

Optimal inference with suboptimal models: Addiction and active Bayesian inference



Philipp Schwartenbeck^{a,b,c,*}, Thomas H.B. FitzGerald^a, Christoph Mathys^a, Ray Dolan^a, Friedrich Wurst^d, Martin Kronbichler^{b,c}, Karl Friston^a

^a The Wellcome Trust Centre for Neuroimaging, UCL, 12 Queen Square, London WC1N 3BG, UK

^b Institute for Cognitive Neuroscience, University of Salzburg, Salzburg, Austria

^c Neuroscience Institute, Christian-Doppler-Klinik, Paracelsus Medical University Salzburg, Salzburg, Austria

^d Department of Psychiatry and Psychotherapy II, Christian-Doppler Hospital, Paracelsus Medical University, Salzburg, Austria

ARTICLE INFO

Article history:

Received 1 September 2014

Accepted 8 December 2014

ABSTRACT

When casting behaviour as active (Bayesian) inference, optimal inference is defined with respect to an agent's beliefs – based on its generative model of the world. This contrasts with normative accounts of choice behaviour, in which optimal actions are considered in relation to the true structure of the environment – as opposed to the agent's beliefs about worldly states (or the task). This distinction shifts an understanding of suboptimal or pathological behaviour away from aberrant inference as such, to understanding the prior beliefs of a subject that cause them to behave less 'optimally' than *our* prior beliefs suggest they should behave. Put simply, suboptimal or pathological behaviour does not speak against understanding behaviour in terms of (Bayes optimal) inference, but rather calls for a more refined understanding of the subject's generative model upon which their (optimal) Bayesian inference is based. Here, we discuss this fundamental distinction and its implications for understanding optimality, bounded rationality and pathological (choice) behaviour. We illustrate our argument using addictive choice behaviour in a recently described 'limited offer' task. Our simulations of pathological choices and addictive behaviour also generate some clear hypotheses, which we hope to pursue in ongoing empirical work.

© 2014 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Introduction

Optimal inference or optimal models?

Recent theoretical attempts to provide a unified understanding of behaviour and brain function in a dynamic and uncertain world cast cognition as probabilistic Bayesian inference [4,47,48,76], culminating in the Bayesian brain hypothesis ([27,36,37,54,65]). In this view, agents represent hidden states of the world as probability distributions and update their beliefs according to Bayes' rule. Importantly, Bayesian inference is based on an agent's generative model of the world, which is used to infer a probability distribution over causes (latent variables) generating sensory input. Formally, a generative model is a joint distribution $p(y, \theta)$ over sensory input (observations) y and subject-specific parameters θ , based on a likelihood function $p(y|\theta)$ and a prior distribution over parameters $p(\theta)$ [72]. Agents can then invert this model to find parameters that maximise the evidence of a given model in data, and thus infer the

causal structure of their environment. Put simply, in order to make sense out the world, it is necessary to entertain a (generative) model of the causal structure that underlies noisy and constantly changing sensory impressions [23]. Therefore, an agent's generative model crucially influences the result of an inference process.

Recent evidence, especially in the field of perception, suggests that the brain may indeed perform (close to) 'optimal' probabilistic inference [18,48,55,82,83]. In other instances, it has been shown that organisms substantially deviate from Bayes-optimality defined with respect to the true generative process [15,16,19,46,51], leading some to question whether humans truly perform Bayesian inference and suggesting that in many situations non-optimal models are employed; for example, simple heuristics when making decisions [17,43,52,77,78]. This question becomes particularly important when assessing individual differences in behaviour or characteristic aberrancies in psychopathology. For instance, how could a model based on 'optimal' Bayesian inference be used to account for characteristic suboptimal behaviour of a specific patient group?

Here, we propose two potential causes of suboptimal behaviour, which have radically different implications for understanding individual differences and psychiatric disease; namely, failures in

* Corresponding author at: Institute for Cognitive Neuroscience, University of Salzburg, Hellbrunnerstraße 34, 5020 Salzburg, Austria.

E-mail address: philipp.schwartenbeck@sbg.ac.at (P. Schwartenbeck).

inference and failures in the models upon which (optimal) inference is based (cf., [5]). We argue that patients showing pathological behaviour can still perform 'optimal' inference based on a particular generative model of the environment, which causes their behaviour to be less well adapted compared to control subjects. Importantly, the inference machinery itself is not broken, but instead a maladaptively constituted generative model causes pathological behaviour – based on (optimal) Bayesian inference. This implies a focus not on what patients perform in place of Bayesian inference but rather how their particular generative model – upon which inference is based – is constituted. Put simply, if the brain truly is a hierarchical Bayesian prediction machine, one can explain both normal and pathological behaviour. Identifying specific parameters or the specific constitution of the generative model in a patient group would have substantial impact on improving diagnosis, prediction of relapse or therapeutic success in clinical practise. Note that this marks a crucial advantage over approaches that focus on the deviation from optimality: investigating the subject-specific generative model of a task – as opposed to noting a subject's deviation from 'optimal' behaviour – allows us (in principle) to identify the origins of aberrant inference in the cerebral hierarchy. In addition, this approach allows for a more nuanced perspective on the causes underlying observed behaviour. It is possible that some psychiatric conditions share characteristic aberrancies in expressed behaviour but that their particular causes are different. While noting the characteristic deviation from 'optimality' does not enable us to differentiate between different mechanisms that might cause this deviation, investigating the generative model underlying observed behaviour allows us to differentiate possible mechanisms that induce similar behaviour.

A prominent example, which we discuss in more detail below, is overly habitual and impoverished goal-directed behaviour both in addiction and obsessive compulsive disorders, which is likely to be induced by high impulsivity (low confidence) and high anxiety, respectively. It is important to note, however, that a computational approach to psychiatry rests on the assumption that some systematic aspects of the models used by a particular patient group can be identified, which may not always be guaranteed in empirical work.

In summary, failing to account for the subject-specific generative models, means that there is only one way of acting optimally (e.g., by maximising the monetary gain in an economic decision-making task) and the only question one can pose is to what extent individuals differ from optimal behaviour as defined *a priori*. This is likely to create the impression that agents substantially deviate from Bayes-optimal inference and show significant inter-individual deviations. If one acknowledges individual variation in representing the environment, however, a second definition of optimality arises, which is relative to an agent's generative model. In this framework, suboptimal or even pathologic behaviour is not understood as broken inference as such but as (optimal) Bayesian inference based on a suboptimal generative model of the world. Put differently, Bayesian inference based on individual models of the world will look suboptimal in a manner proportional to the differences between these models and our beliefs about the true structure of the world.

In the following, we discuss how inter-individual differences and psychiatric conditions can be understood in the context of Bayesian inference. To illustrate our argument, we consider a recently proposed Bayesian model of decision-making that casts choice behaviour as pure inference. We will briefly review the behavioural and physiological validity of this model and then discuss its application to translational work. Specifically, we will show how active inference can be used to characterise choice behaviour in patient cohorts (e.g., chronic alcohol abusers and pathological gamblers) using simulations.

Normative or process models

Under the Bayesian brain hypothesis, a subject's generative model characterises how a cognitive process can be performed by an embodied agent. In computational neuroscience one tries to understand how a certain problem could be solved in the brain to motivate a generative model of a task that might be implemented in the brain in a biologically plausible way.

Normative approaches try to understand cognition and behaviour with respect to optimising a particular quantity, such as maximising reward [74] or minimising surprise [38]. In such approaches, behaviour is governed by an overall principle that compels agents to optimise a specific quantity. Within such a general framework it is possible to formulate models with varying degrees of biological plausibility. For instance, it makes a substantial difference whether one simply describes an optimal solution to a task, such as a solution in a decision-making paradigm based on backwards induction and dynamic programming, or whether one takes into account the actual cognitive resources and neuronal message-passing necessary to perform a process given the computational resources of the brain [40]. In other words, normative approaches provide a state theory (about what the brain does), while biologically plausible implementations of the state theory furnish a process theory (about how it is done).

Characterising generative models allows one to formulate a process model of how cognitive processes are implemented in the brain, in contrast to models that rest on a correlational approach or are agnostic about the relationship between the mechanics of a model and brain functioning. The normative approach provides tools for assessing the extent of deviation from 'optimal' behaviour, given a model of the world that individuals or cohorts display. If normative theories are equipped with process models, they can provide direct insight into (maladaptive) computations that underlie pathology. Crucially, the question about the optimal solution to a task, given the world, becomes irrelevant because the real interest lies in the (optimal) solution, given an agent's model of the world. 'Bayes optimality' can be conceived of as a principled way of performing optimally in a particular task but it can also be conceived of as optimal inference under an agent's generative model of the world. When adopting the latter perspective, Bayesian inference can provide a process model of how agents actually behave – instead of a normative model of how they should behave.

We will see later that the most likely candidates for subject-specific differences in generative models are their prior beliefs. This is important because there are mathematical (complete class) theorems which suggest that any behaviour which can be described by a loss function can always be described as Bayes optimal under some prior beliefs. This means, that (at least mathematically) any 'broken' Bayesian inference, can always be characterised in terms of aberrant prior beliefs. This paper exploits this perspective to suggest that many neuropsychiatric conditions can be understood in terms of false inference under aberrant prior beliefs (including beliefs about beliefs, such as the confidence in – or precision of – one's beliefs about the world).

Individual differences and psychopathology in Bayesian inference

In cognitive neuroscience, we often try to understand inter-individual variability in cognitive processes and brain activity. Formally, this means that we (as observers) invert generative models of our subjects based on their observed responses (cf., [24]). This allows us to ask which parameters may have caused a distinct type of behaviour and enables us to investigate the relationship between specific parameters and personality traits or characteristics in behaviour, such as impulsivity or sensation seeking. This

is the central procedure in the field of computational psychiatry [29,39,41,60,72].

There are several reasons why organisms might entertain generative models that ‘suboptimally’ approximate the true causal structure of the world. One reason is the accuracy-complexity trade-off inherent in Bayesian inference, as studied in machine learning [11,56]. To prevent overfitting of data, models have to optimise the trade-off between accurate predictions and generalizability, which means that good models have to restrict themselves from treating noise as meaningful information. Note that this is closely related to the concept of bounded rationality and satisficing as opposed to maximising [70,71]. Furthermore, specific circumstances may pose time pressure on behaviour, such as in classical fight or flight decisions. The use of parsimonious models may then become necessary because building a perfect model of the environment simply takes too long – or cannot accommodate slight changes in the environment. There could also be uncertainty about which model is best, inducing the need for Bayesian model averaging [35], or incompatible sensory data and priors; for example, in the case of perceptual illusions [42,53,79]. Also, non-ecological environments in experiments may lead to the adoption of inappropriate models since these may be the only models available to the agent based on her past experience. Most importantly, however, there may be phylogenetic constraints on the type of models that can be implemented biologically – as well as differences related to an individual’s ontogenetic development. The latter is of crucial importance for investigating pathologic behaviour – as it is possible that specific experiences in (biological) neurodevelopment may constrain subsequent models of the world, inducing, in extreme cases, a characteristic psychiatric condition. From a mathematical perspective, these considerations speak to approximate Bayesian inference in which the implementational problems of exact Bayesian inference are finessed using assumed forms for posterior beliefs. The particular choice of these forms speaks to the nature of the neuronal code and neuronal message passing within the brain. We will focus on a particular but ubiquitous form of approximate Bayesian inference later; namely, variational Bayesian updating through the minimisation of variational free energy – that provides a biologically plausible process theory.

In summary, characterising and individual’s generative model may be crucial when investigating group differences among different populations. Indeed, identifying particular characteristics of generative models associated with a particular psychiatric disorder marks a central aim of computational psychiatry [2,58,60,72]. An aim of computational psychiatry is to improve the diagnosis of psychiatric conditions and assess therapeutic progress. Assessing the characteristics of a generative model in a patient group allows one to investigate specific, trait-like characteristics of certain model parameters that may underlie pathological behaviour [72]. In doing so, researchers move from a purely descriptive account to understanding and explaining computational phenotypes that cause maladaptive conditions. Prominent examples for such approaches have been proposed in the field of autism [62], schizophrenia [1] or functional motor and sensory symptoms [31].

Decision-making as active inference

As an illustration, we now present a model of choice behaviour in the framework of active Bayesian inference, and describe how abnormalities in this model might account for pathological decision-making in addiction. Crucially, to understand individual variability in choice, we estimated individual parameters of the generative model (of the task) based on observed behaviour, allowing us to identify specific parameters that may serve as biomarkers of behavioural and neuronal characteristics in substance-based and behavioural addiction, relative to a control group.

Recently, a generic model of decision processes has been proposed in which Markov Decision Problems (MDPs) are solved using (active) Bayesian inference based on variational (discrete-time) updates of prior beliefs about current states, action and control [40]. Crucially, agents try to minimise surprise about future states, where surprise is based on a belief-distribution about future states, in which desired states are more likely to be attained and undesired states the converse. Agents then infer an optimal policy by minimising the relative entropy or Kullback–Leibler-Divergence between their belief distribution and a distribution of likely outcomes. Intuitively, they increase the likelihood of visiting preferred states or goals by selecting policies that produce outcomes that are as close as possible to desired outcomes. Effectively, preferred outcomes are ensured by minimising surprise, where goals constitute the least surprising outcome. Crucially, not only beliefs about policies but also the precision (certainty) with which these beliefs are held have to be inferred. Precision can be understood as the confidence that a selected policy will be successful and thus controls the stochasticity or goal-directedness of an agent’s behaviour. Formally, precision is similar to the inverse temperature parameter in classical logistic or softmax choice rules [25,74]: however, there is a crucial difference – precision has a Bayes-optimal solution and has to be inferred and updated at each time-step. The resulting behaviour depends on a generative model of task contingencies and prior beliefs about precision – and therefore allows for an individual assessment of these model parameters and their relationship with behavioural tendencies, as illustrated in the following example.

We performed a behavioural and neuronal validation of this generic model in a functional magnetic resonance imaging (fMRI) experiment [69]. For this study, we developed a novel task in which participants had to infer the optimal time to accept an offer. In brief, subjects were presented with a small initial offer and had to decide how long to wait for a higher offer, with the risk of losing the initial offer and winning nothing (Fig. 1A). If they accepted too early, they precluded a higher offer at a later stage of the game, whereas if they accepted too late they risked losing the initial offer and winning nothing. To emphasise the context-sensitivity of this task, we defined the probabilities of receiving the high offer and losing the low offer as hazard rates, such that it became progressively less likely to receive the high offer and more likely to lose the initial offer (Fig. 1B). We compared long (8 trials) and short (4 trials) and varied the amount of the initial offer (between 9 and 35 pence, whereas the high offer was always 80 pence). In summary, subjects had to infer the optimal time to accept the initial offer taking the specific task contingencies into account, without waiting too long or not long enough.

We performed maximum a posteriori (MAP) modelling of task-relevant parameters (expected precision, hazard rate and sensitivity) by inverting the generative model of each subject, given observed behaviour. These parameters determined how a specific type of behaviour may have been caused. In brief, prior precision reflects the prior belief about the relative probability of competing policies or a “trait-confidence”. This means subjects with a higher prior precision are more confident about receiving the high offer if they deliberately wait. Estimates of expectations of hazard rates tell us how well subjects represented the task contingencies, whereas the sensitivity for the difference in utilities (log likelihood) of outcomes provides important information about how much subjects valued one outcome over others (see Fig. 2 for observed behaviour and model predictions).

Therefore, we were able to characterise the individual generative model (representation of the task) employed by each subject, upon which variational Bayesian updates were based. For instance, we found evidence that subjects who had higher prior precision accepted the initial offer less frequently and waited longer for the high offer. Furthermore, we were able to show that the updates




Fig. 1. Task for model validation. (A) Each game comprised a specified number of trials (discrete time-steps). On each trial, subjects had to decide whether to accept a current offer or wait for the next trial. If subjects decided to wait, the low offer could be withdrawn with probability r , it could be replaced by a high offer with probability q or could be carried over to the next trial with probability $1 - (r + q)$. (B) The probabilities were defined by hazard rates, such that withdrawal probability increased, whereas a high offer became less likely the longer subjects waited. Adapted from [69].

of expected precision – based on the individual estimates of prior precision – were encoded by dopaminergic midbrain regions.

Having established the validity of this (active Bayesian inference) model of subject responses in behavioural and neuronal terms, we now ask whether it can be used to understand and classify (pathologic) choice behaviour in addiction. Specifically, we can now look in detail at how changes in a subject's prior beliefs will produce characteristic patterns of behaviour by simulating choice behaviour in the limited offer game, under different prior expectations.

Hypothesis and evaluation: addictive choice behaviour as optimal inference

In the following, we discuss characteristic abnormalities of behaviour in addiction and how these could be accounted for by an active inference model of choice behaviour. We will use simulations to illustrate maladaptive choice behaviour in our limited offer paradigm in both a qualitative and quantitative sense. These simulations generate some specific predictions, which we are currently testing empirically – using both behavioural and physiological responses.

Addiction is associated with characteristic patterns of maladaptive choice behaviour [49,61,63], and thus appears to be ideally suited for a computational approach to understand the basis of aberrant decision processes. Individuals with addiction have been found to be more impulsive [3,6,32], and neuronal activation associated with impulsivity or response inhibition can be used to predict future substance abuse [57]. Furthermore, addictive choice behaviour has been related to steeper delay discounting [59,64,66,73] – with the implication that the negative future consequences of addictive behaviour are strongly devalued – as well as to reduced risk sensitivity and decreased inhibitory control [21,20,50].

From a computational perspective, it has been argued that addictive behaviour may be caused by a shift from action–outcome to stimulus–response control [29], resulting in a predominance of less flexible, habitual modes of behaviour [32,45,67]. This has also been framed as the gradual transition from ‘impulsion to compulsion’, implying that addictive behaviour is marked initially by seeking hedonic pleasure from a certain substance intake or behaviour, whereas later this addictive behaviour becomes habitual and automatic [33]. Other perspectives invoke a shift from a predominance of liking to wanting in the context of incentive salience [8,9], implying that addicts show compulsive ‘wanting’ of drugs or drug-like behaviour without actually ‘liking’ it [68].

In addiction research there are several hypotheses that explain addictive behaviour. The impulsivity hypothesis states that individuals with addiction are more impulsive and worse at inhibiting a response than controls, which – combined with an increased sensitivity to reward – may induce addictive behaviour [3,10]. A reward-deficiency hypothesis, on the other hand, suggests that addicted individuals are less sensitive to nondrug rewards relative to drug-related rewards, which compels them to seek drugs [12]. Imaging studies provide evidence for both accounts [50]. Specifically, addictive behaviour has been associated with changes in reward processing [59,64] that is often associated with a dysfunction of the dopaminergic system [7,28,30,75] – for which there is also genetic evidence [13,14,22]. Most prominently, addiction is marked by a loss of phasic response of the dopaminergic system [80,81], which resonates with the reward-deficiency hypothesis – as this implies that only drug-related behaviour can elicit phasic dopaminergic responses and thus the feeling of reward.

In summary, addiction appears to be strongly related to impulsivity and a lack of inhibitory control, a preponderance of habitual as opposed to goal-directed or planned behaviour and a hypersensitivity to rewarding, drug-related stimuli. Given the large body of empirical evidence for these characteristics of addictive (choice) behaviour – what can we gain from an active inference perspective on addiction? As discussed earlier, the important question is whether specific parameters (in subject-specific generative models) can be identified that explains group differences in addictive compared to normal choice behaviour: in other words, what sorts of generative model (or model parameters) produce (Bayes optimal) addictive behaviour? Fig. 3 provides an intuitive depiction of the proposed variational message passing scheme in active inference upon which our considerations are based.

Precision reflects the confidence about reaching a desired goal; i.e., the successful implementation of a policy. In active inference, this means that precision reflects the degree of stochasticity or goal-directedness of behaviour. In the context of the limited offer task, precision reflects confidence about receiving the high offer. Thus, subjects with a higher prior precision tend to wait longer for the high offer and accept the initial offer less often – which can be understood as being more confident that the high offer will eventually transpire. Conversely, subjects with a low precision will display more variability in their behaviour and are less confident about waiting for the high offer. Consequently, they will be more impulsive in accepting the initial offer. It is important to appreciate that precision can account for this type of ‘reflection impulsivity’




Fig. 2. Results. Observed and estimated acceptance probabilities as predicted by our model (note that $P_{\text{wait}} = 1 - P_{\text{accept}}$). These estimates are based upon maximum a posteriori values for prior precision, the hazard rate and sensitivity to differences in monetary cues.




Fig. 3. Proposed variational message passing scheme. This figure illustrates the proposed message passing scheme in active inference and its putative mapping to brain function. The scheme rests on variational updates, such that several distinct representations (expectations of states, policies and precision) are updated and inform each other – providing a nice metaphor for functional segregation and integration in the brain. Here, current observations inform inference about hidden states of the environment, which induce updates of precision. Precision influences policy selection which in turn specifies actions – that are sampled from the most likely policies. Note that the connections are recursive because variation updates require a reciprocal exchange of sufficient statistics. Recent empirical work has shown that expected precision is encoded in the dopaminergic midbrain [69].

(i.e., a tendency to respond spontaneously with little deliberation) but not for higher risk-taking: if precision increases, subjects will be very sensitive to differences in the values of policies (e.g., waiting until a later trial before accepting the initial offer) and act

accordingly. In other words, they will not deviate much from their favoured policy over successive games. This means that increasing precision will not compel subjects to wait too long and become too risk-seeking, it will just make them more sensitive to the

advantages of planned and goal-directed behaviour. Conversely, decreasing precision will make subjects less confident in receiving the high offer and thus less goal-directed. This represents a state of poor self-control, causing the subject to display habitual behaviour. Subjects, therefore, will accept the initial offer too early because they cannot postpone their response – and therefore show a lack of inhibitory control (impulsive behaviour). Fig. 4 shows the effects of prior precision on behaviour when the offer is withdrawn or accepted as well as the propensity to accept the initial offer (safe option) simulated in 256 games for different prior expectations about precision.

A possible mechanism underlying impulsivity – and a lack of inhibitory control – could therefore be lower prior precision, reflecting more stochastic (inconsistent) behaviour and less confidence in the successful implementation of long-term policies (and therefore less goal-directedness). This provides an explanation for why some subjects accept earlier than would be optimal for gain maximisation in this task [69]. Note that Schwartenbeck et al. [69] were able to show that precision, as defined here, is encoded in the dopaminergic midbrain, which fits comfortably with the association between impulsive behaviour and dysfunction of the dopaminergic system in addiction.

However, we also know that higher risk-taking – often associated with steeper discounting of future rewards or events – is a key characteristic of addictive decision-making. In the context of our limited offer task, being too risk-seeking would imply that subjects wait too long for the high offer – and thus risk losing the initial offer. How could this be represented in a subject's generative model? As discussed earlier, subjects also have to represent the hazard rate of a withdrawal and the receipt of a high offer. Being too optimistic about winning the high offer – without considering the danger of losing the initial offer – could be explained by a poor expectation of the hazard rate, such that one assumes that receiving the high offer is very likely whereas losing the initial offer is unlikely. In contrast to precision, the individual representation of the hazard rate can account for higher risk-taking or (equivalently) the overconfidence in winning regardless of the risks, a common phenomenon in pathological gambling [21,20]. Fig. 5 illustrates the effects of adopting different hazard rates.

Note that a suboptimal representation of the hazard rate also explain the converse; namely, accepting prematurely because one estimates that a withdrawal is very likely whereas receiving the high offer is not likely at all. However, this form of being overly cautious is not consistent with previous findings of choice




Fig. 4. Effects of precision on behaviour. (A–C) reflect low ($\alpha = 4$), medium ($\alpha = 8$) and high ($\alpha = 16$) prior precision, respectively. The left column shows the inferred hidden state and precision in an example game with initial offer withdrawal at trial three. In each panel, the top row shows the inferred hidden state the agent believes this is in (upper left) and the expected probability of accepting or declining as a function of trials (upper right panel: dotted lines for earlier trials and the solid lines for the last trial). These expectations are updated using expected precision that is shown as a function of trial in the lower row – in terms of the precision at the end of each trial (lower left) and simulated dopamine responses reflecting the changes in precision (lower right). See [69] for further details about these (variational) updates. Here, the initial propensity to wait (stay) depends on prior precision and increases for higher values of precision (upper right panel). Intuitively, the drop of expected precision when the initial offer is withdrawn is larger with higher prior precision. The middle column shows the equivalent updates when a high offer is presented in the fifth trial. Expected precision displays a sharp increase when the high offer is received, where the size of the increase again depends on individual prior precision. The right column shows the observed acceptance probabilities at each time step, simulated in 256 games. It becomes obvious that the acceptance latencies are shifted towards later trials with increasing precision (confidence that the high offer will be received).




Fig. 5. Effects of estimates of the hazard rate on behaviour. (A–C) reflect high (losing the initial offer is likely), medium and low (losing the initial offer is unlikely) estimates of the hazard rate, respectively. This figure uses the same format as Fig. 4: the left and middle column show expectations about hidden states, behaviour and expected precision for a withdrawal of the initial offer (at the third trial) and the receipt of a high offer at the fifth trial, respectively. This time, the behaviour and initial level of precision is identical, but the expected behaviour reflects the differences in hazard rates – with the probability of waiting being smaller if losing the initial offer is likely. The right column predicts that subjects will accept the initial offer early if a loss is thought to be likely.

behaviour in addition, which is marked by higher impulsivity, less goal-directed behaviour and higher risk-taking. One might therefore hypothesise that people suffering from addiction will exhibit lower prior precision, causing behaviour to be more impulsive and stochastic, combined with a lower prior expectation about the hazard rate, inducing higher risk-taking and lower risk-sensitivity, as illustrated in Fig. 6. By investigating the prior expectations in a group of substance-based addiction, behavioural addiction and a control group one should be able to characterise general differences and commonalities on an individual and group level – and associate these differences in particular parameters with differences in observed behaviour and neuronal activation.

Discussion

We have argued that in order to understand inter-individual variability – and pathological behaviour – it is necessary to move from a descriptive or normative account of optimality to a definition of optimal inference, based on an individual's generative model of the world. People are now starting to appreciate that casting atypical behaviour (e.g., planning) as inference does not necessarily imply apparent flaws in an inference process, but calls for a deeper understanding of biologically plausible approximations or bounds on probabilistic inference [2,26,65,72]. To illustrate this point we have used a generic model of choice




Fig. 6. Effects of a low prior precision and a low estimate of the hazard rate on behaviour. Low precision causes subjects to be more stochastic and less patient in their behaviour, whereas a low estimate of the hazard rate causes subjects to become more risk-seeking in waiting for the high offer. See legends for Figs. 4 and 5 for details.

behaviour, which provides a formal characterisation of decision processes in terms of prior expectations causing observed behaviour. We have tried to substantiate this approach by using it to simulate a potential behavioural phenotype of addiction in a 'Limited offer task'.

These simulations also provide a nice illustration of how a computational account of pathologic behaviour can provide the basis for understanding similar behaviour that is caused by different mechanisms. As discussed earlier, obsessive compulsive disorders (OCD) have been associated with habitual and impoverished goal directed behaviour, but may be caused by high anxiety rather than high reflection impulsivity (i.e., a lack of confidence or precision) [33,34,44]. Both Figs. 4A and 5A show behaviour that is less goal-directed – and marked by a propensity to accept an initial offer in early trials – as one would expect to see in addiction and patients with OCD. In the former case, however, overly habitual behaviour is caused by reduced confidence or high impulsivity; i.e., the inability to decline a current offer (also marked by higher dispersion over acceptance latencies). In the latter simulation, the short acceptance latency is caused by pessimistic prior expectations about the hazard rate; in other words, assuming that a withdrawal is very likely. This could be induced by overly cautious or anxious (pessimistic) expectations, and thus might explain the predominance of habitual behaviour in OCD. In short, characterising the generative model underlying suboptimal behaviour provides a principled approach to understanding the origins of maladaptive behaviour as well as the diverse computational phenotypes that present similar 'symptoms'. This further illustrates the diagnostic value of a computational approach to psychiatry. Casting choice behaviour as (active) Bayesian inference may provide a useful tool to do so – and thus help improve our diagnostic assessment and evaluation of therapeutic success.

Note that our scheme allows us to identify likely causes that underlie pathologic behaviour – such as prior expectations that causes behavioural and neuronal responses in addiction. It neither makes strong predictions about the development of an addiction nor was it designed to be (*a priori*) sensitive to addictive choice behaviour. Rather, the framework provides a generic account of decision-making that allows one to assess the individual (generative) model that underlies (Bayes optimal) choice, and thus the individual characteristics that underlie addictive choice. Crucially, this approach posits specific mechanisms that may impair decision-making in addiction. This may be relevant for differentiating addiction from other conditions and may also – more practically – prove useful for assessing the current status of an addict (and their therapeutic progress) or for predicting relapse.

In conclusion, we believe that an active (Bayesian) inference framework may be useful for quantifying maladaptive choice behaviour in addiction, such as being too impulsive and risk seeking or showing a predominance of habitual behaviour when making choices. Having specified our hypotheses and theoretical considerations, we look forward to presenting empirical tests of these ideas in the near future.

Conflict of interest

The authors declare no conflict of interest.

Acknowledgements

Philipp Schwartenbeck is a recipient of a DOC Fellowship of the Austrian Academy of Sciences at the Institute for Cognitive Neuroscience, University of Salzburg. The Wellcome Trust Centre for Neuroimaging is supported by core funding from the Wellcome Trust 091593/Z/10/Z. Christoph Mathys is supported by a Joint

Initiative involving Max Planck Society and University College London on Computational Psychiatry and Aging Research.

References

- [1] Adams R, Perrinet LU, Friston K. Smooth pursuit and visual occlusion: active inference and oculomotor control in schizophrenia. *PLoS ONE* 2012;7(10):e47502. <http://dx.doi.org/10.1371/journal.pone.0047502>.
- [2] Adams R, Stephan KE, Brown HR, Frith CD, Friston KJ. The computational anatomy of psychosis. *Front. Psychiatry* 2013;4:47. <http://dx.doi.org/10.3389/fpsy.2013.00047>.
- [3] Bechara A. Decision making, impulse control and loss of willpower to resist drugs: a neurocognitive perspective. *Nat. Neurosci.* 2005;8(11):1458–63. <http://dx.doi.org/10.1038/nn1584>.
- [4] Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, et al. Probabilistic population codes for Bayesian decision making. *Neuron* 2008;60(6):1142–52. <http://dx.doi.org/10.1016/j.neuron.2008.09.021>.
- [5] Beck JM, Ma WJ, Pitkow X, Latham PE, Pouget A. Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 2012;74(1):30–9. <http://dx.doi.org/10.1016/j.neuron.2012.03.016>.
- [6] Belin D, Mar AC, Dalley JW, Robbins TW, Everitt BJ. High impulsivity predicts the switch to compulsive cocaine-taking. *Science* 2008;320(5881):1352–5. <http://dx.doi.org/10.1126/science.1158136>.
- [7] Berke JD, Hyman SE. Addiction, dopamine, and the molecular mechanisms of memory. *Neuron* 2000;25(3):515–32.
- [8] Berridge KC. The debate over dopamine's role in reward: the case for incentive salience. *Psychopharmacology (Berl)* 2007;191(3):391–431. <http://dx.doi.org/10.1007/s00213-006-0578-x>.
- [9] Berridge KC. From prediction error to incentive salience: mesolimbic computation of reward motivation. *Eur J Neurosci* 2012;35(7):1124–43. <http://dx.doi.org/10.1111/j.1460-9568.2012.07990.x>.
- [10] Bickel WK, Miller ML, Yi R, Kowal BP, Lindquist DM, Pitcock JA. Behavioral and neuroeconomics of drug addiction: competing neural systems and temporal discounting processes. *Drug Alcohol Depend* 2007;90:585–91. <http://dx.doi.org/10.1016/j.drugalcdep.2006.09.016>.
- [11] Bishop C. Pattern recognition and machine learning. Springer; 2006.
- [12] Blum K, Braverman ER, Holder JM, Lubar JF, Monastra VJ, Miller D, et al. Reward deficiency syndrome: a biogenetic model for the diagnosis and treatment of impulsive, addictive, and compulsive behaviors. *J Psychoact Drugs* 2000;32:1–112.
- [13] Blum K, Sheridan PJ, Wood RC, Braverman ER, Chen TJ, Comings DE. Dopamine D2 receptor gene variants: association and linkage studies in impulsive-addictive-compulsive behaviour. *Pharmacogenetics* 1995;5(3):121–41.
- [14] Blum K, Wood RC, Braverman ER, Chen TJ, Sheridan PJ. The D2 dopamine receptor gene as a predictor of compulsive disease: Bayes' theorem. *Funct Neurol* 1995;10(1):37–44.
- [15] Bowers JS, Davis CJ. More varieties of Bayesian theories, but no enlightenment. *Behav Brain Sci* 2011;34(04):193–4. <http://dx.doi.org/10.1017/S0140525X11000227>.
- [16] Bowers JS, Davis CJ. Bayesian just-so stories in psychology and neuroscience. *Psychol Bull* 2012;138(3):389–414. <http://dx.doi.org/10.1037/a0026450>.
- [17] Brighton H, Gigerenzer G. Bayesian brains and cognitive mechanisms: harmony or dissonance? In: Chater N, Oaksford M, editors. *The probabilistic mind: prospects for a Bayesian cognitive science*. New York: Oxford University Press; 2008. p. 189–208.
- [18] Cain MS, Vul E, Clark K, Mitroff SR. A bayesian optimal foraging model of human visual search. *Psychol Sci* 2012;23(9):1047–54. <http://dx.doi.org/10.1177/0956797612440460>.
- [19] Charness G, Levin D. When optimal choices feel wrong: a laboratory study of and affect Bayesian updating, complexity, and affect. *Am Econ Rev* 2014;95(4):1300–9.
- [20] Clark L. Decision-making during gambling: an integration of cognitive and psychobiological approaches. *Philos Trans R Soc London, Ser B* 2010;365(1538):319–30. <http://dx.doi.org/10.1098/rstb.2009.0147>.
- [21] Clark L, Limbrick-Oldfield EH. Disordered gambling: a behavioral addiction. *Curr Opin Neurobiol* 2013;23(4):655–9. <http://dx.doi.org/10.1016/j.conb.2013.01.004>.
- [22] Comings DE, Blum K. Reward deficiency syndrome: genetic aspects of behavioral disorders. *Prog Brain Res* 2000;126:325–41. [http://dx.doi.org/10.1016/S0079-6123\(00\)26022-6](http://dx.doi.org/10.1016/S0079-6123(00)26022-6).
- [23] Conant RC, Ashby RW. Every good regulator of a system must be a model of that system. *Int J Syst Sci* 1970;1(2):89–97. <http://dx.doi.org/10.1080/00207177008920220>.
- [24] Daunizeau J, den Ouden HEM, Pessiglione M, Kiebel SJ, Stephan KE, Friston KJ. Observing the observer (I): meta-bayesian models of learning and decision-making. *PLoS ONE* 2010;5(12):e15554. <http://dx.doi.org/10.1371/journal.pone.0015554>.
- [25] Daw ND, Doherty JPO, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. *Nature* 2009;441(7095):876–9. <http://dx.doi.org/10.1038/nature04766>.
- [26] Dayan P. Rationalizable irrationalities of choice. *Top Cognit Sci* 2014;1–25. <http://dx.doi.org/10.1111/tops.12082>.
- [27] Dayan P, Hinton GE, Neal RM, Zemel RS. The Helmholtz machine. *Neur Comp* 1995;7(5):889–904.

- [28] Diana M. The dopamine hypothesis of drug addiction and its potential therapeutic value. *Front Psychiatry* 2011;2:64. <http://dx.doi.org/10.3389/fpsy.2011.00064>.
- [29] Dolan RJ, Dayan P. Goals and habits in the brain. *Neuron* 2013;80(2):312–25. <http://dx.doi.org/10.1016/j.neuron.2013.09.007>.
- [30] Doyon WM, Dong Y, Ostroumov A, Thomas AM, Zhang T, Dani J. Nicotine decreases ethanol-induced dopamine signaling and increases self-administration via stress hormones. *Neuron* 2013;79(3):530–40. <http://dx.doi.org/10.1016/j.neuron.2013.06.006>.
- [31] Edwards M, Adams R, Brown H, Pareés I, Friston K. A Bayesian account of “hysteria”. *Brain* 2012;135(11):3495–512. <http://dx.doi.org/10.1093/brain/awt129>.
- [32] Everitt BJ, Belin D, Economidou D, Pelloux Y, Dalley JW, Robbins TW. Review. Neural mechanisms underlying the vulnerability to develop compulsive drug-seeking habits and addiction. *Philos Trans R Soc London Ser B* 2008;363(1507):3125–35. <http://dx.doi.org/10.1098/rstb.2008.0089>.
- [33] Everitt BJ, Robbins TW. Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nat Neurosci* 2005;8(11):1481–9. <http://dx.doi.org/10.1038/nn1579>.
- [34] Eysenck M, Derakshan N, Santos R, Calvo M. Anxiety and cognitive performance: attentional control theory. *Emotion* 2007;7(2):336–53. <http://dx.doi.org/10.1037/1528-3542.7.2.336>.
- [35] FitzGerald THB, Dolan R, Friston K. Model averaging, optimal inference and model averaging. *Front Hum Neurosci* 2014;8:457. <http://dx.doi.org/10.3389/fnhum.2014.00457>.
- [36] Friston K. A theory of cortical responses. *Philos Trans R Soc London Ser B* 2005;360(1456):815–36. <http://dx.doi.org/10.1098/rstb.2005.1622>.
- [37] Friston K. The free-energy principle: a rough guide to the brain? *Trends Cognit Sci* 2009;13(7):293–301. <http://dx.doi.org/10.1016/j.tics.2009.04.005>.
- [38] Friston K. A free energy principle for biological systems. *Entropy* 2012;14(11):2100–21. <http://dx.doi.org/10.3390/e14112100>.
- [39] Friston KJ, Dolan RJ. Computational and dynamic models in neuroimaging. *NeuroImage* 2010;52(3):752–65. <http://dx.doi.org/10.1016/j.neuroimage.2009.12.068>.
- [40] Friston K, Schwartenbeck P, FitzGerald T, Moutoussis M, Behrens T, Dolan RJ. The anatomy of choice: active inference and agency. *Front Hum Neurosci* 2013;7:598. <http://dx.doi.org/10.3389/fnhum.2013.00598>.
- [41] Friston K, Stephan KE, Montague R, Dolan RJ. Computational psychiatry: the brain as a phantastic organ. *Lancet Psychiatry* 2014;1(2):148–58. [http://dx.doi.org/10.1016/S2215-0366\(14\)70275-5](http://dx.doi.org/10.1016/S2215-0366(14)70275-5).
- [42] Geisler WS, Kersten D. Illusions, perception and Bayes. *Nat Neurosci* 2002;5(6):508–10. <http://dx.doi.org/10.1038/nn0602-508>.
- [43] Gigerenzer G, Gaissmaier W. Heuristic decision making. *Annu Rev Psychol* 2011;62:451–82. <http://dx.doi.org/10.1146/annurev-psych-120709-145346>.
- [44] Gillan CM, Pappmeyer M, Morein-Zamir S, Sahakian BJ, Fineberg NA, Robbins TW, et al. Disruption in the balance between goal-directed behavior and habit learning in obsessive-compulsive disorder. *Am J Psychiatry* 2011;168(7):718–26. <http://dx.doi.org/10.1176/appi.ajp.2011.10071062>.
- [45] Graybiel AM. Habits, rituals, and the evaluative brain. *Annu Rev Neurosci* 2008;31:359–87. <http://dx.doi.org/10.1146/annurev.neuro.29.051605.112851>.
- [46] Grether DM. Testing Bayes rule and the representativeness heuristic: some experimental evidence. *J Econ Behav Organiz* 1992;17(1):31–57. [http://dx.doi.org/10.1016/0167-2681\(92\)90078-P](http://dx.doi.org/10.1016/0167-2681(92)90078-P).
- [47] Griffiths TL, Tenenbaum JB. Optimal predictions in everyday cognition. *Psychol Sci* 2006;17(9):767–73. <http://dx.doi.org/10.1111/j.1467-9280.2006.01780.x>.
- [48] Griffiths TL, Tenenbaum JB. Predicting the future as Bayesian inference: people combine prior knowledge with observations when estimating duration and extent. *J Exp Psychol Gen* 2011;140(4):725–43. <http://dx.doi.org/10.1037/a0024899>.
- [49] Henden E, Melberg HO, Røgeberg OJ. Addiction: choice or compulsion? *Front Psychiatry* 2013;4:77. <http://dx.doi.org/10.3389/fpsy.2013.00077>.
- [50] Hommer D, Bjork JM, Gilman JM. Imaging brain response to reward in addictive disorders. *Ann NY Acad Sci* 2011;1216:50–61. <http://dx.doi.org/10.1111/j.1749-6632.2010.05898.x>.
- [51] Jones M, Love BC. Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behav Brain Sci* 2011;34(4):169–88. <http://dx.doi.org/10.1017/S0140525X10003134>.
- [52] Kahnemann D, Tversky A. Subjective probability: a judgement of representativeness. *Cognit Psychol* 1972;1(3):430–54.
- [53] Kersten D, Mamassian P, Yuille A. Object perception as Bayesian inference. *Annu Rev Psychol* 2004;55:271–304. <http://dx.doi.org/10.1146/annurev.psych.55.090902.142005>.
- [54] Knill DC, Pouget A. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci* 2004;27(12):712–9. <http://dx.doi.org/10.1016/j.tins.2004.10.007>.
- [55] Lee TS, Mumford D. Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am* 2003;20(7):1434. <http://dx.doi.org/10.1364/JOSA.20.001434>.
- [56] MacKay D. Information theory. Inference and Learning Algorithms: Cambridge University Press; 2003.
- [57] Mahmood OM, Goldenberg D, Thayer R, Miglioni R, Simmons N, Tapert SF. Adolescents' fMRI activation to a response inhibition task predicts future substance use. *Addict Behav* 2013;38(1):1435–41. <http://dx.doi.org/10.1016/j.addbeh.2012.07.012>.
- [58] Maia TV, Frank MJ. From reinforcement learning models to psychiatric and neurological disorders. *Nat Neurosci* 2011;14(2):154–62. <http://dx.doi.org/10.1038/nn.2723>.
- [59] Miedl SF, Peters J, Büchel C. Altered neural reward representations in pathological gamblers revealed by delay and probability discounting. *Arch Gen Psychiatry* 2012;69(2):177–86. <http://dx.doi.org/10.1001/archgenpsychiatry.2011.1552>.
- [60] Montague PR, Dolan RJ, Friston KJ, Dayan P. Computational psychiatry. *Trends Cognit Sci* 2012;16(1):72–80. <http://dx.doi.org/10.1016/j.tics.2011.11.018>.
- [61] Monterosso J, Piray P, Luo S. Neuroeconomics and the study of addiction. *Biol Psychiatry* 2012;72(2):107–12. <http://dx.doi.org/10.1016/j.biopsych.2012.03.012>.
- [62] Pellicano E, Burr D. When the world becomes “too real”: a Bayesian explanation of autistic perception. *Trends Cognit Sci* 2012;16(10):504–10. <http://dx.doi.org/10.1016/j.tics.2012.08.009>.
- [63] Peters J, Büchel C. The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends Cognit Sci* 2011;15(5):227–39. <http://dx.doi.org/10.1016/j.tics.2011.03.002>.
- [64] Peters J, Miedl SF, Büchel C. Formal comparison of dual-parameter temporal discounting models in controls and pathological gamblers. *PLoS ONE* 2012;7(11):e47225. <http://dx.doi.org/10.1371/journal.pone.0047225>.
- [65] Pouget A, Beck JM, Ma WJ, Latham PE. Probabilistic brains: knowns and unknowns. *Nat Neurosci* 2013;16(9):1170–8. <http://dx.doi.org/10.1038/nn.3495>.
- [66] Redish AD, Jensen S, Johnson A. A unified framework for addiction: vulnerabilities in the decision process. *Behav Brain Sci* 2008;31(4):415–37. <http://dx.doi.org/10.1017/S0140525X0800472X>.
- [67] Redish AD, Jensen S, Johnson A, Kurth-Nelson Z. Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychol Rev* 2007;114(3):784–805. <http://dx.doi.org/10.1037/0033-295X.114.3.784>.
- [68] Robinson TE, Berridge KC. The neural basis of drug craving: an incentive-sensitization theory of addiction. *Brain Res Rev* 1993;18(3):247–91.
- [69] Schwartenbeck P, FitzGerald THB, Mathys C, Dolan R, Friston K. The dopaminergic midbrain encodes the expected certainty about desired outcomes. *Cereb Cortex* 2014. <http://dx.doi.org/10.1093/cercor/bhu159>.
- [70] Simon HA. A behavioral model of rational choice. *Q J Econ* 1955;69(1):99. <http://dx.doi.org/10.2307/1884852>.
- [71] Simon HA. Bounded rationality and organizational learning. *Organiz Sci* 1991;2(1):125–34. <http://dx.doi.org/10.1287/orsc.2.1.125>.
- [72] Stephan KE, Mathys C. Computational approaches to psychiatry. *Curr Opin Neurobiol* 2014;25:85–92. <http://dx.doi.org/10.1016/j.conb.2013.12.007>.
- [73] Story GW, Vlaev I, Seymour B, Darzi A, Dolan RJ. Does temporal discounting explain unhealthy behavior? A systematic review and reinforcement learning perspective. *Front Behav Neurosci* 2014;8:76. <http://dx.doi.org/10.3389/fnbeh.2014.00076>.
- [74] Sutton R, Barto A. Reinforcement learning. An introduction. Cambridge, MA: MIT Press; 1998.
- [75] Taber KH, Black DN, Porrino LJ, Hurley RA. Neuroanatomy of dopamine: reward and addiction. *J Neuropsychiatry Clin Neurosci* 2012;24(1):1–4. <http://dx.doi.org/10.1176/appi.neuropsych.24.1.1>.
- [76] Tenenbaum JB, Griffiths TL. Generalization, similarity, and Bayesian inference. *Behav Brain Sci* 2001;24(4):629–40. discussion 652–791. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12048947>.
- [77] Tversky A, Kahneman D. Availability: a heuristic for judging frequency and probability. *Cognit Psychol* 1973;5:207–32.
- [78] Tversky A, Kahneman D. Judgment under uncertainty: Heuristics and Biases. *Science* (New York, N.Y.) 1974;185(4157):1124–31. <http://dx.doi.org/10.1126/science.185.4157.1124>.
- [79] Weiss Y, Simoncelli EP, Adelson EH. Motion illusions as optimal percepts. *Nat Neurosci* 2002;5(6):598–604. <http://dx.doi.org/10.1038/nn858>.
- [80] Whalley K. Addiction: dopamine dips as drug use increases. *Nat Rev Neurosci* 2014;15(6):349. <http://dx.doi.org/10.1038/nrn3752>.
- [81] Willuhn I, Burgeno LM, Groblewski PA, Phillips PEM. Excessive cocaine use results from decreased phasic dopamine signaling in the striatum. *Nat Neurosci* 2014;17(5):704–9. <http://dx.doi.org/10.1038/nn.3694>.
- [82] Wolpert DM. Probabilistic models in human sensorimotor control. *Hum Mov Sci* 2007;26(4):511–24. <http://dx.doi.org/10.1016/j.humov.2007.05.005>.
- [83] Yu AJ. Adaptive behavior: humans act as Bayesian learners. *Curr Biol* 2007;17(22):977–80. <http://dx.doi.org/10.1016/j.cub.2007.09.007>.