an important role in the adaptation of biological agents to unexpected changes in their world.

To put our valence-based meta-learning scheme to a test, we compared two competing agents in a non-stationary environment. The SPV agent with valence replicated the behaviour of the DP agent that explicitly estimated the volatility of the environment [10]. Nevertheless, the adaptive fitness of the SPV agent to unexpected changes was achieved with the representation of only two hidden states $x = \{x_1, x_2\}$ and two parameters $\chi = \{\tau, \lambda\}$, whereas the DP agent required three hidden states $x = \{x_1, x_2, x_3\}$ and three parameters $\chi = \{\kappa, \omega, \vartheta\}$. More importantly, the two parameters λ and τ of the SPV agent have a clear psychological interpretation in terms of sensitiveness to emotional valence and mood, respectively. The mood τ was shown to be important for
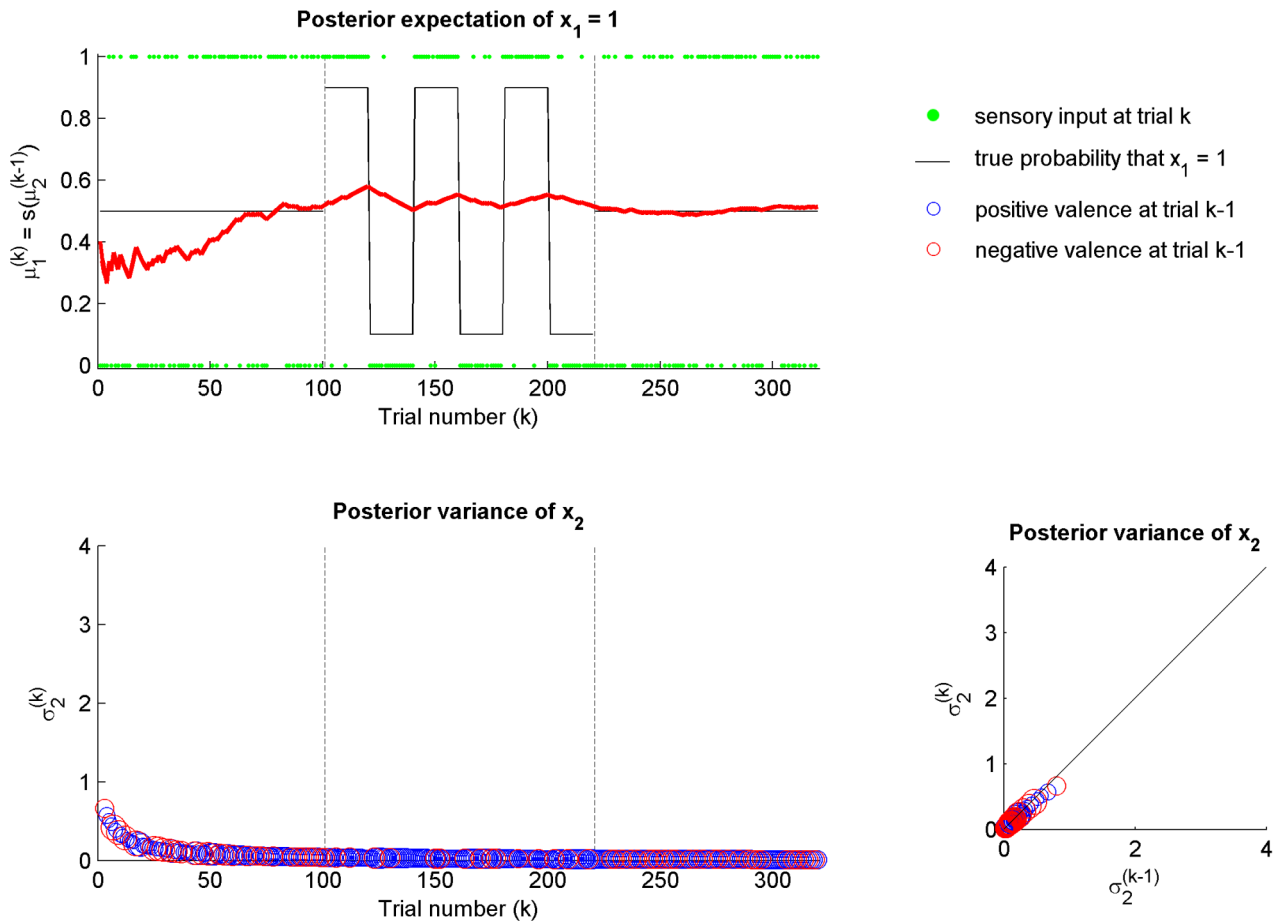
Figure 3. Static perceptual model: $\lambda = 0$ **and** $\tau = 0$**.** The agent is exposed to the same sequence of sensory inputs described in the reference scenario (see Figure 2 for legends). (top) The posterior expectation of $x_1 = 1$ converges to 0.5, which is the true probability of $x_1 = 1$ across the three (low and high volatility) stages. The agent is unaware of unexpected changes in the environment. (bottom left) The posterior variance (estimation uncertainty) σ_2 of x_2 asymptotically converges to zero across trials k. Negative (red circle) and positive (blue circle) valences are indicated when elicited over the trial. The size of the circles is proportional to the surprise of the sensory input u at trial k. (bottom right) The change in the posterior variance of x_2 from trial k to trial $k+1$ as a function of negative (red circle) and positive (blue circle) valences.
doi:10.1371/journal.pcbi.1003094.g003

tracking slow and continuous changes in the environment, known as volatility, whereas the sensitiveness λ was shown to be crucial for tracking fast and discrete changes, known as unexpected uncertainty. The proposed scheme is very general and does not rely on any particular generative model of how sensory inputs are caused, meaning that it can account for any internal model of the world that defines a particular agent (see [7]).

We also investigated the relationship between estimation uncertainty and *factive* (*happiness* as well as *unhappiness*) and *epistemic* (*hope* and *fear*) emotions. Although psychologists and philosophers have traditionally relied on degrees of belief (uncertainty) in their analyses of these families of emotion [28–31], we alternatively relied only on the dynamics of free-energy. In agreement with these more traditional analyses, we found that *epistemic* emotions are indeed more related to higher levels of (estimation) uncertainty than *factive* emotions. However, at the algorithmic level, we reiterate our claim that the computational quantity that unambiguously distinguishes between *factive* and *epistemic* emotions is not degrees of belief, as previously proposed [31], but rather the temporal dynamics of free-energy.

More important for psychological perspectives on emotion, the trajectory invariant representation of emotions in the state space

defined by the first and second time-derivatives of free-energy also recapitulates the dimensional view of emotion [39]. Although the first time-derivative of free-energy $F'(t)$ has been intuitively related to the dimension of valence, it is still unclear how to interpret the second time-derivative $F''(t)$ in terms of a psychological construct. The emergence of some forms of emotion, tentatively labelled as *happiness, unhappiness, hope, fear, disappointment* and *relief*, also provides support for the notion of basic emotions [40], in the sense that these emotions are exclusively related to very precise dynamics of free-energy. Furthermore, our scheme also encompasses important aspects of cognitive models of emotion [31,41,42], in the sense that states of the world (e.g., agents, objects, events), which are relevant for the diversity and complexity of human emotions, can be accounted for within the hierarchical generative model entailed by the agent. To illustrate, happiness (unhappiness) has been related to the negative (positive) first time-derivative and the positive (negative) second time-derivative of the free-energy of some state in the generative model. When the state under consideration is the fate of another person, this can be understood as a specific form of happiness (unhappiness) usually known as 'joy for another' (pity) [31].
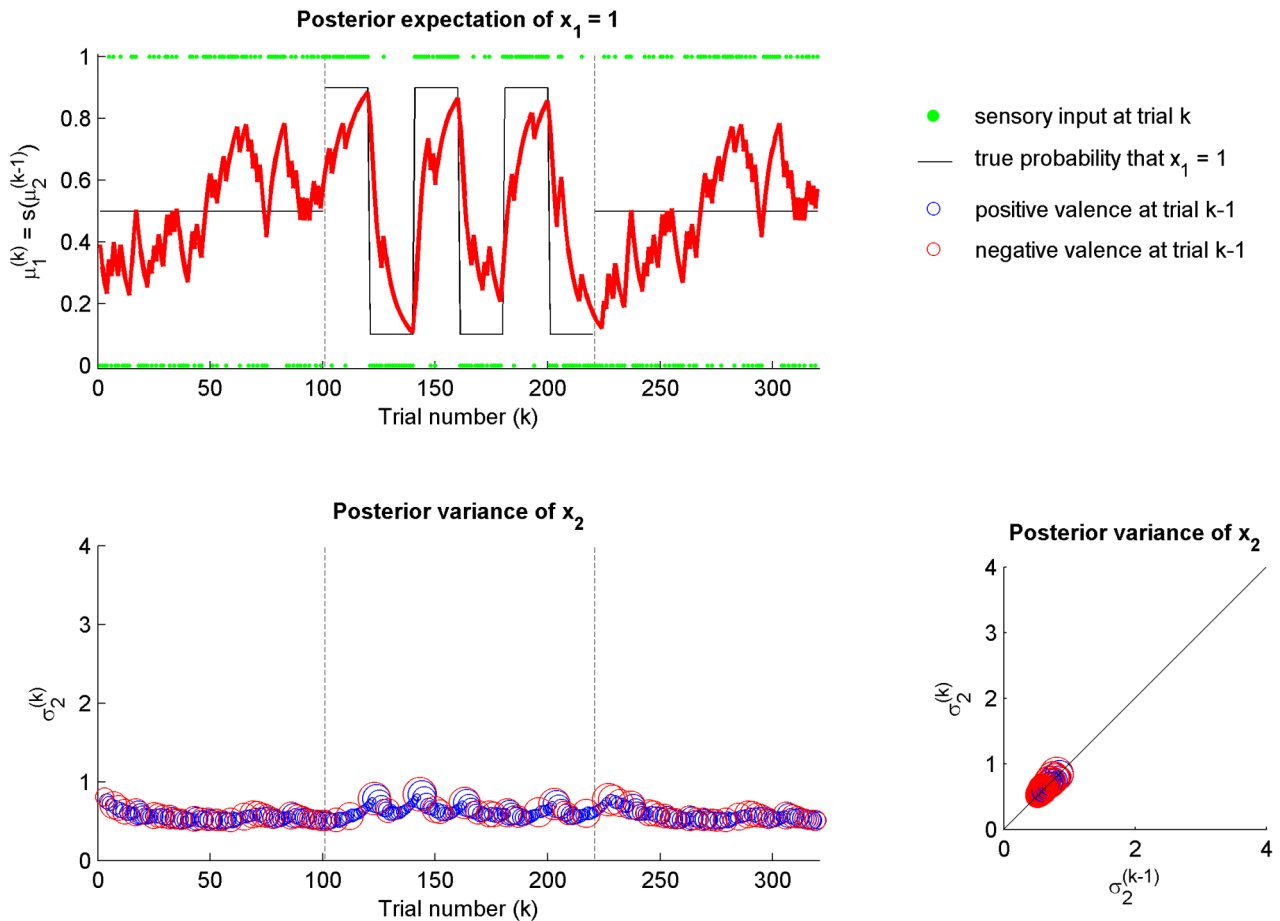
Figure 4. Static perceptual model with valence: $\lambda = 0$ **and** $\tau = -0.13$**.** The agent is exposed to the same sequence of sensory inputs described in the reference scenario (see Figure 2 for legends). Now, the agent becomes reactive to unexpected changes in the environment. (top) The posterior expectation of $x_1 = 1$ fluctuates around the true probability of $x_1 = 1$ at each stage in a manner similar to the dynamic perceptual model (see Figure 2). (bottom left) The posterior variance (estimation uncertainty) σ_2 maintains a constant baseline during the first and third (low volatility) stages mainly defined by the mood, but starts to show a tendency to fluctuate more freely during the second (high volatility) stage. (bottom right) The change in the posterior variance of x_2 from trial k to trial $k+1$ as a function of negative (red circle) and positive (blue circle) valences is quite similar to the standard static model (see Figure 3), except for a small offset defined by the mood.
doi:10.1371/journal.pcbi.1003094.g004

The concept of value has been largely related to valence in social and affective psychology (see [43]). Our definition of emotional valence in terms of the rate of change of free-energy also provides a formal distinction between valence and value. In the free-energy principle, value is the complement of free-energy in the sense that minimizing free-energy corresponds to maximizing the probability that an agent will visit valuable states, where the evolutionary value of a phenotype is the negative surprise averaged over all the (interoceptive and exteroceptive) sensory states it experiences [2]. This formulation parallels a recently proposed reinforcement learning theory for homeostatic regulation [44], which attempts to integrate reward (valence) maximization with the minimization of departures from homeostasis (free-energy).

Our scheme is also broadly compatible with the predictive coding model of conscious presence [45], which claims that interoceptive inference is the constitutive basis of the *subjective experience of emotions*. Although our formulation treats interoceptive and exteroceptive predictions (and their uncertainty) on an equal footing, one might imagine that prediction of interoceptive states would be a particularly important target for emotional regulation. This is because, from an evolutionary perspective, it is important to maintain a physiological homeostasis and respond adaptively to

any unpredicted changes in the internal milieu. Furthermore, the putative emphasis on interoception provides a close link between (literally) 'gut feelings' and the computational (inferential) role of emotion that we have described above.

An apparent paradox that might emerge from our definition of emotional valence is related to the common sense notion that both the violation and the fulfilment of expectations can be either positive or negative. As we stated before, according to our scheme, the fulfilment of expectations must always elicit positive emotions whereas the violation of expectations must always elicit negative emotions. Therefore, how can the *subjective experience* of *positive surprises* and *negative expectations* be accounted for within our scheme?

In our perspective, these experiences emerge from a confound between the fulfilment and the violation of expectations across different levels of the hierarchical generative model. To illustrate this, we first need to recall that in the Bayesian brain formulation, agents encode a hierarchical generative model of the causes of their sensations, where states of the world of increasing complexity and abstraction are encoded in higher levels of the hierarchy and sensory data per se are encoded at the lowest level. Let us imagine the case of an old friend who suddenly steps in our door. This unexpected visit can be intuitively related to the experience of a
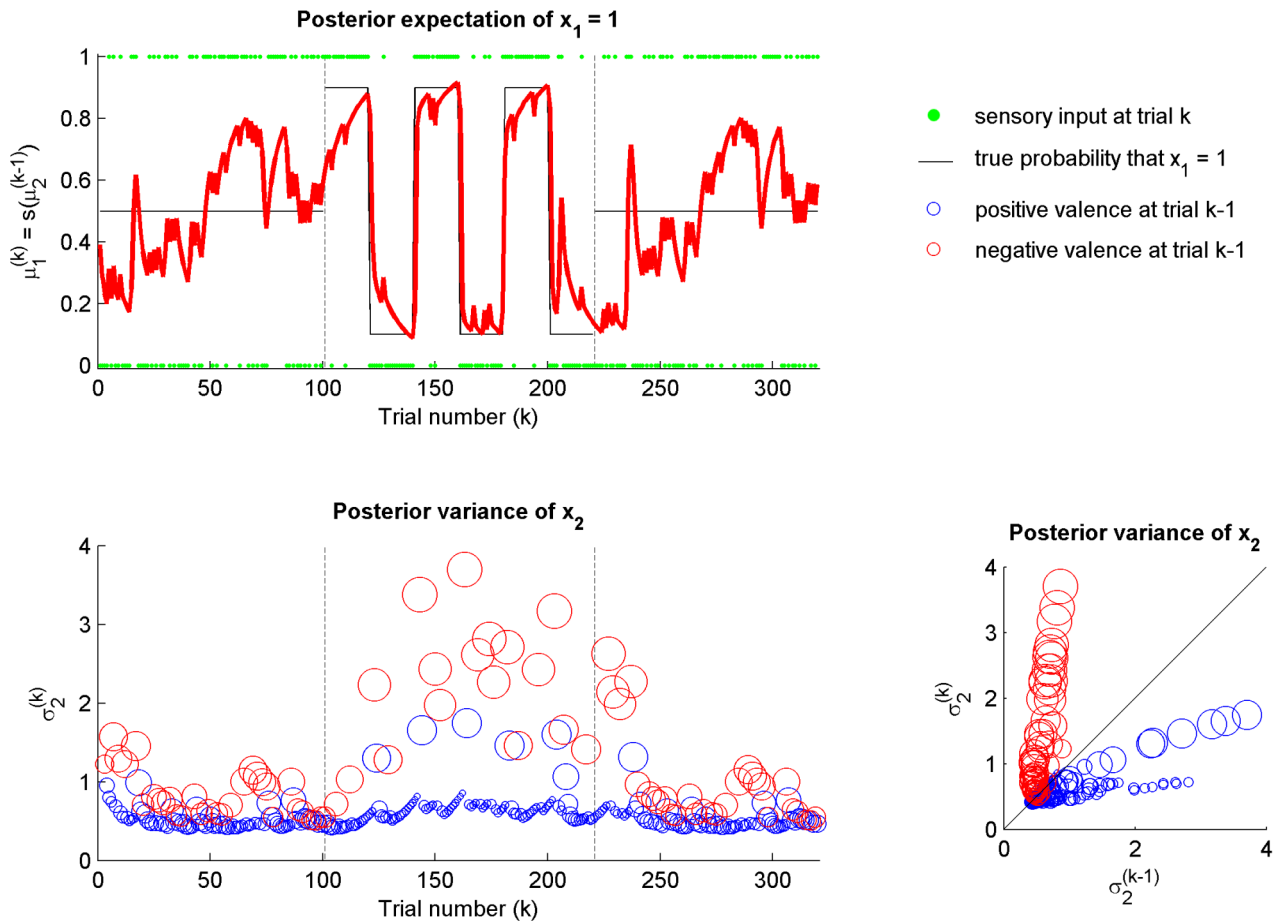
Figure 5. Static perceptual model with valence: $\lambda = 0.8$ **and** $\tau = -0.13$**.** The agent is exposed to the same sequence of sensory inputs described in the reference scenario (see Figure 2 for legends). Now, the agent becomes extremely reactive to unexpected changes in the environment. (top) The posterior expectation of $x_1 = 1$ changes more quickly and is closer to the true probability of $x_1 = 1$ at each stage. (bottom left) The posterior variance (estimation uncertainty) σ_2 maintains a constant baseline during the first and third (low volatility) stages mainly defined by the mood, but it fluctuates more widely during the second (high volatility) stage. This clarifies the distinction between the low and high volatility stages. Negative (red circle) and positive (blue circle) valences are clearly associated with increases and decreases in uncertainty, respectively, and they become more intense during the second (high volatility) stage. (bottom right) The posterior variance of x_2 from trial k to trial $k+1$ increases after negative valence but decreases after positive valence.
doi:10.1371/journal.pcbi.1003094.g005

very positive and surprising emotion. However, a more careful analysis can unveil which aspects of this experience are indeed surprising and which are just as expected, given a hierarchical generative model of how sensations are caused. Assuming that our friend has moved to a distant city many years ago, the sudden apparition of this friend certainly violates any expectation about the physical causes of sensations. It would be very surprising to meet a friend at our door when they are expected to be miles away - no matter how beloved they might be. Such a surprising sensation should elicit unpleasantness at the corresponding levels of the model where physical causes of sensations are encoded. Concomitantly, this same sensation should also fulfil more abstract expectations that we might have of being close to beloved ones. The fulfilment of these expectations should conversely elicit pleasantness at higher levels of the generative model where these more abstract causes of sensations are probably represented. With the formalism of a hierarchical generative model, the causes of sensations can be clearly defined and their respective valence properly investigated. In the example above, we would thus consider it more precise to say that 'we are surprised about the unexpected visit of a friend but happy to be close to a beloved

one'. Here, our explanation rests upon the assumption that the *subjective experience of emotion* usually confounds the increasing fulfilment (pleasantness) and violation (unpleasantness) of expectations across different levels of the hierarchical model.

In another example, the reasoning above also can help us to explain how our scheme may account for sensations that are expected but of negative valence (e.g., the expectation of an eminent injury). Let us imagine the case of someone who is walking on the street and suddenly sees a cyclist riding a bicycle dangerously. As the cyclist gets closer, the person becomes increasingly confident that they will be hit by the bicycle. In this situation, the movement of the bicycle fulfils the expectations of the person about how physical bodies should move in the world and, therefore, it elicits pleasantness at those levels of the generative model. Indeed, it would be very surprising (and unpleasant at these levels) if the bicycle suddenly disappeared or made an unexpected movement that violated the physical laws of motion. Nevertheless, the approach of the bicycle also violates other expectations regarding the safety of walking down the street, which are probably represented at different levels of the hierarchical model. At these levels, the approach of the bicycle
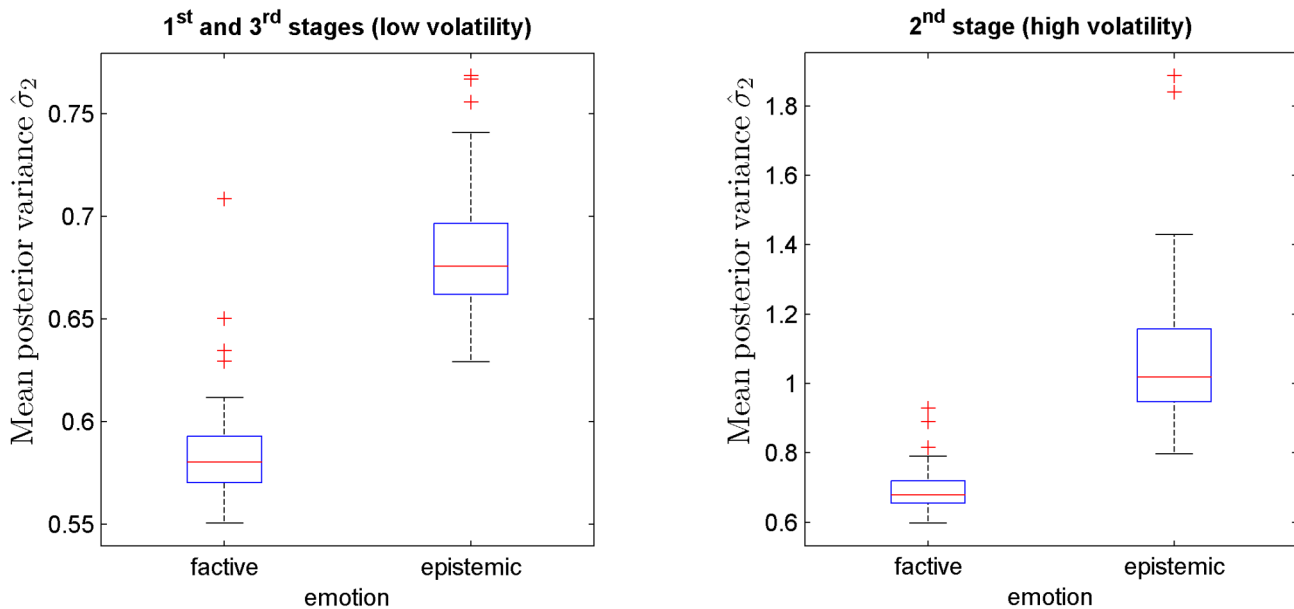
Figure 6. Boxplots of the mean posterior variance $\hat{\sigma}_2$ of state x_2 after the elicitation of *factive* (*happiness* or *unhappiness*) and *epistemic* (*fear* or *hope*) emotions and before the observation of the next sensory input. (left) Mean posterior variance $\hat{\sigma}_2$ during the low volatility stages of the reference scenario. (right) Mean posterior variance $\hat{\sigma}_2$ during the high volatility stages of the reference scenario. The mean $\hat{\sigma}_2$ was computed for each of 100 simulations of the reference scenario. In both the low and high volatility stages, the mean $\hat{\sigma}_2$ was on average higher for the *epistemic* (low volatility: M = 0.68, SD = 0.03; high volatility: M = 1.07, SD = 0.19) than the *factive* (low volatility: M = 0.58, SD = 0.02; high volatility: M = 0.69, SD = 0.06) emotions and it was also on average higher during the high (M = 0.88, SD = 0.24) than the low volatility (M = 0.63, SD = 0.06) stages.
doi:10.1371/journal.pcbi.1003094.g006

is very unpleasant and becomes even more unpleasant when the person is indeed injured by the bicycle. Again, in this case, we would consider more precise to say that 'the person expects the bicycle to hit them - under such environmental conditions - but they do not expect to be injured when walking down the street'.

The flexibility of our scheme to accommodate different generative models may raise some concerns regarding the falsifiability of our theory. However, we would like to clarify that hypotheses derived from our theory should be tested conditional on a particular generative model. Especially given the known diversity of phenotypes in nature, we consider that this flexibility is more a strength than a weakness. Furthermore, generative hierarchical models and free-energy minimization provide a principled way to represent the relationship between hidden states and to understand their dynamics. Nevertheless, further empirical work is still required to better understand at which levels of the hierarchical generative model the violation of expectations might be more closely related to the *subjective experience* of surprise and emotional valence. Our intuition is that the *subjective experience of surprise* is more closely related to violations at lower levels of the hierarchy, whereas the *subjective experience of emotional valence* is more closely related to violations at higher levels.

The distinction between violation and fulfilment of expectations across different levels of the generative model might also help us to further disambiguate the *subjective experience* of other emotions such as *fear* and *anxiety*, which have an important role in psychopathology. One of the ways in which cognitive theories of emotion have distinguished *fear* from *anxiety* is based on the physical and existential aspect of their causes. *Fear* involves threats that are concrete and sudden, whereas *anxiety* is related to threats that are more symbolic, existential and ephemeral [41,42]. Nevertheless, both *fear* and *anxiety* are related to the prospect of visiting unpleasant states in the future, which in our scheme has been

related to a 'faster and faster' increase of free-energy over time. To illustrate, let us imagine the case of a spider-phobic person who is presented with a spider. The *subjective experience* of *fear* in this case could be explained as the product of (1) a 'slower and slower' increase in the violation of the expectations about the more physical causes of sensations, which encodes the physical recognition of the spider, eliciting unhappiness at these levels; and (2) a 'faster and faster' increase in the violation of the expectations about more abstract causes of sensations, such as the increasing probability of being bitten by the spider, eliciting fear at these levels. However, in the case of *anxiety*, there seems to be incongruence between the violation of expectations about the physical and the existential causes of sensations. Therefore, in our perspective, the *subjective experience* of *anxiety* should be expressed as the product of (1) a stationary violation of the expectations about the physical causes of sensations (i.e., the environment is physically perceived as usual) bringing neutrality to these levels, and (2) a 'faster and faster' increase in the violation of the expectations about more abstract/existential causes of sensations, eliciting fear at these levels. This incongruence of violation across levels of the generative model could explain the difficulty that anxious people have to attribute concrete causes to their fears.

Our formulation of emotional valence might also be of importance in the investigation of affective and other mental disorders, such as depressive and anxiety disorders [46]. For instance, when we use our model to explain major depressive disorder (MDD), which is a complex debilitating psychiatric condition that is largely characterized by persistent low mood and decreased interest or pleasure in usually enjoyable activities [47], we immediately find the crucial role played by our mood model parameter τ. In our meta-learning scheme, when mood is low ($\tau < 0$), the estimation uncertainty of environmental states is overestimated and top-down predictions become under confident.

Theoretical computational simulations has shown that pathological under confidence in top-down predictions can impair behaviour due to a failure in eliciting sufficient sensory prediction errors [48]. Consequently, the agent reacts less vigorously toward, or away from stimuli that might have been previously evaluated as pleasant or unpleasant. In fact, several studies have reported that clinically depressed individuals spend significantly more time looking at negative stimuli [49–52]. A subsequent, and cyclical, increase in mood ($\tau > 0$) could eventually explain manic episodes in bipolar disorders [53]. Manic episodes are characterized by a distinct period during which patients experience abnormally and persistently elevated, expansive, or irritable mood [54]. In fact, a pathological increase in the precision of top-down predictions has also been shown to induce perseverative behaviours [48]. It would be interesting to investigate how mood induction in healthy subjects might affect their performance on tasks where tracking volatility is necessary. According to our theory, we would predict that subjects with mood levels below and above the optimum for tracking some particular level of environmental volatility should benefit from positive and negative mood induction, respectively. More precisely, an inverted U-shaped performance curve is predicted with depressed and manic patients found at the lowest and highest extremes of the mood range.

A reasonable approach to test hypotheses derived from our theory would be to invert a generative model (i.e., estimate the unknown model parameters) for the experimental task at hand using variational Bayes [55]. The free-energy computed during this inversion process can then be exploited to estimate the emotions at different levels of the hierarchical generative model according to our scheme. A complete characterization of the generative model could eventually be relaxed if a direct measure of the free-energy or, under simplifying assumptions, prediction error is also available. Indeed, the quantity that matters for testing our emotional valence hypothesis is the rate of change of free-energy rather than the generative model itself.

Future empirical work should investigate the correlation between the estimated emotional valence (i.e., the first time-derivative of free-energy) and verbal-reports of valence for a variety of experimental conditions. As previously mentioned, free-energy is an upper bound on surprise and its minimization also entails prediction error reduction. In this perspective, recording prediction error signals in the brain, computing their temporal derivatives and correlating them to verbal-reports of valence could be a suitable procedure. Human neuroimaging studies have shown that the orbitofrontal cortex plays an important role in linking different types of reward to hedonic experience (see [56]). Orchestrated with the striatum [57], which has been traditionally implicated in reward prediction error [58] and saliency [59], those two regions might be of particular relevance to the investigation of our scheme in the brain. In biologically plausible implementations of free energy minimisation, precision (i.e., the inverse of uncertainty) is encoded by the gain of cells reporting prediction error [2]. This directly implicates the classical ascending neuro-modulatory transmitter systems like dopamine, acetylcholine and norepinephrine in the encoding of uncertainty. The diverse and complex interactions between these neurotransmitters and their role in encoding different forms of uncertainty are still far from being clearly understood [11,60,61]. Future work will address how our meta-learning scheme, which links the rate of change of free-energy (prediction error) to estimation uncertainty (precision), can help in elucidating the complex interaction between these neurotransmitters and the activity in their target brain areas.

To conclude, by providing a general framework in which different perspectives on emotion can be formally interrelated, and by demonstrating how emotional valence can dynamically regulate uncertainty, we hope to contribute to paving the way for future computational studies of emotion in learning and uncertainty.

Acknowledgments

Author Contributions

Conceived and designed the experiments: MJ. Performed the experiments: MJ. Analyzed the data: MJ. Contributed reagents/materials/analysis tools: MJ GC. Wrote the paper: MJ GC.

References

1. Friston K, Kilner J, Harrison L (2006) A free energy principle for the brain. Journal of Physiology, Paris 100: 70–87.
2. Friston K (2010) The free-energy principle: a unified brain theory? Nat Rev Neurosci 11: 127–138.
3. Friston K (2003) Learning and inference in the brain. Neural Networks: The Official Journal of the International Neural Network Society 16: 1325–1352.
4. Friston K (2005) A theory of cortical responses. Philosophical Transactions of the Royal Society B: Biological Sciences 360: 815–836.
5. Friston K, Daunizeau J, Kiebel SJ (2009) Reinforcement learning or active inference? PLoS ONE 4: e6421.
6. Friston K, Daunizeau J, Kilner J, Kiebel SJ (2010) Action and behavior: a free-energy formulation. Biological Cybernetics 102: 227–260.
7. Friston K (2008) Hierarchical models in the brain. PLoS Comput Biol 4: e1000211.
8. Yoshida W, Ishii S (2006) Resolution of uncertainty in prefrontal cortex. Neuron 50: 781–789.
9. Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS (2007) Learning the value of information in an uncertain world. Nature Neuroscience 10: 1214–1221.
10. Mathys C, Daunizeau J, Friston KJ, Stephan KE (2011) A bayesian foundation for individual learning under uncertainty. Frontiers in Human Neuroscience 5: 1–20.
11. Yu AJ, Dayan P (2005) Uncertainty, neuromodulation, and attention. Neuron 46: 681–692.
12. Payzan-LeNestour E, Bossaerts P (2011) Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. PLoS Comput Biol 7: e1001048.
13. Kim S, Shephard N, Chib S (1998) Stochastic volatility: Likelihood inference and comparison with ARCH models. The Review of Economic Studies 65: 361–393.
14. Courville AC, Daw ND, Touretzky DS (2006) Bayesian theories of conditioning in a changing world. Trends in Cognitive Sciences 10: 294–300.
15. Daunizeau J, den Ouden HEM, Pessiglione M, Kiebel SJ, Friston K, et al. (2010) Observing the observer (II): deciding when to decide. PLoS ONE 5: e15555.
16. Nassar MR, Wilson RC, Heasly B, Gold JI (2010) An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment. Journal of Neuroscience 30: 12366–12378.
17. Yu AJ, Cohen JD (2009) Sequential effects: Superstition or rational behavior? In: Volume 21, Advances in Neural Information Processing Systems. pp. 1873–1880.
18. Steyvers M, Brown S (2006) Prediction and change detection. In: Volume 18, Advances in Neural Information Processing Systems. pp. 1281–1288.
19. Frijda NH, Manstead ASR, Bem S (2000) The influence of emotions on beliefs. In: Frijda NH, Manstead ASR, Bem S, editors. Emotions and Beliefs: How Feelings Influence Thoughts, New York: Cambridge University Press. pp. 1–9.
20. Colombetti G (2005) Appraising valence. Journal of Consciousness Studies 12: 103–126.
21. Charland LC (2005) The heat of emotion: Valence and the demarcation problem. Journal of consciousness studies 12: 810.
22. Russell JA (2003) Core affect and the psychological construction of emotion. Psychological Review 110: 145–172.
23. Schwarz N, Clore GL (2007) Feelings and phenomenal experiences. In: Higgins ET, Kruglanski AW, editors. Social Psychology: Handbook of Basic Principles, New York: The Guilford Press. pp. 385–407.
24. MacKay DJC (2003) Information Theory, Inference and Learning Algorithms. Cambridge: Cambridge University Press.
25. Rescorla RA, Solomon RL (1967) Two-process learning theory: Relationships between pavlovian conditioning and instrumental learning. Psychological review 74: 151–182.

26. Sutton RS, Barto AG (1998) Reinforcement Learning: An Introduction. Cambridge: MIT Press.
27. Bentham J (1907) An introduction to the principles of morals and legislation. Oxford: Clarendon Press.
28. Davis W (1981) A theory of happiness. American Philosophical Quarterly 18: 111–120.
29. Gordon RM (1990) The Structure of Emotions: Investigations in Cognitive Philosophy. New York: Cambridge University Press.
30. Green OH (1992) The emotions: a philosophical theory. Dordrecht: Kluwer Academic Publishers.
31. Reisenzein R (2009) Emotional experience in the computational belief-desire theory of emotion. Emotion Review 1: 214–222.
32. Textor M (2012) States of a_airs. In: Zalta EN, editor, The Stanford Encyclopedia of Philosophy. Summer 2012 edition. Available: http://plato.stanford.edu/archives/sum2012/entries/states-of-affairs/.
33. Carver CS, Scheier MF (1990) Origins and functions of positive and negative affect: A control-process view. Psychological Review 97: 19–35.
34. Hsee CK, Abelson RP (1991) Velocity relation: Satisfaction as a function of the first derivative of outcome over time. Journal of Personality and Social Psychology 60: 341–347.
35. Batson CD, Shaw LL, Oleson KC (1992) Differentiating affect, mood, and emotion: Toward functionally based conceptual distinctions. In: No. 13, Emotion. Review of personality and social psychology. Thousand Oaks (California): Sage Publications, Inc. pp. 294–326.
36. Ben-Ze'ev A (2001) The Subtlety of Emotions. Cambridge: MIT Press.
37. Macknik SL, Martinez-Conde S (2009) Real magic: future studies of magic should be grounded in neuroscience. Nature Reviews Neuroscience 10: 241–241.
38. Forabosco G (2008) Is the concept of incongruity still a useful construct for the advancement of humor research? Lodz Papers in Pragmatics 4: 45–62.
39. Barrett LF, Russell JA (1999) The structure of current affect: Controversies and emerging consensus. Current Directions in Psychological Science 8: 10–14.
40. Ekman P (1999) Basic emotions. In: Dalgleish T, Power MJ, editors, Handbook of Cognition and Emotion. Chichester: John Wiley & Sons Ltd. pp. 45–60.
41. Ortony A, Clore GL, Collins A (1990) The Cognitive Structure of Emotions. Cambridge: Cambridge University Press.
42. Lazarus RS (1994) Emotion and adaptation. New York: Oxford University Press.
43. Higgins ET (2007) Value. In: Kruglanski AW, Higgins ET, editors, Social Psychology: Handbook of Basic Principles. 2nd edition. New York: The Guilford Press. pp. 454–472.
44. Keramati M, Gutkin BS (2011) A reinforcement learning theory for homeostatic regulation. In: Shawe-Taylor J, Zemel RS, Bartlett P, Pereira FCN, Weinberger KQ, editors, Advances in Neural Information Processing Systems 24. p. 8290.
45. Seth AK, Suzuki K, Critchley HD (2012) An interoceptive predictive coding model of conscious presence. Frontiers in Consciousness Research 2: 395.
46. American Psychiatric Association (2000) Diagnostic and Statistical Manual of Mental Disorders DSM-IV-TR. 4th edition. Arlington: American Psychiatric Association.
47. Foland-Ross LC, Gotlib IH (2012) Cognitive and neural aspects of information processing in major depressive disorder: an integrative perspective. Frontiers in psychology 3: 489.
48. Friston K (2012) Policies and priors. In: Gutkin B, Ahmed SH, editors, Computational Neuroscience of Drug Addiction. New York: Springer. pp. 237–283.
49. Matthews GR, Antes JR (1992) Visual attention and depression: Cognitive biases in the eye fixations of the dysphoric and the nondepressed. Cognitive Therapy and Research 16: 359–371.
50. Eizenman M, Yu LH, Grupp L, Eizenman E, Ellenbogen M, et al. (2003) A naturalistic visual scanning approach to assess selective attention in major depressive disorder. Psychiatry research 118: 117–128.
51. Caseras X, Garner M, Bradley BP, Mogg K (2007) Biases in visual orienting to negative and positive scenes in dysphoria: An eye movement study. Journal of abnormal psychology 116: 491–497.
52. Kellough JL, Beevers CG, Ellis AJ, Wells TT (2008) Time course of selective attention in clinically depressed young adults: an eye tracking study. Behaviour research and therapy 46: 1238–1243.
53. Mller-Oerlinghausen B, Berghfer A, Bauer M (2002) Bipolar disorder. The Lancet 359: 241–247.
54. American Psychiatric Association (2002) Practice guideline for the treatment of patients with bipolar disorder (revision). The American journal of psychiatry 159: 1–50.
55. Daunizeau J, Friston K, Kiebel S (2009) Variational bayesian identification and prediction of stochastic nonlinear dynamic causal models. Physica D: Nonlinear Phenomena 238: 2089–2118.
56. Kringelbach ML (2005) The human orbitofrontal cortex: linking reward to hedonic experience. Nature reviews Neuroscience 6: 691–702.
57. Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A (2008) Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. The Journal of Neuroscience 28: 5623–5630.
58. Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275: 1593–1599.
59. Zink CF, Pagnoni G, Chappelow J, Martin-Skurski M, Berns GS (2006) Human striatal activation reflects degree of stimulus saliency. NeuroImage 29: 977–983.
60. Friston KJ, Shiner T, FitzGerald T, Galea JM, Adams R, et al. (2012) Dopamine, affordance and active inference. PLoS Comput Biol 8: e1002327.
61. Dayan P, Yu AJ (2006) Phasic norepinephrine: a neural interrupt signal for unexpected events. Network (Bristol, England) 17: 335–350.