

Searching for an anchor in an unpredictable world: A computational model of obsessive compulsive disorder

Isaac Fradkin^a
The Hebrew University of Jerusalem

Rick A. Adams^{bcde}
University College London

Thomas Parr^f
University College London

Jonathan P. Roiser^b
University College London

Jonathan D. Huppert^a
The Hebrew University of Jerusalem

In this article, we develop a computational model of obsessive compulsive disorder (OCD). We propose that OCD is characterized by a difficulty in relying on past events to predict the consequences of patients' own actions and the unfolding of possible events. Clinically, this corresponds both to patients' difficulty in trusting their own actions (and therefore repeating them), and to their common preoccupation with unlikely chains of events. Critically, we develop this idea on the basis of the well-developed framework of the Bayesian brain, where this impairment is formalized as *excessive uncertainty regarding state transitions*. We illustrate the validity of this idea using quantitative simulations and use these to form specific empirical predictions. These predictions are evaluated in relation to existing evidence, and are used to delineate directions for future research. We show how seemingly unrelated findings and phenomena in OCD can be explained by the model including: a persistent experience that actions were not adequately performed and a tendency to repeat actions; excessive information gathering (i.e. checking); indecisiveness and pathological doubt; overreliance on habits at the expense of goal-directed behavior; and over-responsiveness to sensory stimuli, thoughts, and feedback. We discuss the relationship and interaction between our model and other prominent models of OCD, including models focusing on harm-avoidance, not-just-right experiences, or impairments in goal-directed behavior. Finally, we outline potential clinical implications and suggest lines for future research.

Analysis code for all simulations: <http://doi.org/10.17605/OSF.IO/2X8VD>

Keywords: Obsessive compulsive disorder; computational psychiatry; Bayesian brain; active inference; Prediction error; Habits

^aThe Hebrew University of Jerusalem, Mt. Scopus, Jerusalem, Israel.

^bInstitute of Cognitive Neuroscience, University College London, 17 Queen Square, London, United Kingdom

^cDivision of Psychiatry, University College London, 149 Tottenham Court Road, London, United Kingdom

^dMax Planck-UCL Center for Computational Psychiatry and Ageing Research, 10-12 Russell Square, London United Kingdom

^eDepartment of Computer Science, University College London, 90 High Holborn, London, United Kingdom

^fWellcome Centre for Human Neuroimaging, Institute of Neurology, 12 Queen Square, University College London, London, United Kingdom

Corresponding author: Isaac Fradkin, Department of Psychology, The Hebrew University of Jerusalem, Mt. Scopus, Jerusalem, 91905 Israel, Phone 972-2-5883376; Fax 972-2-5881159, Email: itzik.fradkin@mail.huji.ac.il

Preparation of this manuscript was supported by the Israel Science Foundation (grant # 1698/15) to Jonathan D. Huppert. Some of these ideas and simulation results were presented as a poster in the 2018 meeting of the Israeli Society for Cognitive Psychology meeting, the 2018 conference of the Israeli Society of Biological Psychiatry conference, and the 2019 conference of the Society for Psychopathology Research conference. Conflict of interest: None.

Obsessive compulsive disorder (OCD) is a psychiatric disorder characterized by intrusive thoughts, urges or images (i.e., obsessions) that cause marked distress or anxiety, and repetitive behavioral or mental rituals (i.e., compulsions; American Psychiatric Association, 2013). Some prominent theories explain compulsivity as resulting from a neurocognitive impairment in goal-directed control, leading to repetitive, rigid and habitual behavior (Gillan, Kosinski, Whelan, Phelps, & Daw, 2016; Robbins, Gillan, Smith, de Wit, & Ersche, 2012; Voon et al., 2015). Conversely, traditional cognitive-behavioral models consider compulsions to be goal-directed behaviors, specifically driven by attempts to prevent potential harm implicated by obsessions (i.e. *harm-avoidance*; Foa & Kozak, 1996a; Salkovskis, 1999). Yet other models emphasize the role of *not just right experiences* (NJREs) in OCD, defined as subtle experiences that specific actions, movements, or external stimuli are currently not 'just-right' (e.g., "the picture is not hanging perfectly straight"; e.g., Coles and Ravid, 2016). Here, we aim to integrate these previous theories in a novel computational model of OCD.

OCD patients often report excessive uncertainty and pathological doubt, which have been described as hallmark features of the disorder (Dar, 2004; Dar, Rish, Hermesh, Taub, & Fux, 2000; Frost & Shows, 1993; Pitman, 1987b; Sarig, Dar, & Liberman, 2012). Indeed, OCD was called '*folie a doute*' – the doubting disorder. Clinically, it has been suggested that uncertainty and doubt in OCD are driven by inferential processes overestimating the diversity of possible chains of events (i.e. obsessing about everything that *could* happen; Aardema, O'Connor, Péliissier, & Lavoie, 2009). Indeed, it has been shown that eliciting such 'what could happen' scenarios in non-clinical participants induces OCD-like cognitions (Giele, Van den Hout, Engelhard, Dek, & Hofmeijer, 2011). Normally, people can use prior experience to predict both the consequences of their actions, and the evolution of external events. Here we suggest that OCD is characterized by an impairment in this prediction mechanism (especially – but not uniquely – relating to the consequences of one's own actions).

To illustrate, consider yourself leaving home, inserting your wallet into your bag. Normally this behavior would allow you to form a highly confident belief that your wallet is there (given prior experiences with similar actions). In many cases you will have no need of sensory information (e.g., looking inside the bag or making sure that you can feel the wallet).

However, consider a situation in which you knew that your hand is trembling, or alternatively, that your bag has a hole in it. In such cases you will be insecure that your wallet is in your bag, despite placing it there. Therefore, you will rely more heavily on your senses, and probably will also periodically check that your wallet is still in your bag.

Formally, alternative locations of the wallet may be thought of as alternative states. The term state refers to a configuration of the environment at a given time. Clearly, the true state of the total environment is very complex, and the location of the wallet is a small part of it. Our model is based on the Bayesian brain framework (Friston, 2010; Knill & Pouget, 2004), where the true state of the environment is assumed not to be directly observable (i.e. a hidden/latent state). Rather, the brain is assumed to probabilistically infer states (e.g., the location of the wallet) by integrating different sources of information: sensory input and predictions regarding how states evolve (i.e. state transitions). First, sensory input is assumed to have a close but probabilistic relationship with the hidden state. For example, the feeling of a bulge in the bag can be used as evidence that the wallet is likely there, but is not conclusive (and different sensory modalities have different reliabilities). Another source of information is beliefs regarding the probabilistic relationship between actions and their consequences (we use the term *action-dependent transitions* to refer to this relationship). For example, inserting the wallet in the bag makes it likely to be in the bag. Finally, one can use previous knowledge about the way states evolve over time (henceforth: *action-independent transitions*) as another source of information. For example, having seen the wallet in the bag in the past makes it likely to be in the bag now. Both sorts of predictions can help infer present states (from past states and actions; particularly when current sensory input is ambiguous), and predict future states (from present states and current or future actions).

The Bayesian brain framework further assumes that the perceived uncertainty of different sources of information determines the relative weights they are given in this inferential process (Knill & Pouget, 2004). This allows us to formalize the process in which not being able to rely on past states or one's actions (as in the hole in the bag or trembling hand examples above) leads to increased gathering of sensory information (e.g., actively feeling the bulge in your pocket to verify that the wallet is there; Figure 1). Using this formalization, we propose that patients with OCD have difficulties in using actions and past states as reliable

sources of information even in predictable and relatively stable situations, leading to excessive reliance on sensory information (e.g., checking to verify that the wallet is there even when the bag has no hole in it; returning to the kitchen to verify that the stove is off, although it was turned off before leaving the kitchen), as well as to overall uncertainty regarding the current state (i.e. current location of wallet or whether the stove is off). Thus, we suggest "*excessive uncertainty regarding state transitions*" (which we term also "*transition uncertainty*") to be a core feature of OCD. In other words, we suggest that there is a difficulty in being able to understand and predict changes (or the lack thereof) from one state to the next, particularly but not exclusively, in the case in which such changes are influenced by the individual's actions.

Using this computational perspective allows us to formulate specific, testable hypotheses. Specifically, we show that increased transition uncertainty should:

- 1) impair the ability to rely on action planning and execution as sources of information regarding the successful completion of actions, creating a recurring experience that (compulsive) actions were not done 'right' ;
- 2) lead to excessive checking and foraging behaviors, and, under some conditions – overreliance on habits.
- 3) increase the weighting of immediate feedback at the expense of predictions about state transitions, which can potentially explain different types of obsessions, including intrusive thoughts and patients' sensory over-responsiveness.

We begin by providing a more elaborate introduction to the Bayesian brain framework (a glossary of major computational terminology is presented in Appendix A). Then, we investigate the above hypotheses via a review of relevant findings and computational simulations designed to establish the quantitative explanatory strength and specificity of the proposed model. Finally, we discuss the relationship of our model to other models of OCD, as well as the model's specificity to OCD in relation to other disorders which have related impairments.

The Bayesian brain and active inference

Predictions shape our experiences and functioning in the world. For example, having your car go straight has a completely different meaning if you expect it to turn left (e.g., after turning the steering wheel). The computational quantity underlying such an experience is termed a prediction error (PE). Predictions and PEs have been implicated in a broad variety of cognitive and behavioral processes. The highly influential predictive coding framework suggests that perception entails the integration of top-down predictions (e.g., the expected motion of the car after turning the steering wheel) and bottom-up information (e.g., sensory input regarding the car's actual motion; Friston, 2010; Rao & Ballard, 1999; Shipp, Adams, & Friston, 2013). The predictions depend on representations that have a greater spatiotemporal scale or degree of abstraction (e.g. representations of the occluded parts of a visual objects, or of the next word in a sentence) and are formed via experience. Bottom-up sensory input is processed only to the extent that it generates a PE (e.g., a surprising next word in a sentence, having the car to go straight despite turning the wheel).

At the neurobiological level, the brain was suggested to constitute a hierarchical model in which higher levels send predictions (about the causes of sensory inputs – represented by neural activity) to lower levels. Lower levels compute PEs (the difference between the predicted and actual activity at that level) and return these PEs to the higher levels, which then use them to refine the original predictions (for an illustration see Edwards, Adams, Brown, Parees, & Friston, 2012; Figure 1 and 3). Predictive coding is efficient in the sense that it prioritizes unexpected stimuli (hence most likely to be informative) in the allocation of resources. Furthermore, the incorporation of previous experience improves the accuracy of perception by minimizing the impact of sensory noise.

Computationally, predictive coding models implement Bayesian inference (Bogacz, 2017; Friston, 2010; Knill & Pouget, 2004). Bayes theorem provides the optimal way of integrating prior beliefs (predictions) and current (sensory) evidence, by weighting each according to its uncertainty (or noise; mathematically measured as variance). This is also known as *precision-weighting* of the PEs, because computationally, the weight of the PEs is determined by the relative precision (inverse variance) of the prior belief and the current evidence. For example, as sensory input becomes more informative (less noisy), PEs are given a higher weight.

The notion of the weight of a PE is important. Due to sensory and neuronal noise, no stimulus is ever fully predictable. The precision-weighting of the PE has a key role in accounting for the expected noise when 'deciding' the extent to which PEs should update inference and learning. Greater weight could make even trivial deviations from predictions appear important and salient, rather than disregarded as mere noise. This, in turn could result in abnormal cognitive, behavioral and emotional responses (Fletcher & Frith, 2009).

Bayesian inference in a dynamic world with active agents

Usually, with time and learning, predictions (derived from past data) become increasingly reliable. However, in dynamic environments, these predictions should become more uncertain when it is likely that the environment has changed – which is often the case when one acts on the environment. Figure 1 illustrates this effect, showing that the relative weight of sensory information increases when the uncertainty regarding state transitions increases. In other words, when one believes that the environment is likely to change (possibly as a consequence of one's actions) in an unpredictable fashion (i.e., under high transition uncertainty), current information and sensory evidence should be considered more reliable than previously accumulated knowledge (which is likely no longer relevant). This results in an increase in the weight assigned to the PEs. It is important to note that increased transition uncertainty also leads to increased uncertainty regarding inferred hidden states.

Predictions regarding action-dependent state transitions (in addition to sensory feedback) can serve as information regarding the consequences (e.g. the completion) of an action. Such predictions will be given greater weight when sensory feedback is missing or unreliable (J. Moore & Haggard, 2008). Consider the case of washing one's hands after going to the bathroom. In this case, there is no sensory feedback that can guarantee that one's hands are in a germ-free state. For most people, a conclusion of 'sufficient cleanliness' may be reached given handwashing, as the probability of transitioning from a 'dirty' to a 'clean' state, given handwashing, is high. However, if the probability for this transition is more uncertain, this inference will be less likely, and sensory evidence may be unable to compensate. We return to this example below.

Active inference

The above examples describe the role of Bayesian inference in action and perception. The active inference framework integrates action planning and selection within the same Bayesian model (Friston, 2010, 2011). For technical details, please see the Supplemental Material, Bogacz (2017) and Buckley, Kim, McGregor and Seth (2017). A basic tenet of active inference is that action planning can be performed using Bayesian inference on *expected* (rather than actual) observations (Friston, 2011; Friston et al., 2016), and that actions occur to resolve discrepancies between these expectations and actual observations. For example, from previous experience one can predict that when pressing the light switch, the light will turn on, and one will be able to see the objects in the room. Thus, using Bayesian inference on the *expected* observation involved in seeing the objects in the room, one can infer the optimal sequence of actions to be performed (i.e., *policies*, referring to a planned sequence of actions). Note that this involves a reformulation of goals as high-level expected states (priors; Friston, 2011). In that sense, high-level expectations (e.g., finding an object in the room), propagate downwards and induce lower-level expected observations (e.g., a lit room), used to infer the actions predicted to realize these expectations.

The process of selecting one of several (discrete) courses of action (i.e., policies) is formalized in active inference as a process of Bayesian model comparison, balancing accuracy and uncertainty. That is, actions are selected not only based on their expected consequences (i.e. whether one can either gain or learn something from taking the action), but also on the basis of how (un)certain these consequences are (FitzGerald, Dolan, & Friston, 2014; Friston et al., 2016). Transition uncertainty makes it hard to predict which action will lead to the best consequences (see Figure 1). This challenge can be dealt with in several ways. First, it has been shown that high transition uncertainty increases information gathering behaviors (Parr and Friston, 2017), designed to improve the predictions regarding the consequences of potential actions (by constantly updating these predictions). Second, such a challenge in goal-directed action selection can lead one to use simpler behavioral strategies, which do not require planning ahead (FitzGerald et al., 2014; Friston et al., 2016). Computationally, such behaviors have been argued to depend either on the simple, accumulated values of the different actions, while disregarding possible future changes in the environment (i.e. model-free learning; Daw, Niv, & Dayan, 2005), or in even simpler cases – on the degree to which these actions were repeated in

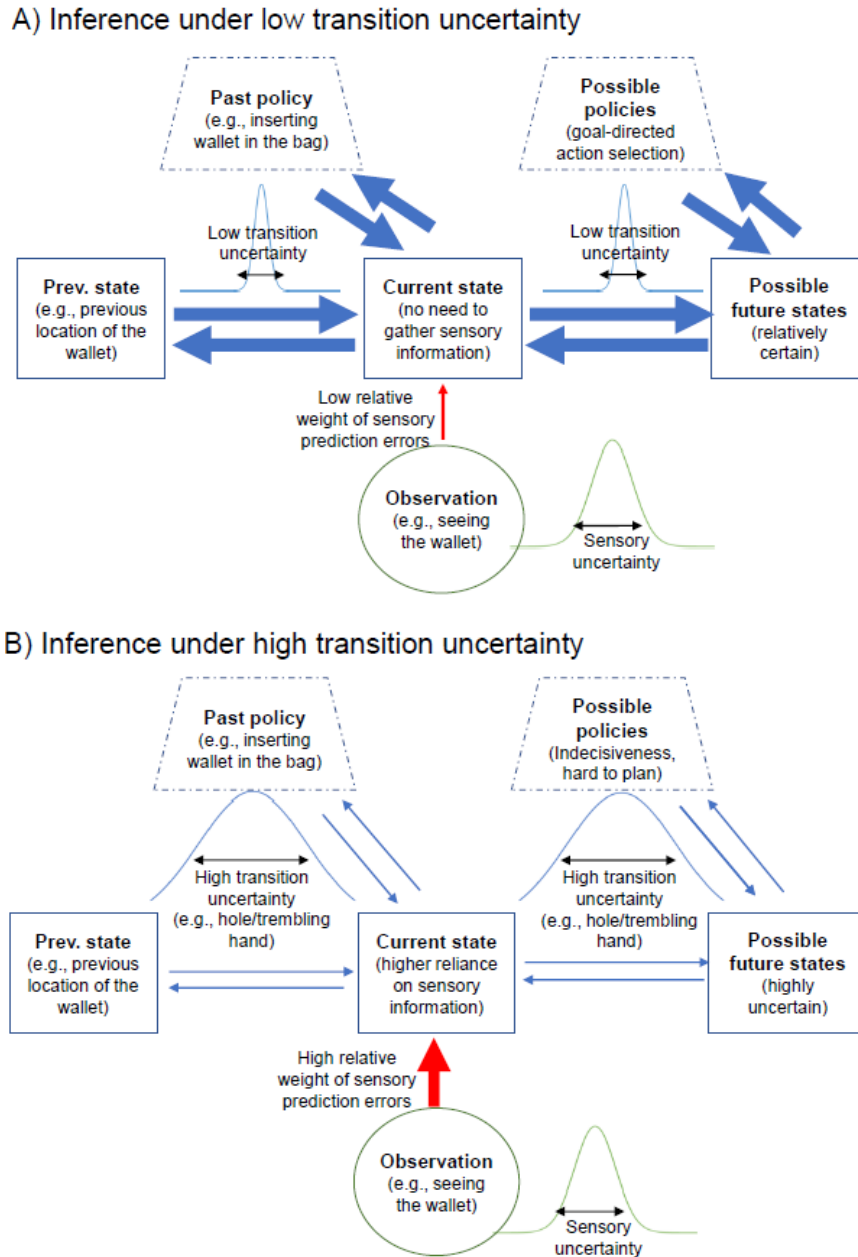


Figure 1 – An illustration of the effect of transition uncertainty on Bayesian inference. Inference regarding current states is based on precision-weighted (with precision denoted by width of arrows) integration of sensory information, and predictions derived from previous states and policies (i.e., sequences of actions; policies are not involved in the case of action-independent state transitions). Inference regarding future states and possible actions is based on current states and knowledge regarding the transitions (i.e. relationship) between the three. Panel A shows that when state-state and action-state transitions are known with sufficient certainty, the relative weight of sensory information (i.e., relative weight of sensory prediction errors) is reduced, and goal-directed action selection is easier. Panel B shows that when transitions are unreliable, sensory information is more important, and planning is much more difficult. Note that sensory uncertainty is equal in both panels. Backward arrows between states refer to the fact that knowledge about current states can be used to inform knowledge about previous states and actions.

the past (i.e. habitual behavior; FitzGerald et al., 2014), or simpler still – on innate (Pavlovian) responses (Dorfman & Gershman, 2019). In these cases, behavior is governed solely by the past, while ignoring potential

changes in the environment that could render the past irrelevant. Thus, because high transition uncertainty can greatly increase the uncertainty of goal-directed policies, simpler heuristic strategies may be prioritized

over potentially more effective (but uncertain) goal-directed ones.

Excessive transition uncertainty in OCD

We propose *excessive uncertainty regarding state transitions to be a core computational impairment in OCD*, that interacts with other key mechanisms (e.g., harm-avoidance). We discuss both action-dependent and action-independent transitions, as both may be relevant to OCD symptomatology, although the former might have a more specific role, given the importance of repetitive, compulsive behavior in the clinical syndrome. We now review existing empirical data and present novel simulations pertaining to the hypotheses (presented above) derived from this computational impairment.

High transition uncertainty leads to persistent experiences that actions were not adequately performed.

The repetitive nature of compulsions makes patients look as if they are stuck in a loop, unable to achieve a sense of completion (Pitman, 1987b, 1987a; Summerfeldt, 2004). Based on this phenomenology, Pitman (1987b) suggested that "the obsessive-compulsive's internal comparator mechanism is faulty. No matter what perceptual input it receives, it continues to register mismatch [...]. It may be that in fact the action was well done, but the defective comparator cannot register it" (p. 340). According to his framework, obsessions are conceptualized as excessive mismatch signals, whereas compulsions are the behavioral output designed to minimize these errors signals, though without long-term success. This idea has been highly influential in inspiring theoretical and empirical work (e.g., Coles & Ravid, 2016; Gehring, Himle, & Nisenson, 2000).

When integrated within the Bayesian brain framework, such excessive mismatch signals can be formalized as highly weighted action-related PEs (relative to the weighting of predictions). Indeed, high uncertainty regarding action-dependent state transitions can reduce reliance on action execution (and the accompanying predictions regarding its consequences in the world) in assessing whether the action's goal was achieved. Such an impairment increases the relative weighting of sensory information in inferences about the consequences of actions, highlighting even trivial deviations from predictions (experienced as doubts, obsessions or NJREs). It also leads to an overall uncertainty regarding the completion of actions.

Several studies have shown that OCD is related to increased weighting of sensory feedback when asked to reproduce a specific muscle tension, even when the sensory feedback was false (Lazarov, Dar, Liberman, & Oded, 2012a, 2012b; Lazarov, Liberman, Hermesh, & Dar, 2014; Zhang et al., 2017). Critically, these findings were specific to OCD, and did not generalize to anxiety disorders (Lazarov et al., 2014). Here, we argue that these results can be explained by decreased confidence in predictions regarding action-dependent transitions (due to increased transition uncertainty), rendering patients more reliant upon sensory feedback. Indeed, our explanation is supported by a study manipulating the availability of such predictions (Ezrati, Sherman, & Dar, 2018). In that study, participants were first asked to either rotate their heads to a specific angle (active condition), or let the experimenter rotate their heads (passive condition). Participants were then asked to reproduce the rotation angle. Higher (subclinical) OC symptoms were associated with less accurate reproduction of the rotation angle in the active condition but not in the passive condition. Because the two conditions differ only in the availability of predictions regarding action-dependent transitions as an additional source of information in the active condition, this implies increased transition uncertainty.

Self-generated movements are known to be associated with down-weighting of their (expected) sensory consequences (i.e. *sensory attenuation*). A familiar example is our inability to tickle ourselves (Blakemore, Wolpert, & Frith, 2000). Indeed, additional evidence for patients' impairments in utilizing predictions regarding action-dependent transitions comes from studies showing decreased attenuation of sensory consequences of self-generated actions in OCD (Gentsch, Schütz-Bosbach, Endrass, & Kathmann, 2012; Rossi et al., 2005). This suggests that even when the sensory feedback of an action should be ignored, it is interpreted as critical information, being registered as a 'not just right' signal.

Illustrative computational simulations. To demonstrate more concretely how increased uncertainty regarding action-dependent transitions leads to overreliance on sensory feedback, and the obsessional experience that an action was not done 'just right', we conducted an illustrative simulation of Bayesian integration of information during hand-washing. We applied a




Figure 2 – Results of the hand-washing simulations, for simulated 'healthy controls' with optimal transition uncertainty (A), and 'OCD patients' with overestimated transition uncertainty (B). The black curve represents the true state (hand dirtiness), green dots represent the sensory feedback regarding dirtiness. The blue bars represent the posterior estimates of dirtiness. The red line represent the subjective criteria for 'sufficient cleanliness', and black dots on the top represent simulated hand-washing episodes. This figure shows how an agent with increased transition uncertainty relies less on predictions stemming from action execution (action-dependent state transitions), leading to higher reliance on sensory feedback, higher uncertainty, a paradoxical increase in uncertainty following compulsive actions, and more handwashes.

particle filter, which is a flexible Bayesian inference algorithm suitable for the integration of (continuous) observations and state dynamics (Sanborn, 2017; Speekenbrink, 2016). The particle filter was used to examine the effects of uncertainty regarding the prediction that hand-washing decreases dirtiness. To focus specifically on action-dependent transitions, we assumed that actual and predicted dirtiness change only when the agent washes their hands. To simplify, we assumed that agents use a univariate sensory signal (integrating all available sensory information) which is, objectively and subjectively, highly noisy (i.e. we usually do not have precise sensory information regarding how dirty/contaminated our hands are). Hand dirtiness was measured on an arbitrary scale, from 0 (absolutely clean) to 100 (very dirty). At each timepoint, the agent integrates predictions and sensory observations to estimate the current level of dirtiness, and its uncertainty. The focus of these simulations was on the perceptual and inferential qualities of hand-washing, rather than on the generation of action itself. Thus, we applied a

simplistic, illustrative control algorithm, in which hand washes were assumed to be discrete events. The probability of handwashing at any (discrete) timepoint is a function of the agent's confidence (i.e. the 95% credible interval) that its hand dirtiness surpasses a predefined threshold. A more detailed and realistic account of actions under transition uncertainty is delineated in the active inference simulations below. The code used for all the simulations presented in this paper can be found in <http://doi.org/10.17605/OSF.IO/2X8VD>.

The results of this simulation are presented in Figure 2. Panel A represents optimal estimation of transition uncertainty (i.e. a healthy control agent), and panel B corresponds to overestimated transition uncertainty (i.e. an agent 'with OCD'). Figure 2 demonstrates how excessive transition uncertainty results in inferred states that are more uncertain (wider blue credible intervals) and more affected by the noisy sensory input (green dots). For this reason, they are less representative of the real levels of hand dirtiness (black line). This

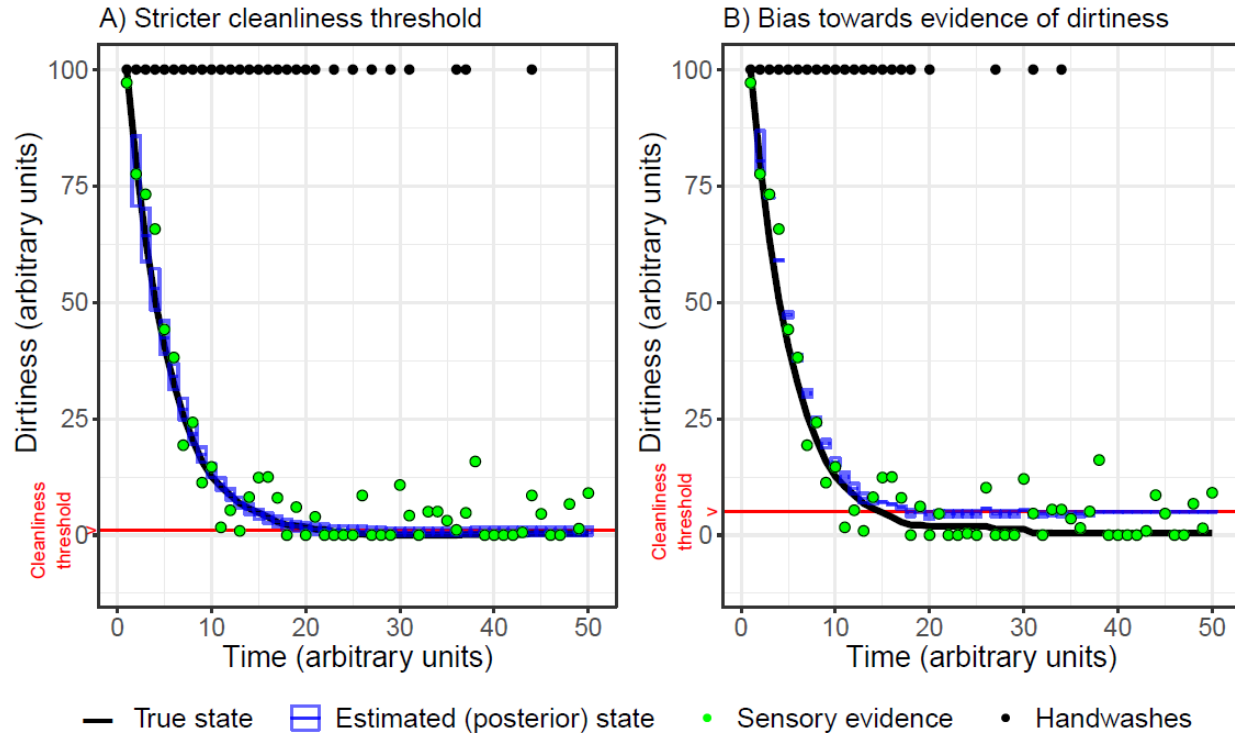


Figure 3 – Results of the hand-washing simulations for alternative possible explanations for OCD-like behavior. Panel A corresponds with an agent with a strict criterion for cleanliness (corresponding with importance of being clean or general perfectionism). Panel B corresponds with an agent with an affective bias, giving a higher weight for negative evidence (i.e. dirty hands). This was accomplished by using a skewed likelihood distribution, such that the agent believes that positive outcomes 'cannot be trusted'. Although both types of alternative impairments increase washing behavior, they do not result in increased uncertainty or the paradoxical effect of action on increasing uncertainty depicted in Figure 2.

results in more instances in which the inferred state surpasses the goal or threshold (red line), thus resulting in more washing behavior (black dots on the top of the plot). In other words, the agent remains uncertain as to whether their hands are clean enough, and is unable to ignore meaningless not-just-right (not clean-enough) sensory experiences – and therefore washes again. Notably, greater action-dependent transition uncertainty led to an increase in the overall posterior uncertainty only after washing. This effect can partially explain the well-documented paradoxical effect of compulsions on increasing uncertainty and additional compulsive urges (for a review see van den Hout, van Dis, van Woudenberg, & van de Groep, 2018).

This setup has also allowed us to examine how altering other parameters of the model, which relate to other previously suggested etiological factors in OCD, affects inference. We examined the effects of: a) making the cleanliness threshold stricter, corresponding to a higher importance of being clean (Figure 3A); and b) introducing an affective bias in which negative stimuli (i.e. my hands are dirtier) have a higher weight (Figure

3B), by using a left-skewed distribution to represent the relationship between hidden states and observations (corresponding with the belief that 'my hands look cleaner than they really are'). Both cases resulted in more washing behaviors (first row, black dots), and the affective bias manipulation also resulted in a biased inference (i.e. a bias towards the belief that one's hands are dirty). However, the pattern of compulsive behavior was different in these cases, as it seemed to gradually decrease with time. In contrast, in the case of transition uncertainty, washing behavior led to an increase in uncertainty and the probability for additional washing. More generally, uncertainty regarding the posterior dirtiness was increased only in the case of transition uncertainty (Figure 2B). Stated otherwise, these alternative potential impairments predict an experience of *being certain that one's hands are too dirty*, whereas transition uncertainty leads to an experience of *having doubts that one's hands might be dirty*. Indeed, the large majority of patients with OCD have insight into their symptoms, and acknowledge that they are not certain about their obsessions (e.g., Foa and Kozak, 1995).

Thus, if increased uncertainty is indeed a key characteristic of OCD, these alternative potential impairments seem to have a lower explanatory power for that aspect of OCD phenomenology. Of course, an alternative is that these processes interact or occur simultaneously.

It is important to consider the necessary conditions for obtaining these simulation results. First, the assumption that objective sensory noise is considerably larger than objective transition uncertainty is critical to ascertain that predictions (and their uncertainty) have a considerable impact on normal inference. For example, a substantial reduction in objective sensory noise eliminated the effects described above, because sensory information regarding dirtiness became sufficiently precise. This questions the relevance of our model to compulsive situations in which precise sensory feedback is available (e.g., seeing that the stove is off). To account for situations in which one does receive precise feedback it is important to consider that (Bayesian) inference is not limited to the time in which feedback is obtained, but likely occurs when feedback is no longer (or not yet) available. For example, it is mainly when one can no longer see any evidence for the state of the stove/door (e.g., after leaving the room) that doubt and checking will be likely to emerge. Similarly, transition uncertainty will also affect inference regarding future feedback. Below we present active inference simulations demonstrating how such mechanisms can lead to the emerge of excessive checking and other symptoms.

In contrast to the effects of sensory noise, the simulation results were robust to violations of other assumptions. For example, whereas using a normal distribution to model action-dependent transitions implies some chance that handwashing might increase dirtiness, using a trimmed distribution instead did not eliminate the effect of transition uncertainty on compulsive handwashing. Similarly, the results were not substantially altered when using a different Bayesian inference algorithm (the Kalman filter). Finally, allowing actual and predicted dirtiness to change even when no washing was performed (i.e., including action-independent transitions) resulted in compulsive checking behavior, but eliminated the paradoxical effect of washing demonstrated above.

To summarize, in this section we suggested that uncertainty regarding action-dependent transitions underlies OCD patients' obsessional experience or feeling "that actions that they perform are incompletely achieved [...] though to observers these actions may appear to have been performed perfectly well" (Pitman, 1987a, p. 226). Specifically, as demonstrated by the simulations above, uncertainty regarding the expected

consequences of actions makes one both overly sensitive to noisy sensory feedback (which produces transient experiences that the action did not reach its goal), and highly uncertain about the environment's (or their own) state. The predictions resulting from this simulation are supported by empirical findings regarding increased uncertainty and overreliance on sensory feedback during simple actions in OCD. Furthermore, our simulations provide insight into the paradoxical effect of compulsions on increasing uncertainty and thereby stimulating yet more compulsions (van den Hout et al., 2018). That is, actions (including compulsive actions) necessarily result in a change in the environment, which, combined with increased uncertainty regarding such changes will have a destabilizing effect on one's beliefs regarding the state of the environment (or their own). Interestingly, other previously suggested etiological factors in OCD (i.e., a stricter threshold for cleanliness, and an affective, negative bias) generated a different pattern, not involving increased uncertainty or such a paradoxical effect of compulsions.

High transition uncertainty leads to compulsions, and to the predominance of habits over goal-directed behavior

We propose above that increased transition uncertainty can lead to PEs of increased weight, evident in the form of an obsessional experience that actions were not well executed. Here, we discuss the relevance of increased transition uncertainty to the understanding of compulsive behaviors. We focus on three common conceptualizations of compulsions: i) as a response to obsessions; ii) as exaggerated checking and information gathering behavior; and iii) as excessive habitual behavior.

First, some compulsions can be conceptualized as a response to, and an attempt to minimize highly weighted sensory PEs. This idea was already illustrated in the hand-washing simulations presented above, where increasing transition uncertainty resulted in higher reliance upon sensory input, leading to transient experiences (i.e. obsessions or NJREs) that things are not as expected (see Figure 2). Then, because the agent was assumed to have a prior belief (i.e., goal) regarding how things should be (e.g., having clean hands), such obsessions lead to compulsive washing.

Second, compulsions in OCD are often formalized as attempts to reduce uncertainty and doubts by various checking behaviors. Numerous studies have demonstrated increased perceptual information gathering behavior in OCD even in neutral contexts, evident in

longer search times and more fixations in visual search tasks (Toffolo, van den Hout, Engelhard, Hooge, & Cath, 2016; Toffolo, van den Hout, Hooge, Engelhard, & Cath, 2013), and more checking behavior in change detection tasks (Clair et al., 2013; Jaafari et al., 2013; Rotge et al., 2008). Furthermore, such increased information gathering appears to be specific to OCD, and is not evident in anxiety disorders (Toffolo et al., 2016).

In active inference, action serves the role of sampling the environment as a method of improving one's model of the world, thus minimizing potential future PEs. Indeed, it has been shown that greater transition uncertainty leads to increased (visual) information gathering (Parr & Friston, 2017). Considering the wallet example above, this makes perfect sense: If one cannot rely on past states and actions to infer present states, the optimal behavior is to continue gathering information. From this perspective, excessive checking behavior, evident even in neutral environments, could be interpreted as evidence for increased uncertainty regarding state transitions.

Third, a considerable body of literature conceptualizes OCD as an impairment in the balance between goal-directed and habitual behavior (Gillan et al., 2014, 2011; Gillan & Robbins, 2014). More generally, compulsivity has been characterized by maladaptive behavior that appears to have no goal-directed value, instead reflecting stimulus-driven, habitual behavior that often results in adverse consequences (Robbins et al., 2012). However, whereas overreliance on habits usually entails perseverative, inflexible behavior, a recent meta-analysis showed that OCD patients exhibit impaired performance not only when required to shift behavior (indicating perseveration), but also when required not to shift behavior (potentially indicating overly exploratory responses; Fradkin et al., 2018). Furthermore, it is unclear whether overreliance on habits in OCD entails a difficulty in inhibiting habits, stronger habit learning, or in fact an impairment of goal-directed control (Gillan et al., 2016; Gillan & Robbins, 2014). Critically, early studies have shown that the acquisition of automatic stimulus-response associations (mostly in the context of procedural learning) is impaired in OCD patients, who instead seem to over-recruit controlled, goal-directed strategies (Deckersbach et al., 2002; Rauch et al., 1997; Rauch et al., 2007), which might suggest overly deliberate goal-directed behavior.

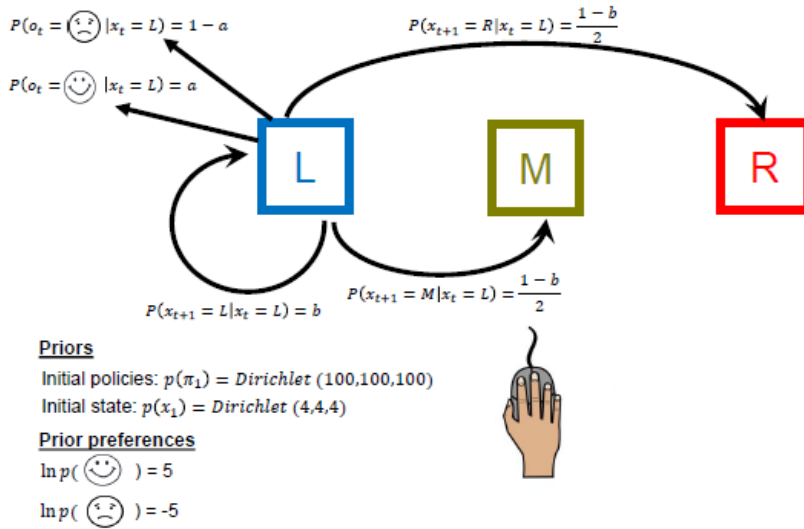
These apparent contradictions suggest that there is a need for a principled account as well as boundary conditions, explaining when OCD patients over-rely on habits, and when they exhibit overly deliberate or exploratory goal-directed behavior. We suggest that given

the active inference conceptualization of the arbitration between goal-directed and habitual behavior, both types of behavior flow naturally from increased transition uncertainty. High transition uncertainty results in higher uncertainty regarding goal-directed policies because it is harder to predict future states resulting from these policies. Thus, since transition uncertainty makes goal-directed behavior computationally expensive and highly uncertain, habits become Bayes-optimal actions.

Furthermore, when conceptualizing habits as learned expectations regarding one's actions in a given state (i.e., "how do I usually behave in this context"; FitzGerald et al., 2014; Friston et al., 2016), one could hypothesize that exaggerated habitual behavior will emerge in OCD only for responses that were repeated with sufficient consistency to form such expectations. In contrast, in accordance with the results of a recent meta-analysis of cognitive flexibility in OCD (Fradkin et al., 2018), tasks that do not supply consistent feedback supporting the repetition of a specific policy (such as probabilistic reversal learning tasks) are expected to elicit more exploratory and volatile behavior, rather than more habitual behavior, in OCD (see Hauser et al., 2017; Vaghi et al., 2017 for compatible results). This could also explain why most habitual, repetitive compulsions occur in everyday situations (e.g., hand-washing, door-locking, etc.), rather than in unusual situations (e.g., having to make major life decisions), where uncertain state transitions induce indecisiveness. For example, choosing between dishes on a unfamiliar restaurant menu could easily paralyze an OCD patient with indecision (e.g., Frost and Shows, 1993). In contrast, this can be avoided by restricting the range of restaurants and ordering the same dishes.

Illustrative computational simulations. To make these ideas more concrete we conducted several simulations of decision making in two different reversal learning tasks. Reversal learning tasks are suitable for investigating transition uncertainty because they require participants to infer whether past cue-outcome contingencies predict present cue-outcome contingencies (corresponding with action-independent transition uncertainty). In the first task, resembling classic tasks commonly administered to OCD patients to measure perseveration and habit learning (Fradkin et al., 2018), participants are required to detect contingency changes via trial-and-error. Here, we used a task (see Figure 4A) where the (simulated) agent was expected to choose, on each trial, between three cues. The agent knows that only one cue is correct at any time, and that the correct

A) generative model for the reversal learning task



B) generative model for reversal learning + explicit checking




Figure 4 – a representation of the agent's generative models in the active inference MDP models.

Panel A corresponds with a simple reversal learning task, with three possible cues, only one of which is correct (leading to positive outcome with probability a) at each trial.

Panel B corresponds to a reversal learning task with two cues, and one 'checking' cue allowing the agent to reveal the current context (i.e. correct cue), without the need to click on the exploitative cues directly (i.e. avoiding the risk for a punishment). In both panels, b corresponds with perceived transition uncertainty, and a corresponds with actual cue-outcome contingencies.

cue's identity (i.e. context) can change occasionally. The agent was provided with rewarding/punishing feedback with respect to whether the correct cue was chosen.

In the second task, we examined exploratory behavior more directly by allowing agents to choose a cue that reveals the current context (i.e., correct cue) directly, without risking loss. This more closely reflects checking behavior in real life. In this task (corresponding to the three-arm maze, as described in Friston et al., 2016; see Figure 4B), only two of the cues had utilitarian value (i.e. gain/loss), whereas a third cue (i.e. the checking, green colored, cue) could be used to check

which of the 'utilitarian' cues is currently correct. On each trial, the agent was required to plan two steps ahead, under the assumption that checking (in the first step) usually serves the role of informing a decision (in the second step). Checking behavior had no cost other than missing the opportunity for a gain in this time step. To discriminate exploratory behavior from simple avoidance of a possible loss, the agent was also given the option of doing nothing.

Following a key assumption in active inference (Friston et al., 2016; Parr & Friston, 2017), in both tasks participants were assumed to have a motivation to

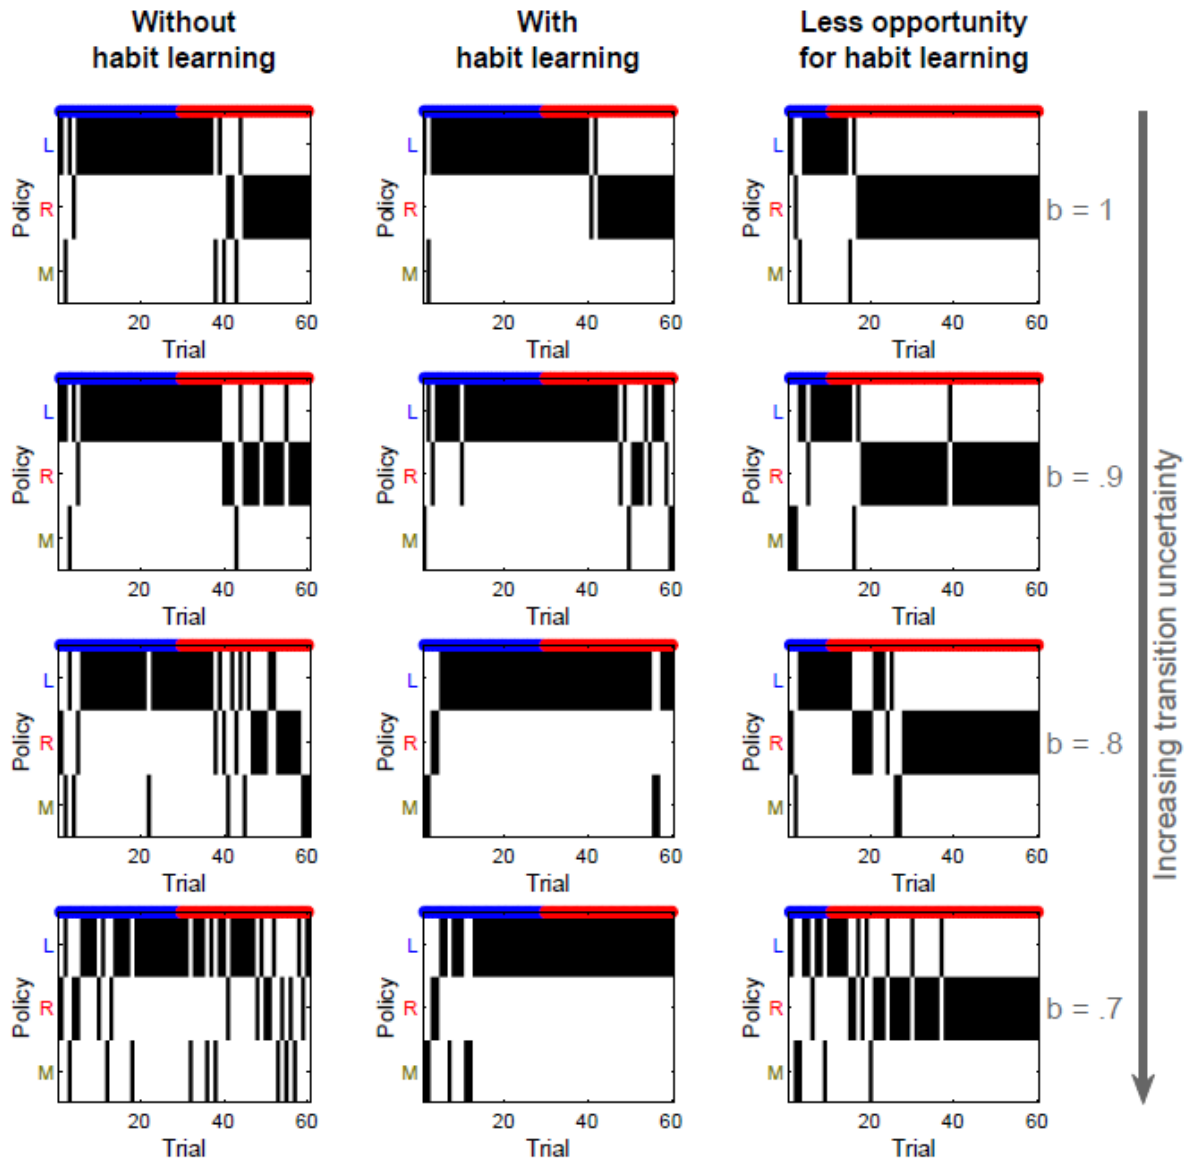


Figure 5 – Probabilities for chosen policies in a reversal learning task with three cues (L – left, R – right, M – middle; see Figure 4A), when increasing transition uncertainty (decreasing b), in an active inference MDP scheme. State-outcome contingencies (a) are set to 1 across these simulations (i.e. a deterministic reversal learning task). For each figure, the actual states on each trial are presented as colored lines at the top (blue = left cue, red = right cue). The left panel depicts the results of this simulation when not allowing the agent to learn habits, whereas in the middle panel – habit learning is allowed. The right panel represents a case where habit learning is allowed, but reversal occurs at the 10th trial (rather than the 30th, as in the other simulations), thus impeding the possibility to learn habits (due to insufficient exposures to this behavior). Increasing transition uncertainty leads to more ‘exploratory’ and volatile behavior in the left and right panels where habit learning is removed or reduced, but to more rigid and perseverative behavior in the middle panel where it is allowed.

both maximize rewards (and minimize losses) and minimize uncertainty and potential future PEs. Another cardinal assumption delineated above is that the relative weight given to goal-directed policies and habits is determined by their relative uncertainty, such that hab-

its emerge under two conditions: high uncertainty regarding the ability of goal-directed policies to reach their goals, and a sufficient opportunity to learn habits (i.e. form a prior regarding one’s behavior).

The agent's performance on these tasks was modeled using well-developed active inference Markov decision process schemes (MDPs; Friston et al., 2016, 2015; Parr & Friston, 2017). MDPs (in this case, partially observable MDPs) provide a probabilistic, mathematical framework suitable for modeling behavior where an agent may select between distinct actions at particular timepoints (i.e. a discrete state-space). Hidden states are assumed to evolve over time according to a state transition probability distribution, and to generate observable outcomes according to a likelihood distribution (describing the relationship between states and observations). Here we briefly describe the generative model for hidden *context* states (i.e., the currently relevant cue), and for observable *affektive* outcomes (e.g., reward/punishment). In Figure 4, a represents the (actual and estimated) probability for a rewarding outcome when choosing the correct cue. For simplicity we assume that *estimated* transition precision is represented by a constant b value for all states and trials (e.g., $b = 0.9$ corresponds with the agent's belief that the probability of a state x remaining unchanged at the next time-step is 0.9). The full generative model (see Supplemental Material) also includes hidden states and observable outcomes relating to the position where the agent has clicked (this is necessary to allow for action in active inference).

Three additional probability distributions were used in specifying the generative model: First, a uniform prior probability over states at $t=1$ was specified (the D vector in the Supplemental Material), meaning the agent doesn't know which cue is correct before starting the task. Second, the agent's preferences regarding outcomes (also defined as prior beliefs, in the C vector) give rewards a very high probability (i.e. the agent expects to make actions that lead to rewards). Specifically, there was a difference in log-probability of 10 ($e^5 / (e^{-5} + e^5) = 0.9999$) between rewarding and unrewarding outcomes, making the former overwhelmingly more desired. Third, to model habit learning, agents were allowed, in some simulations, to learn their own behavior (stored in the E vector), which then determined the prior probability of policies, irrespective of their expected outcome. The full generative model, and mathematical principles underlying learning and action in this scheme can be found in the Supplemental Material.

Illustrative simulation results for the first task (simple reversal learning), with a set to 1 (i.e. a deterministic cue-outcome relationship) are presented in Figure 5, depicting the probabilities for different actions (here

choosing the left, right or middle cue), whereas the correct cues' colors are presented on the top part of the figure (i.e. the 'left/blue' cue is correct in the first part of the task, and the 'right/red' cue is correct in the second part). The left column demonstrates how increasing transition uncertainty (i.e., decreasing b), when not allowing agents to learn habits, produces more exploratory and volatile behavior. This happens because high transition uncertainty means that, even if the agent believes a specific cue is correct at the first time-step, it is still uncertain whether this cue will still be correct at the next time-step. The middle column shows that when allowing the same agent to learn habits, transition uncertainty causes an increased reliance on habits, manifesting as perseverative selection of the left cue even after it is no longer advantageous (i.e. after the reversal). Note that even in this scenario, transition uncertainty also increases exploration before reversal, such that it takes the agent more time to learn habits. Finally, the rightmost column demonstrates the effects of one possible boundary condition for precluding habit learning, by introducing the reversal on an earlier stage (10th vs. 30th trial). This blocks the effect of transition uncertainty on perseveration, because in this case the agent does not have sufficient time to learn habits.

Figure 6 presents the average results from this simulation (averaged over 500 repetitions) in terms of the proportion of the different types of errors frequently reported in reversal learning tasks. Panel A replicates the results of Figure 5, showing that when impeding (left plot) or reducing (via early reversal; right plot) habit learning, increasing transition uncertainty results in more errors, with no evidence for a specific increase in perseveration. Rather, both choices of the previous cue (perseverative errors) and choices of the third (middle, gold) cue that is not reinforced either before or after the reversal (non-perseverative errors) likely indicate increased exploratory behavior. In contrast, when allowing for habits to form, higher transition uncertainty increases only perseverative errors after reversal. In all three scenarios, high transition uncertainty also increases errors before reversal.

Panel B focuses on the scenario where habit learning is allowed (with a reversal at $t=30$), while examining the effects of increasing the actual uncertainty of state-outcome mappings (reducing a). Notably, whereas under deterministic contingencies ($a = 1$), increased transition uncertainty induces more perseveration (mirroring the top-middle figure), under probabilistic contingencies, it reduces perseveration while increasing non-perseverative errors. The reason for this

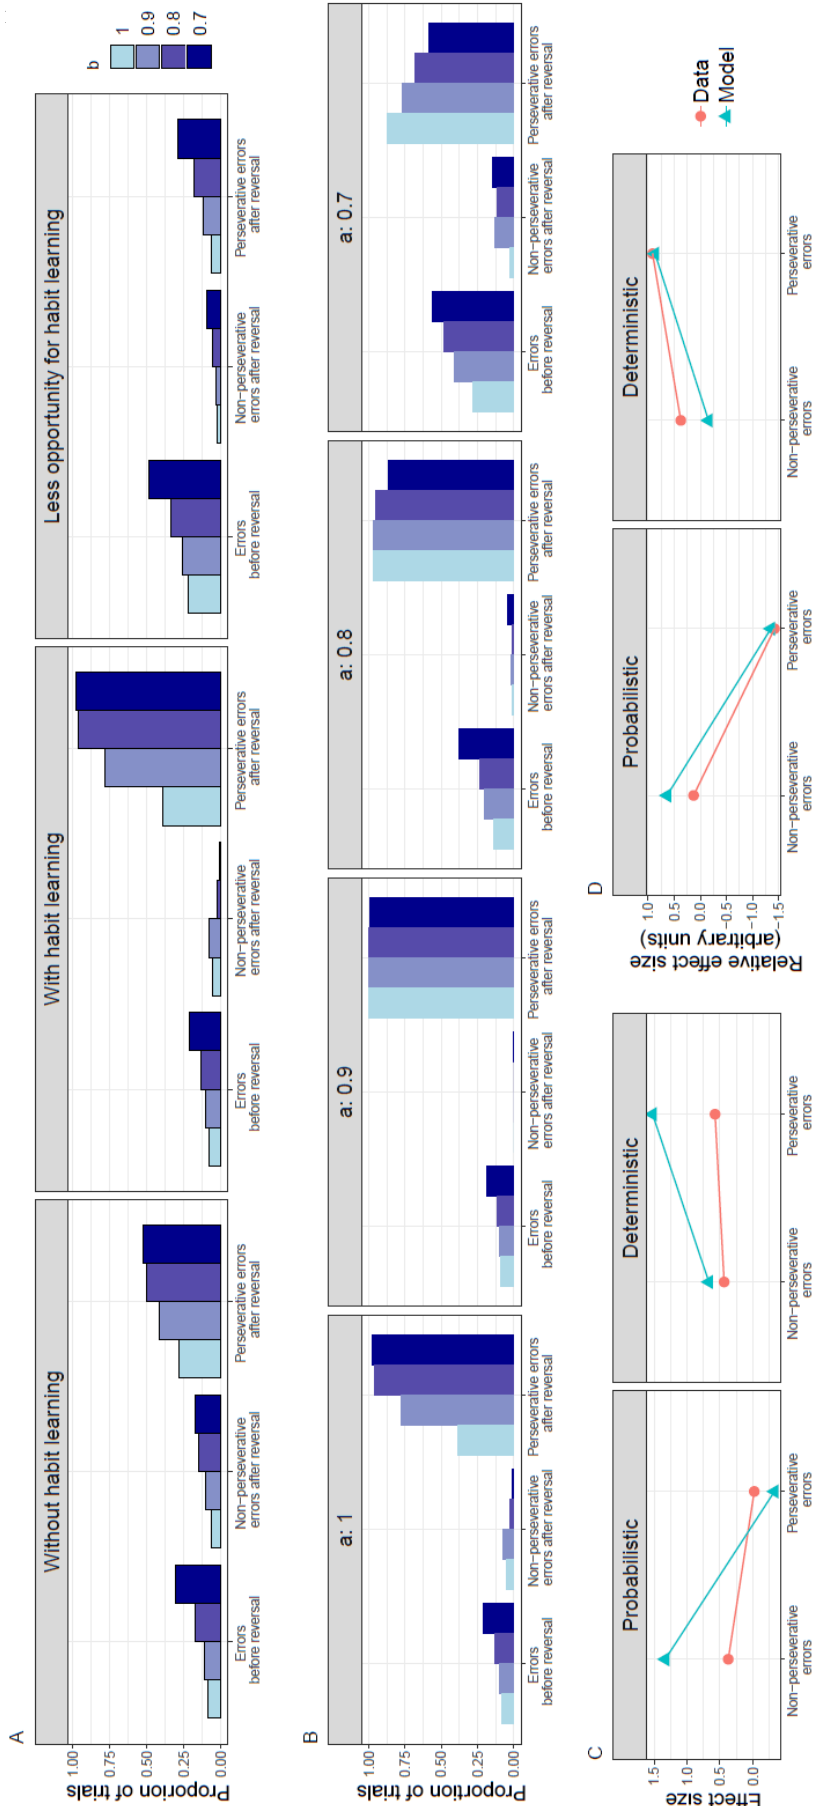


Figure 6 – This depicts the average results (over 500 repetitions of the simulation) from an active inference MDP scheme in a reversal learning task with three cues (see Figure 4A, and Figure 5). This figure focuses on three types of possible errors: errors performed before the reversal (indicating a difficulty in establishing or maintaining the initial context); perseverative errors performed after the reversal (A common measure of perseveration, corresponding with choosing the initially correct context after it is no longer correct); and non-perseverative errors done after the reversal (indicating choices of a cue that has not been correct either before the reversal or after it). Panel A represents the results for deterministic state-outcome contingencies (i.e. the simulation settings depicted in Figure 5 but averaged across 500 simulations). Panel B represents the results when allowing the agent to learn habits, with a reversal occurring at the 30th trial, while manipulating the state-outcome contingencies (*a*), thus examining the moderating effects of using a deterministic vs. a probabilistic reversal learning task. This figure shows that transition uncertainty increases perseveration (under transition uncertainty) only when previous responses were consistently reinforced (and repeated). In contrast, when such reinforcing feedback is less consistent (B) or short-lived (A, rightmost plot), non-perseverative errors are more common. Panels C and D compare the results of this simulation to real data obtained from a recent meta-analysis of inflexibility in OCD (*Fradkin et al., 2018*). In panel C the actual effect sizes (OCD-controls) obtained in that meta-analysis are shown, whereas panel D shows relative effect sizes (calculated by standardizing the raw effect sizes derived from *Fradkin et al., 2018*).

havior makes behavior less consistent, and thus less likely to become habitual.

Figures 6C and 6D depict a qualitative comparison of the effect sizes predicted by the model (calculated as the ratio of % errors for simulated high vs. low transition uncertainty agents), to real effect sizes obtained from a recent meta-analysis of flexibility in OCD (Fradkin et al., 2018). Whereas the magnitudes of the effect sizes cannot be directly compared (due to design and scaling differences), the general pattern in these results is consistent (see Figure 6D - effect sizes were standardized to provide a relative measure). As predicted by the model, OCD patients made both more perseverative and non-perseverative errors in deterministic tasks (Figure 6C), although effect sizes for the former are more pronounced (Figure 6D). By contrast, on probabilistic tasks, non-perseverative errors are more common, and in fact there is no clear evidence for perseveration (Figure 6C). One inconsistency is that the model predicts slightly fewer perseverative errors in OCD patients, whereas real data suggests a null effect. However, it should be noted that fewer perseverative errors were found in some probabilistic tasks included in the meta-analysis (Fradkin et al., 2018, Figure 3), and these inconsistencies might reflect differences in design or the operationalization of perseverative errors.

Figure 7 depicts the probability for different policies in second task: when the agent is given the option for explicit checking (i.e. choosing a cue that cannot contain reward or loss – only information regarding context). The features of the generative model for this are shown in Figure 4B. In this task, there may be reward at either the left or right cues, whereas the bottom (checking) cue provides information about the current position of the reward. On each trial, the agent was required to choose a two-step policy, where in each step the options were – left cue (L), right cue (R), check (C), or do nothing (N). The top left plot on Figure 7 – using a similar format to Figure 5 – replicates previous results (Friston et al., 2016), demonstrating a transition from an exploratory to a strictly exploitative strategy (when the agent is no longer checking before choosing the left cue), and falling back to exploration after the reversal. Consistent with our model, increased transition uncertainty (descending the left column, Figure 7) increases checking behavior, by preventing the agent from shifting to an exploitative strategy. Thus, an agent with high transition uncertainty always has to verify that the context did not change.

Next, we examined what happens (in the second task) when making checking behavior obsolete, by exposing the agent to the correct context at the beginning

of each trial from trial 45 onwards (yellow shaded area, Figure 7 middle column). Unsurprisingly, this results in purely exploitative behavior after trial 45, except for very high transition uncertainty, where checking continues (note that in this case the policy of first staying at the initial/neutral position can be considered as checking rather than just doing nothing). Furthermore, when allowing the agent to learn habits (right column, Figure 7), even more moderate values of transition uncertainty ($b=0.8$) cause checking behavior to continue, even when it serves no epistemic goal, because in this case checking behavior itself became a habit. This can be thought of as evidence for increased *epistemic habits* under high transition uncertainty, i.e. habitual information gathering even when nothing more can be learned. This resonates with the OCD-related phenomenology of continuing checking even when there is strong evidence that everything is OK.

This setup also allowed us to explore the role of harm-avoidance (representing either a trait, or a context variable) in the generation of OCD-related behavior. To do this, we used the same paradigms as before, but varied the aversiveness of incorrect decisions. Technically, this is achieved by increasing the prior precision of losses/punishments by a substantial amount (converting the log prior probability of losses to -10 instead of -5; decreasing this number further did not alter the conclusions below). As shown in Figure 8, harm-avoidance alone (when $b = 1$) cannot explain the ‘over-exploratory’ and habitual behaviors predicted by increased transition uncertainty, as observed in previous reversal learning studies of OCD (Fradkin et al., 2018). In fact, an agent with both high transition uncertainty and increased harm-avoidance was slightly less exploratory (fewer errors before reversal, and fewer non-perseverative errors after reversal) and less habitual (fewer perseverative errors when allowing agents to learn habits) than an agent with high transition uncertainty alone. This was true in a reversal learning task where checking the current state of the environment entailed the risk of obtaining a loss (see Figure 4A). However, in a setting that allowed the agent to check without risking loss (i.e. in the second task in which the checking option was separated from the stimuli leading to reward or loss; see Figure 4B), greater harm-avoidance did result in increased checking behavior (compare the first row of Figure 9, incorporating harm-avoidance when $b = 1$, to that of Figure 7 which does not include harm-avoidance). However, greater transition uncertainty (in the presence of habit learning) *uniquely* causes habitual checking when checking itself is devalued. Furthermore, increasing harm avoidance alone (even to much

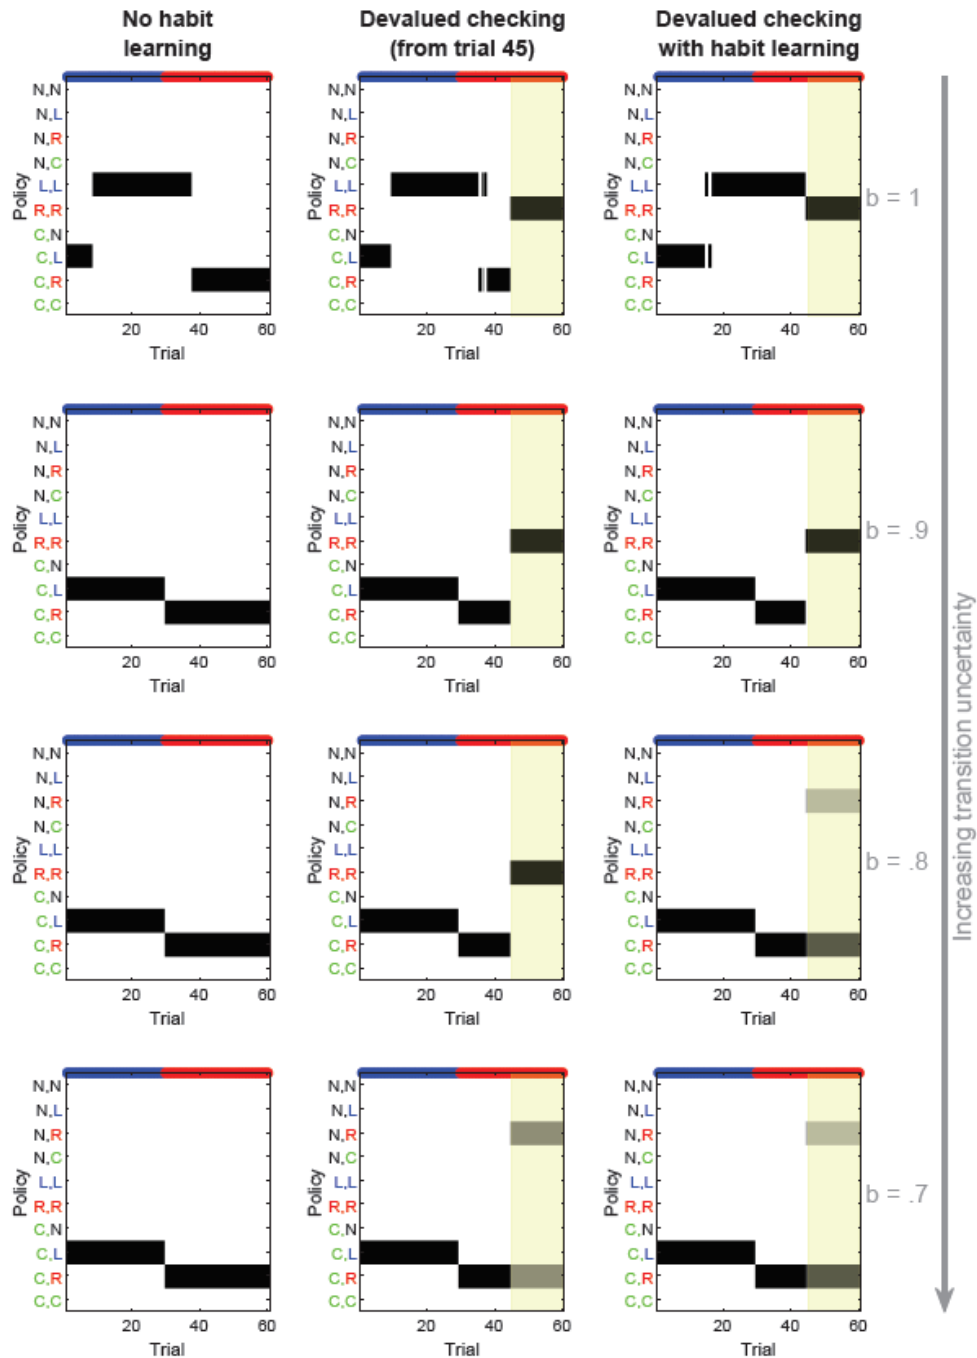


Figure 7 – Probabilities for chosen policies in a reversal learning task with two cues (L – left, R – right) and an additional explicit 'checking' cue (C; see Figure 4B), when increasing transition uncertainty (decreasing b), under an active inference MDP scheme. On each trial the agent chooses a two-step policy, corresponding with two sequential actions. Each action (within a policy) can be one of the following: choosing the left or right cue, choosing the checking cue, or doing nothing (N). Note that at some instances the simulation shows that several alternative policies can be chosen (depicted by grayscale rectangles representing the relative probability of these policies). The left and middle panels depict the results of this simulation when not allowing the agent to learn habits. In the middle and right panels, the epistemic value of checking behavior is devalued from trial 45 onwards, because from this point the agent is exposed to the current context at the beginning of each trial. In the right panel, habit learning is allowed. It can be seen that increased transition uncertainty leads to increased checking behavior. When habits are allowed to form, habitual checking that has no epistemic value emerges.

higher levels – not presented here) never resulted in strictly exploratory policies (i.e. exclusive checking behavior – where the agent checks in both time-steps, depicted as 'CC' policies in Figure 9). In contrast, such behavior was evident when either increasing transition uncertainty to extreme levels (e.g., $b = 0.51$; not presented here), or when increasing both harm-avoidance and transition uncertainty (Figure 9, bottom row). Together, this suggests that although it could clearly contribute to some symptoms in OCD (including the preponderance of harm-related obsessions), harm-avoidance alone cannot explain findings of both more volatile behaviors and increased habitual behaviors in reversal learning tasks and daily life.

To summarize, in this section we demonstrate how, under active inference, high transition uncertainty entails many of the kinds of behavior observed in OCD patients both in the lab (e.g., reversal learning experiments) and in day-to-day life (i.e. compulsions).

Compulsions can be conceptualized as responses to obsessions or NJREs, as excessive information gathering behavior driven by doubts and uncertainty, and as exaggerated habitual responding. Using simulations, we have shown how transition uncertainty leads to indecisiveness and volatile behavior that can be replaced by inflexible (habitual) behavior when the setting allows the agent to use simple behavioral strategies (e.g., to form habits). These simulations focused on action-independent transition uncertainty, because it plays a central role in tasks measuring inflexibility and overreliance on habits by introducing action-independent contingency shifts (e.g., reversal learning tasks). However, action-dependent transition uncertainty is also expected to cause a difficulty in planning that may be avoided by relying on habits and other heuristic strategies. Note that, whether action dependent or not, transitions describe a mapping from states to states. Action becomes involved through contextualizing this mapping.

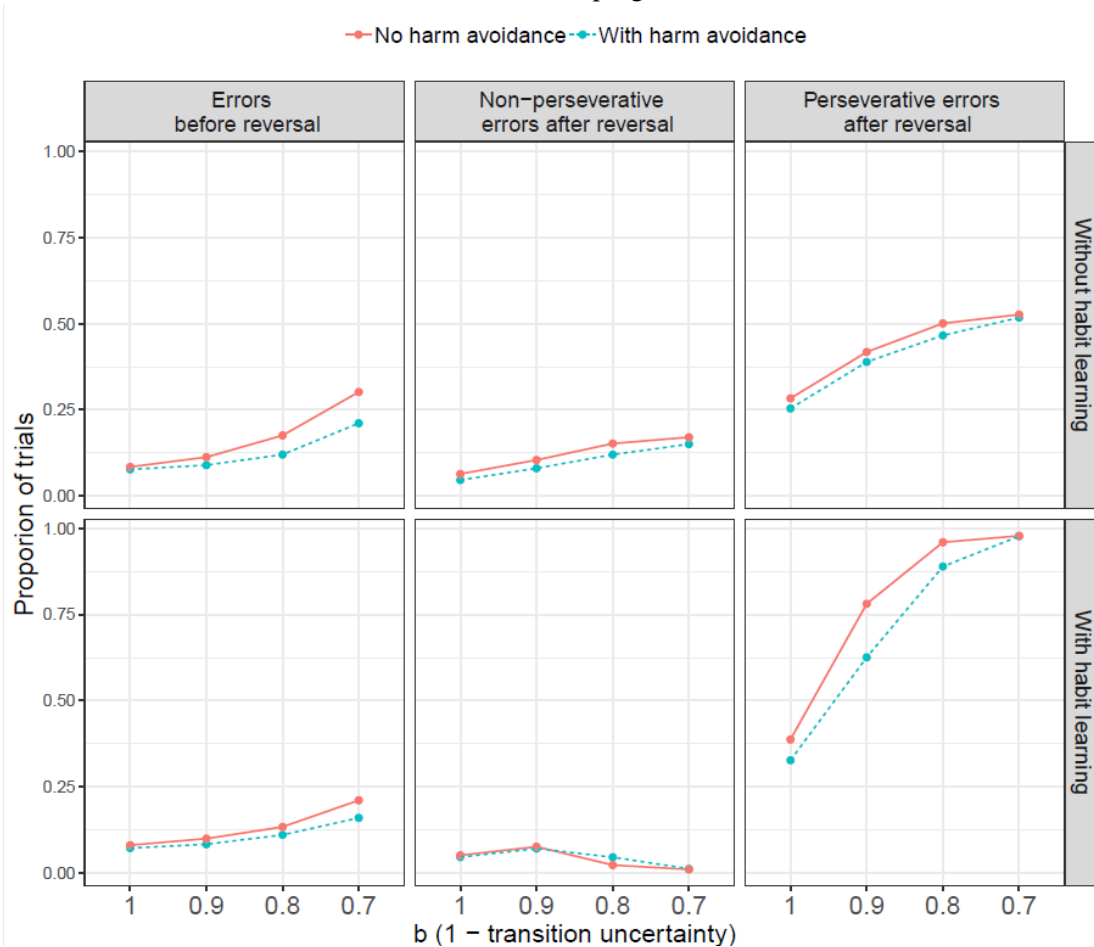


Figure 8 – The effects of harm-avoidance and different levels of transition uncertainty (as b decreases, transition uncertainty increases) on the types of errors agents are expected (averaged over 500 simulations) to make in a reversal learning task (see Figure 6). Harm avoidance counteracts the effects of transition uncertainty seen in Figure 6: i.e. it almost never leads to an increase in errors in this task.

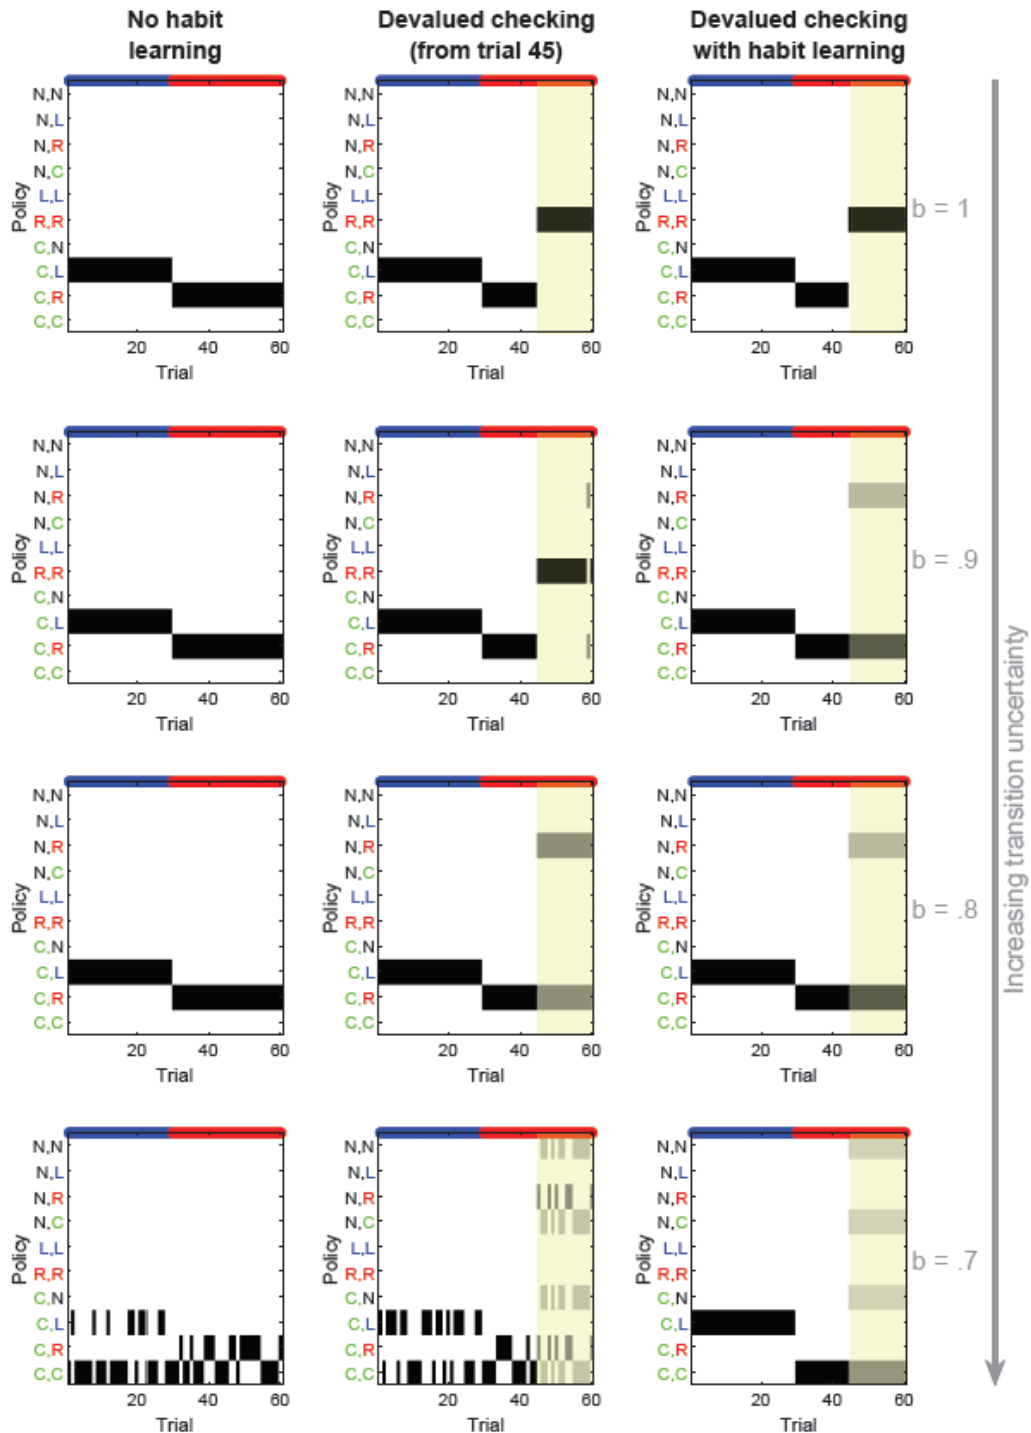


Figure 9 – Probabilities for chosen policies in a reversal learning task with two cues (L – left, R – right) and an additional explicit 'checking' cue (C; see Figure 4B), for a harm-avoidant agent, with a concurrent increase in transition uncertainty. These results should be compared to Figure 7, which does not include harm-avoidance. In this task, the effects of harm-avoidance are consistent with those of transition uncertainty, whereas the combination of the two results in exclusive checking behavior (i.e. C,C policies).

We also show that traditional accounts of OCD focusing on harm-avoidance cannot explain poor performance in reversal learning tasks. However, in real-life

situations (which are often similar to the second task), the effects of transition uncertainty on behavior will be magnified in potentially harmful situations, where the

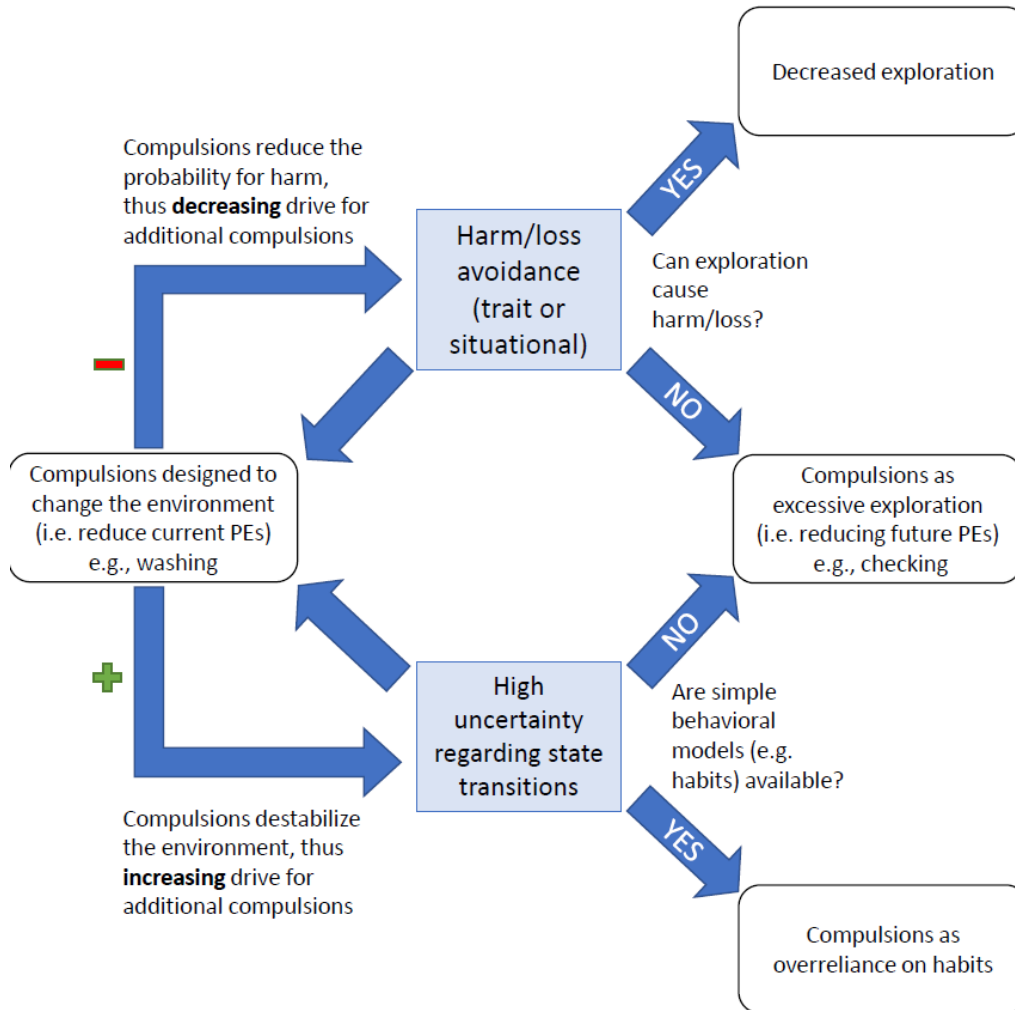


Figure 10 – A graphical summary of the postulated mechanisms leading to different classes of compulsions. Particularly the figure delineates the interactions between transition uncertainty, harm-avoidance, and overreliance on habits as predicted by the simulations.

need to avoid potentially negative transitions is emphasized. Figure 10 summarizes the predictions ensuing from the results of this simulation (and the hand-washing simulation).

High transition uncertainty can explain abnormalities in the processing of sensory, thought-related and reward prediction errors in OCD.

In this section we review previous findings related directly to the processing of PEs in OCD. Above we focused mainly on the effects of transition uncertainty on the planning and perception of actions, including the processing of action-related PEs. Here, we examine whether OCD is characterized by an increased weighting of PEs across different domains, and

whether this can contribute to the understanding of OCD phenomenology that does not involve action, including sensory over-responsiveness and intrusive thoughts.

Sensory prediction errors in OCD. At the phenomenological level, increased PE weight predicts increased behavioral and emotional responses to mundane sensory stimuli (i.e. *sensory over-responsiveness*). Indeed, numerous studies have shown sensory over-responsiveness (and related phenomena) in OCD and OC-related disorders (Ben-Sasson, Dickstein, Lazarovich, & Ayalon, 2017; Ben-Sasson & Podoly, 2017; Dar, Kahn, & Carmeli, 2012; Lewin, Wu, Murphy, & Storch, 2015; Wu, Lewin, Murphy, & Storch, 2014).

Critically, sensory over-responsiveness per se can result from various causes. More direct evidence for an

impairment in predictive processes in perception comes from studies demonstrating impaired sensory gating in OCD (Ahmari, Risbrough, Geyer, & Simpson, 2012; Hashimoto et al., 2008; Hoenig, Hochrein, Quednow, Maier, & Wagner, 2005; Swerdlow, Benbow, Zisook, Geyer, & Braff, 1993; but see de Leeuw, Oranje, van Megen, Kemner, & Westenberg, 2010). ‘Sensory gating’ is a historical term for the finding that predictable stimuli evoke weaker behavioral, attentional and neuronal responses than unpredictable stimuli. Impaired sensory gating was operationalized in the above studies as increased electrophysiological (e.g., P50) or motor (e.g., startle reflex in a prepulse inhibition paradigm) responses to predictable stimuli. These findings are consistent with the idea that sensory over-responsiveness in OCD is the result of increased weighting of sensory PEs.

Although imprecise predictions make predictable stimuli more salient, they also make unpredictable stimuli relatively less salient (Adams, Stephan, Brown, Frith, & Friston, 2013; Friston, 2005). This is often measured by ERP components sensitive to the violation of regularities, e.g. the mismatch negativity (MMN), and the P300. Only two studies have examined the MMN in OCD, with no evidence for a difference (from controls) in amplitude, but some evidence for a different topography (Oades, Zerbin, Dittmann-Balcar, & Eggers, 1996; Towey et al., 1994). The P300 has been examined in a larger number of studies, with most of them demonstrating lower P300 amplitudes for oddballs in OCD (Kim et al., 2003; Malloy, Rasmussen, Braden, & Haier, 1989; Sanz, Molina, Martin-Loeches, Calcedo, & Rubia, 2001; Towey et al., 1994; but see Mavrogiorgou et al., 2002). Interestingly, in two of these studies, OCD patients exhibited higher P300 amplitudes for regular stimuli (Malloy et al., 1989; Towey et al., 1994), replicating the effects regarding sensory gating discussed above.

Another line of evidence consistent with this hypothesis was reported in a preliminary, but intriguing, study (Ray Li et al., 2000) using a binocular rivalry paradigm. A different object was projected to each eye, and participants were asked to report what they perceived. In such experiments, participants usually report *perceptual switching* (i.e., alternations between perceiving the two objects). Bayesian models explain perceptual switching as resulting from the brain's attempts to choose the best model (i.e. posterior distribution) that minimizes PEs (i.e. explains the data), and that due to the contradictory sensory evidence, considerable PEs remain for each model (Weilnhammer, Stuke, Hessel-

mann, Sterzer, & Schmack, 2017). This study found increased perceptual switching in OCD (Ray Li et al., 2000), which is predicted by increased weighting of PEs.

Intrusive thoughts as prediction errors. Many obsessions have no relationship with sensory input, instead corresponding with thoughts, urges or images (sometimes called autogenous obsessions; Lee, Kwon, Kwon, J., & Telch, 2005). Interestingly, several theories have suggested that predictive processes might be involved in our trains of thought and how we experience them (Gallagher, 2000; Martin & Pacherie, 2013; Sterzer, Mishara, Voss, & Heinz, 2016). Indeed, several studies have shown that OCD is related with experiencing one's thoughts as being less related to both the current external context (Julien, O'Connor, & Aardema, 2009; Seli, Risko, Purdon, & Smilek, 2016) and the previous train of thought (Fradkin, Eitam, Strauss, & Huppert, 2019b; Fradkin & Huppert, 2018). In addition, a recent study has shown that participants with high obsessive compulsive symptoms can experience even predictable thoughts (where predictability was operationalized by manipulating semantic context) as intrusive (Fradkin et al., 2019b).

Thus, unpredictable transitions between thoughts may result in unusual thoughts (given a context) and thoughts experienced as intrusive (even if they are objectively predictable). Although promising, these results are limited by their subjective nature, and future research should attempt to devise neural and computational measures of PEs in thought processes.

Reward prediction errors in OCD. Reward PEs, referring to the response to feedback that is either worse or better than expected, can be conceptualized as another type of PE (which affects confidence regarding policies; FitzGerald, Dolan, & Friston, 2015). Consistent with our hypothesis, OCD patients exhibit increased reward PE signals in the anterior cingulate cortex (Hauser et al., 2017), and a stronger influence of reward PEs on action (Vaghi et al., 2017). Interestingly, in the latter study, the authors report that patients acted as if the contingencies had to be learned anew on each trial. As discussed above, such behavior indicates a belief that present states (and therefore outcomes) cannot be reliably predicted by past states – thus providing evidence for increased transition uncertainty (although this was reported only with regards to a latent 'Bayesian belief' driving action, but not in a model of meta-cognitive self-reported confidence).

Some studies have attempted to study reward PEs indirectly by focusing on a related ERP component, feedback related negativity (FRN). Results have been

inconsistent across studies (Endrass and Ullsperger, 2014), although the modulation of these effects by expectancies (as in the case of the P300 above) has not been examined to our knowledge. Another ERP component that has been studied intensively in OCD (and in other anxious groups) is the error related negativity (ERN), with most studies showing increase in OCD (Endrass and Ullsperger, 2014). Although the ERN is evoked by erroneous responses (and not by feedback) it has been argued to express a PE (Holroyd & Coles, 2002). However, ERP studies in OCD have yet to formalize the computational processes underlying these ERP signals. Such analysis is essential for the FRN and ERN which are related to multiple processes, involving the evaluation of the feedback valence and magnitude (Sambrook & Goslin, 2015), or conflict monitoring (Yeung, Botvinick, & Cohen, 2004), respectively.

Prediction errors and uncertainty. It is important to note that increased weighting of PEs could, theoretically, result also from underestimation of sensory noise. However, there is an important difference between this case and the case of increased transition uncertainty: Greater transition uncertainty leads to greater uncertainty in posterior beliefs, whereas an increase in sensory precision would lead to overconfidence. Consistent with the former proposition, there exists a plethora of clinical and experimental evidence supporting the notion of elevated uncertainty in OCD. Specifically, OCD patients have been shown to exhibit increased uncertainty in episodic memory (Cogle, Salkovskis, & Wahl, 2007; Hermans et al., 2008; Hermans, Martens, De Cort, Pieters, & Eelen, 2003), and other domains (Dar, 2004; Dar et al., 2000; Tuna, Tekcan, & Topçuoğlu, 2005; but see Tekcan, Topçuoğlu, & Kaya, 2007), greater decision thresholds (Banca et al., 2014) and indecisiveness (Frost & Shows, 1993). In that sense, excessive doubt, which can be defined as increased uncertainty regarding the current (or past) state of the world, is a natural consequence of increased transition uncertainty. Furthermore, indecisiveness in OCD has been associated with greater reliance on external feedback (Sarig et al., 2012), as expected in the case of increased weighting of PEs combined with greater uncertainty.

To summarize, we suggest that increased weighting of PEs in OCD explains findings of increased responses to predictable stimuli. Phenomenologically, it could account for sensory over-responsiveness in OCD, as well as the experience of intrusive thoughts. Importantly, overweighting of sensory information does not imply better (or less noisy) processing. On the contrary, depending on the context, overweighting of sensory data

often implies an impairment in processing, as it leads to a failure of the use of prior information to contextualize new data, and to better attenuate sensory noise. Finally, this increased weighting of PEs seems more likely to be the result of high transition uncertainty than underestimation of sensory noise. Indeed, in two studies (Ray Li et al., 2000; Vaghi et al., 2017), increased weighting of PEs was related to perceiving the world as highly unstable.

Discussion

In this theoretical paper, we suggest that obsessive compulsive pathology can be traced back, at the computational level, to excessive uncertainty regarding state transitions. This means that patients have a difficulty in predicting how events and sensations unfold (particularly but not exclusively as a consequence of their own actions). This makes the world appear less predictable and less controllable. More specifically, high transition uncertainty results in i) increased relative weighting of PEs – manifesting in patients' experiences that actions were not performed correctly, obsessional thoughts and sensory over-responsiveness, and ii) compulsions – conceptualized as responses to obsessions, excessive checking, and overreliance on habits (including habitual checking formalized as epistemic habits). The different effects of transition uncertainty on patients' symptoms and cognition are summarized in Figure 11.

We first demonstrated the effects of uncertainty regarding action-dependent transitions (using a particle filter) on increased reliance on noisy sensory feedback (potentially resulting in obsessions and NJREs), and decreased confidence (particularly, a paradoxical effect of compulsions on increasing uncertainty) during the estimation of current states (e.g., how clean my hands are). We then demonstrated the effects of transition uncertainty on the process of planning and action selection (using an active inference MDP model), leading to indecisiveness and increased information gathering, and a compensatory overreliance on habits. Compulsions are implicated as responses to current PEs and uncertainty in the former, and as a means of preventing future PEs in the latter. Interestingly, these two temporal orientations may correspond to the difference between compulsions driven by certain intrusive thoughts (i.e. autogenous obsessions) or NJREs (e.g. the experience that a specific action was not done right), and compulsions driven by attempts to avoid potential future harm. Importantly, these two apparently distinct modelling approaches appeal to the same (Bayesian)

Predicted effects of transition uncertainty on symptoms and cognition

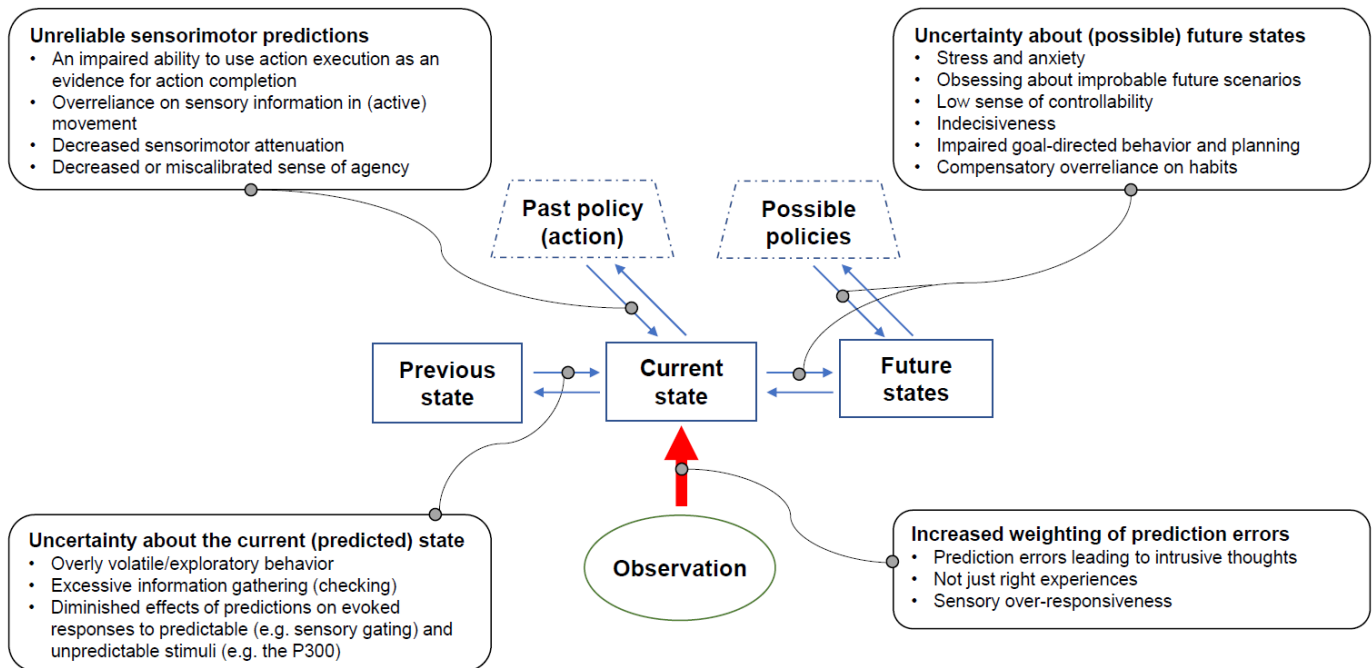


Figure 11 – A summary of the different postulated effects of transition uncertainty on OCD patients' symptoms, experience and cognition. As shown in Figure 1, excessive transition uncertainty affects several aspects of Bayesian inference. The current figure shows how these different aspects lead to the different findings and simulation results delineated in the paper.

generative modelling approach and examine the influence of the same parameter (transition uncertainty) in two different domains.

Notably, whereas previous accounts have suggested mild but wide-ranging learning impairments in OCD (e.g., Fradkin et al., 2018; Leplow, Murphy, & Nutting, 2002), our model and simulations predict specific types of errors in specific experimental contexts, due to the very specific computational quantity that we propose is altered (transition uncertainty). According to this model, OCD patients are expected to exhibit a pattern of increased vigilance and indecisiveness, resulting in errors of changing one's behavior (and over-exploration) even in the absence of objective contingency shifts. However, in highly stable environments, in which simple behavioral strategies (e.g., habits) can be temporarily adaptive, we predict a compensatory and inflexible over-reliance on such strategies (which should result in perseverative errors; see Figure 10). Moreover, given the hypothesized impairment in reliably predicting action-dependent transitions (even in stable environments), our model predicts that actions (e.g., compulsions) result in a paradoxical increase in uncertainty and thus lead to further compulsive action (see Figure 10). Notwithstanding this specificity, our

model can accommodate a gradient of behavioral impairments, as a function of both the level of transition uncertainty (which may vary across different individuals with OCD; see Figures 5-9 demonstrating such a gradual decrease in performance) and potential compensatory mechanisms (e.g., habits).

This perspective complements several previous models of OCD. First, prominent models highlight the role of impaired goal-directed control in OCD (e.g., Gillan et al., 2016). Our model is consistent with these ideas, while suggesting excessive transition uncertainty as the specific mechanism driving impaired goal-directed control, and delineating the task-related moderators explaining different behavioral consequences of this impairment. Thus, at the most basic level, patients experience fundamental indecisiveness, difficulty in predicting the consequences of potential behaviors, and over-sensitivity to new information. However, certain environments (e.g., stable environments or experimental contexts inducing over-training) allow patients to avoid the computational, cognitive and emotional costs of these impairments by relying on simpler strategies. In this sense, our model attempts to explain how seemingly incompatible behaviors (both indecisive and repetitive) can co-exist in OCD.

Although we focused on the case in which such simpler strategies involve habitual behavior (leading to perseverative behavior), the same principles can be conceptually extended to other types of heuristic behavioral strategies, such as model-free reinforcement learning schemes. In both cases, the difficulty in predicting the future leads patients to over-rely on the past. A related mechanism might be in play in the case of the development of compulsive rules (e.g., washing hands exactly 10 times). Using such a rule might help patients avoid goal-directed strategies (e.g., wash hands in a way that will lead to a certain reduction in estimated dirtiness). The development of the specific rule (i.e. number of times) might reflect a habitual policy that acquires a specific number of (for example) hand-washes simply because it was done previously, and this policy (and number) is reinforced each time it occurs. Alternatively, we could see this as a higher level inference about the transition uncertainties. That is, when one is unsure that a single handwash will reduce estimated dirtiness, a policy of repeated handwashes can be used to obtain sufficient certainty *a priori*.

Second, the ‘seeking proxies for internal states’ (SPIS) model (Lazarov et al., 2014) focuses on patients’ increased weighting of objective sensory evidence, at the expense of internal information. Specifically, these authors argue that OCD patients have attenuated access to, and greater uncertainty regarding ‘internal states’, including perception, memory, emotional states, muscle tension, proprioception and more. The SPIS model suggests that patients seek external proxies (e.g., environmental stimuli, rules or behaviors) to compensate for this uncertainty. Whereas we agree that patients rely more on external information, our model suggests a different reason underlying this phenomenon. We propose that the internal information patients have trouble relying upon corresponds with predictions regarding state transitions (including predictions regarding external states). Therefore, the proposed model builds upon the SPIS model while focusing on a specific computational impairment based on the well-established Bayesian brain framework. This creates some divergence in predictions from the SPIS model. For example, proprioception is considered an (attenuated) internal state in the SPIS model, but a (highly weighted) source of sensory information in our model. Paradigms comparing patients’ performance in active vs. passive movement (the latter weakening action-dependent predictions, yet leaving proprioceptive feedback intact) can be used to support this divergent prediction (Ezrati et al., 2018).

Third, Szechtman and Woody (2004) suggested that repetitive behavior in OCD stems from an inability to

utilize signals for the completion of tasks related to wellbeing and security. The authors also speculate that because many OCD-related behaviors have no objective satiety signal (e.g., hands are clean enough), normally such information relies on the execution of action itself. Whereas we agree that action execution serves as information for action completion (a process that is disrupted in OCD), our account bases this idea on established Bayesian accounts of action perception. Indeed, a related account linking the Szechtman and Woody (2004) model to impaired Bayesian inference (and the overweighting of improbable hypothetical, threat-related scenarios) has been suggested by Moore (2015). In contrast to these accounts, we suggest a domain-general core impairment, whereas prominent symptoms are expressed in specific domains that carry evolutionary or personal importance (see below).

Fourth, recently, Levy (2018) has suggested that OCD is related to increased precision for sensory PEs and actions, leading to increased attention to normally automatic processes. Whereas we concur with the general idea, we argue that this is the specific result of excessive uncertainty regarding state transitions.

Finally, classic cognitive behavioral models of OCD focus on the effects of anxiety (following obsessions) on motivating compulsions, which are reinforced by anxiety reduction (Foa & Kozak, 1996b). Whereas we address the emergence of anxiety in OCD below, we emphasize here that we view transition uncertainty as a basic vulnerability factor that likely interacts with other key factors (e.g. negative reinforcement of compulsions due to anxiety reduction) contributing to the clinical presentation.

Cognitive behavioral models of OCD also focus on various high-level meta-cognitive beliefs regarding obsessions and compulsions such as inflated responsibility, thought-action fusion and intolerance of uncertainty. We believe an important future step is to theoretically and empirically examine the potential link between excessive transition uncertainty and the emergence of such beliefs (see Levy, 2018, who discusses some potential directions). For example, intolerance of uncertainty is a potentially multifactorial concept that could result from increased intolerance of typical levels of uncertainty (e.g. about threat), and/or typical intolerance of increased levels of uncertainty (in any one of multiple domains). Thus, intolerance of uncertainty in OCD can be the result of increased transition uncertainty. Indeed, several studies have demonstrated that OCD patients show both less consistent behavior (as predicted in the case of increased transition uncertainty), and increased avoidance (implying intolerance)

of ambiguous choices (characterized by uncertainty regarding states and state transitions; Pushkarskaya et al., 2017, 2015). Furthermore, due to its potentially multifactorial nature, intolerance of uncertainty may result from different mechanisms in other disorders.

Uncertain state transitions, anxiety and harm-avoidance in OCD

OCD has been long considered an anxiety disorder, and despite its recent exclusion from this category (APA, 2013), it is hard to deny the involvement of stress, anxiety and threat-related contents in many OCD subtypes. The link between stress, anxiety and transition uncertainty is likely not coincidental, as unpredictable environments have been long known to precipitate anxiety (Grupe & Nitschke, 2013). Indeed, a recent account of stress suggests that "stress occurs if we are surprised by our sensations, and we are uncertain about what to do to safeguard our physical, mental and social wellbeing" (Peters, McEwen, & Friston, 2017, p. 184). Clearly, exaggerated uncertainty regarding state transitions will increase one's uncertainty regarding the optimal strategies to ensure wellbeing. Thus, although high transition uncertainty is proposed here to be a basic vulnerability factor, even manifested in non-threatening circumstances, its behavioral and affective consequences should be especially prominent in contexts relevant for survival and wellbeing. Indeed, the results of the simulations incorporating both harm-avoidance and transition uncertainty have shown that in some cases (when checking itself is not dangerous) harm-avoidance, here interpreted as a moderating contextual variable (i.e. situations potentially involving serious harm), amplifies the effects of transition uncertainty: leading to more instances of excessive checking and exploration (see Figure 10). Therefore, it is not surprising that OCD symptoms are usually (although not always) related to negative, threat-related content.

The common clinical focus of OCD on threat-related contents also raises the question of whether excessive perception of threat or excessive trait harm-avoidance can explain OCD more parsimoniously than our model. We used the simulations above to investigate this question, operationalizing harm-avoidance as an excessive need to avoid either punishments or contamination, or as a negatively biased (instead of more uncertain) perception of reality. These simulations show that harm-avoidance can explain repetitive washing and checking but cannot explain other important findings in OCD. These include excessive doubt and

uncertainty that are paradoxically increased after action, increased epistemic habits (i.e. continuing to check even when there is no information to gain), as well as over-exploratory and habitual responses in reversal learning tasks. Harm-avoidance also cannot explain other findings not included in the simulations, such as overreliance on exteroception during a muscle tension reproduction task, greater information gathering in non-threatening tasks, and greater perceptual switching. However, this does not preclude that possibility that both transition uncertainty and harm-avoidance are elevated in OCD. Interestingly, whereas in some cases these two factors had an additive effect on checking behavior, in other cases (e.g. when checking itself had a cost), increased harm-avoidance counteracted the effects of transition uncertainty. This generates specific, testable predictions regarding the possible interaction between these two factors, which need to be tested in future studies (see Figure 10).

Although harm-avoidance cannot fully explain OCD phenomenology, the agents in the simulations were assumed to have the goals of having clean hands or obtaining rewards. In active inference such goals take the form of high-level priors. Thus, whereas transition uncertainty explains the forms of behaviors observed in OCD, high-level priors can help explain their specific contents. Common themes of many compulsions (e.g. washing) could be explained by priors that have arisen through evolution, due to their value in promoting survival (Szechtman & Woody, 2004). Furthermore, some predetermined policies (i.e. habits) to threatening stimuli could themselves be determined (in terms of the prior probability of these policies) via processes developed during evolution (see Boyer and Liénard, 2006). Other more idiosyncratic symptoms could be explained by priors acquired through personal experience (e.g., Huppert and Zlotnick, 2012).

Whereas we have outlined evidence supporting our prediction that elevated transition uncertainty should be evident in neutral in addition to threatening (symptom-related) contexts, a non-intuitive prediction of our model is that OCD patients might also overestimate the probability of transitioning to positive states. Experimentally, this could relate to the interaction between transition uncertainty and task domain (e.g., rewards vs. losses). The literature is inconsistent in regard to whether over-exploration and habitualization in learning tasks in OCD appears in reward conditions, loss conditions or both (e.g., Hauser et al., 2017; Pushkarskaya et al., 2017), and research investigating transition uncertainty in different task domains is needed. Illusions of control (e.g., Reuven-Magril et al., 2008)

may also be related to being overly optimistic regarding action-dependent transitions. Clinically, such a mechanism could help account for magical thinking, which is elevated in OCD, and the persistent belief that a ritual will work, despite not having a clear relationship with the expected results. These speculations could be tested in future experiments, that emphasize positive (yet unpredictable) state transitions (e.g., by artificially resulting in unpredictable outcomes matching participants' goals; see Metcalfe and Greene, 2007).

Finally, an interesting and as yet unresolved issue concerns the specificity of these ideas of OCD compared to anxiety disorders, such as generalized anxiety disorder, which have been argued to result from uncertainty regarding future threats (Grupe & Nitschke, 2013). Indeed, a compelling reinforcement learning model of anxiety proposes that anxiety itself could result from pessimistic resolution of various forms of uncertainty: about future action choices, environmental state transitions, or the current state (Zorowitz, Momennejad, & Daw, 2019). We hypothesize that excessive uncertainty in anxiety disorders is less related to action-dependent state transitions (i.e. a difficulty in predicting the consequences of one's actions), which is a key component in our model of OCD. Furthermore, as we show in the hand-washing simulations, such an impairment leads to a paradoxical increase in uncertainty and additional compulsive urges following (compulsive) actions – a specific characteristic of OCD. A second distinction, partially related, is that in our model of OCD, transition uncertainty relates to both threatening and neutral states, whereas anxiety is more specifically related to threatening states. Direct comparisons of performance in experiments with patients with OCD vs. those with anxiety disorders support these distinctions. For example, excessive checking in neutral tasks was found only in OCD but not in anxiety (Toffolo et al., 2016). Similarly, impairments in predictive processes in thoughts (both negative and neutral) was found in high OC participants, but not in highly anxious participants (Fradkin et al., 2019b). Finally, increased weighting of exteroceptive evidence was found in OCD but not in anxiety disorders (Lazarov et al., 2014). However, future studies are needed to test the hypotheses we raise here.

Future directions

Having established the face validity of the proposed model in relation to the empirical data and phenomenology of OCD, we hope to characterize subject-specific phenotypes by estimating the parameters of the

model from behavioral data (e.g., reversal learning tasks). Crucially, we would estimate not just the transition uncertainty but a range of other, potentially interacting parameters. These could include prior preferences (e.g., harm-avoidance), prior beliefs about policies (e.g., habits) and sensory uncertainty. By correlating these parameter estimates with individual symptoms and cognitions (e.g, transition uncertainty as a global impairment that interacts with increased precision of losses – i.e., harm avoidance – in specific OCD subtypes, such as checking and washing), we hope to be able to understand how the interaction between subject-specific beliefs and changes in transition uncertainty induce a range of OCD phenotypes (and related phenomena, such as excessive doubt and intolerance of uncertainty).

Transition uncertainty was defined here as one's (in)ability to predict present and future states from actions or past states. Whereas we suggest that both types of transitions are less predictable in OCD, we also speculated that action-dependent transitions might play a more specific role. Future research should examine the contribution of these two processes to OCD pathology and anxiety. One way of investigating this question is to use a reversal learning task in which actions do not change the environment (e.g., Yu and Dayan, 2005). Alternatively, one can examine patients' expectancies when observing a different agent performing such a task.

The ideas presented here deal with several domains of active inference, and touch upon the complex, hierarchical nature of the Bayesian brain. These include the sorts of generative models required for continuous motor-control, and those needed to plan over successive time-steps, and deal with sequences in time. While these are different sorts of inference, they can be unified through the use of hierarchical generative models (for recent examples of this, see Friston, Parr, & Vries, 2017; Parr & Friston, 2018). These accounts rest upon the idea that the brain may represent dynamics over time as a sequence of short trajectories. Uncertainty regarding transitions at one level likely influence the other levels, and PEs at lower levels can propagate to higher levels. Indeed, from a clinical perspective, compulsions themselves range from a repetition of short, basic sub-actions (e.g., repeating simple movements due to a not-just-right experience) to more elaborate behavioral sequences (e.g., repeating a full sentence, or a relatively complex washing behavior). Nonetheless, future work is needed to better localize the computational pathology in the Bayesian hierarchy.

Neural predictions. Classical neuroanatomical models of OCD implicate the cortico-striato-thalamic-cortical circuit (CSTC), and specifically the orbitofrontal cortex (OFC), anterior cingulate cortex (ACC) and the striatum, although other areas are involved as well (e.g., the dorsolateral prefrontal cortex, the cerebellum, amygdala; for reviews see Chamberlain & Menzies, 2009; Menzies et al., 2008; Milad & Rauch, 2012). Interestingly, recent findings and models suggest that the OFC has a pivotal role in the representation of hidden states (Behrens et al., 2018; Hampton, Bossaerts, & O’Doherty, 2006), particularly those that are related to current task or goals or have a significant biological meaning (Parr, Rikhye, Halassa, & Friston, 2019; Schuck, Cai, Wilson, & Niv, 2016; Wikenheiser & Schoenbaum, 2016). Furthermore, a recent study has shown that rapid changes in firing patterns in the OFC are consistent with the representation of transitions between hidden states during a probabilistic learning task (Nassar, Bruckner, & Frank, 2019). The striatum, in turn, is thought to be associated with the selection and regulation of behavioral policies, and the integration of goal-directed and habitual policies (Graybiel & Grafton, 2015; Parr & Friston, 2018).

The involvement of the OFC and the striatum in the coding of hidden states and policies, respectively, suggests that CSTC abnormalities in OCD might correspond to the types of computations we focus on in our model. However, previous theoretical approaches explained these same CSTC abnormalities as pertaining solely to an imbalance between goal-directed and habitual behavior (Gillan & Robbins, 2014). To directly examine the biological plausibility of our model, future studies should focus more specifically on potential neuroanatomical markers of state transitions in OCD. One direction would be to follow Nassar and colleagues (2019) by examining rapid fluctuations in BOLD activity in the OFC (using representational similarity analysis) as a marker of state transitions. Another direction could be to search for a more stable marker of the uncertainty of state transitions. One caveat here is that other than evidence for the involvement of noradrenaline in the coding of volatility (Yu & Dayan, 2005), the neuroanatomical substrates of transition uncertainty are not well understood. One plausible idea is that if OFC activity represents hidden states, OFC intrinsic connectivity may correspond with characteristics of the distribution of the relationship between states, i.e., state transitions. Indeed, a recent study demonstrated a loss of numerous excitatory synaptic protein transcripts in OFC in post mortem OCD patients (Piantadosi, Chamberlain, Glausier, Lewis, & Ahmari, 2019): a potential

cause of transition uncertainty. A different direction is to examine whether manipulations of OCD-related circuitry results in the types of behavior seen in the simulations above. In a seminal study, Ahmari and colleagues (2013) reported that repeated stimulation of the CSTC circuit in mice resulted in 'compulsive-like' perseverative grooming. Interestingly, acute stimulation resulted in an excessive locomotion, which might reflect increased exploration, but further work is required to establish this.

Impairments in Bayesian inference in other disorders.

Bayesian models have been suggested to account for various psychiatric disorders, including psychosis (Adams et al., 2013; Fletcher & Frith, 2009; Sterzer et al., 2018) autism (ASD; Cruys et al., 2014; Lawson, Rees, & Friston, 2014; Palmer, Lawson, & Hohwy, 2017; Pellicano & Burr, 2012), and tic disorders (Rae, Critchley, & Seth, 2019). This is not surprising: if the brain performs (or approximates) Bayesian inference, many neuropsychiatric disorders will likely be explicable in Bayesian terms. The deeper question is whether these computational models can show how the diverse but partially overlapping risk factors of these disorders (O’Connell, McGregor, Lochner, Emsley, & Warnich, 2018) lead to both shared and distinct abnormalities, generating disorder-specific experimental and computational predictions.

Autistic spectrum disorders. Several theoretical accounts (for a review see Palmer et al., 2017), suggest that abnormal sensory processing in ASD results from overweighting of sensory PEs at the expense of priors (Cruys et al., 2014; Pellicano & Burr, 2012), possibly explained by excessive estimates of volatility (or meta-volatility, i.e. uncertainty regarding how unstable state transitions are; Lawson et al., 2017). However, a more recent study has shown reduced weighting of more recent evidence in ASD, suggesting reduced estimates of volatility (Lieder et al., 2019).

Along with this similarity, one important distinction is that current models and findings in ASD do not touch upon action-dependent transition uncertainty, but rather focus on Bayesian inference in general (e.g., Pellicano & Burr, 2012), or volatility estimates (e.g., Lawson et al., 2017). Furthermore, ASD is a developmental disorder which is characterized by social impairment. Indeed, Bayesian inference impairments in ASD have been suggested to stem from a more specific impairment in the interplay between social cues and the pre-

cision of interoceptive PEs (Atzil, Gao, Fradkin, & Barrett, 2018; Quattrocki & Friston, 2014). Relatedly, a plausible hypothesis is that when Bayesian inference impairments originate at a very early developmental stage (as it likely occurs in ASD), a different clinical picture will emerge. However, there are yet no data or simulations regarding the way in which the onset of impairments in Bayesian inference (e.g., transition uncertainty) influences the resulting clinical picture.

Schizophrenia. Perhaps the first application of Bayesian inference models to psychopathology was in the context of psychotic symptoms (for a review see Sterzer et al., 2018). We have recently suggested that the computational differences between OCD and schizophrenia may be quantitative in their basis (Fradkin, Eitam, Strauss, & Huppert, 2019a). In OCD, the imbalance between sensory and higher-level precision might be not as great as in schizophrenia, such that sensory PEs (and the unease they induce) can be resolved by actions or beliefs regarding the stability or controllability of the environment, rather than by changing more fundamental beliefs about the self and reality (in schizophrenia; see Corlett et al., 2018; Fradkin et al., 2019a). This suggests that the distinctions between psychiatric disorders might be related to differences in the type and level of priors used to explain away highly weighted sensory PEs. This hypothesis could be tested empirically by experimentally manipulating higher level priors. For example, in a recent study from our group (Fradkin et al., 2019b), task-instructions were used to induce a prior belief regarding thought insertion. This caused non-psychotic, high OC participants to interpret their excessive thought-related PEs as inserted thoughts.

Other fronto-striatal disorders. Recent simulation work has shown that transition uncertainty is also one potential cause of impulsive behavior, specifically delay discounting (Mirza, Adams, Parr, & Friston, 2019). Impulsivity is one of the core features of attention deficit hyperactivity disorder (ADHD), which also exhibits partial epidemiological and neurobiological overlap with OCD (Brem, Grünblatt, Drechsler, Riederer, & Walitza, 2014). To our knowledge, there is no work on Bayesian inference in ADHD. However, increased behavioral variability and exploration have been implicated in a recent model (Hauser, Fiore, Moutoussis, & Dolan, 2016). Interestingly, recent computational work has proposed that both OCD and ADHD patients exhibit increased exploration and increased weighting of the most recent outcome, whereas only ADHD patients exhibit increased reward sensitivity (Norman et al.,

2018). This might help explain why ADHD is characterized by increased risk-taking and novelty seeking, and OCD with risk avoidance. That is, a difficulty in predicting the consequences of actions should lead to more risk taking and distractibility when one assigns higher importance or probability to rewards than to losses.

Tic disorders (TDs), showing considerable comorbidity with OCD (Sheppard, Bradshaw, Purcell, & Pantelis, 1999), have been recently suggested to involve abnormal precision over policies, leading to the formation of excessive stereotyped, habitual movements (tics; Rae et al., 2019). Common to both disorders may be an imbalance of precision between goal-directed policies, and habits. Where the disorders are distinct, the primary pathology may differ: increased 'movement' precision in TDs, versus transition uncertainty in OCD (which affects policies but has additional, cognitive and affective components).

Implications for treatment.

Formalizing OCD under the active inference framework provides an opportunity to consider computationally-informed targets for interventions. Notably, although we suggest a core pathology in transition uncertainty, the alleviation of the phenomenological effects of this pathology can be produced by focusing on other, interacting components. One such target is the subjective importance (i.e. precision) of prior preferences and goals (Peters et al., 2017). So, for example, having more flexible criteria for cleanliness is expected to alleviate the effects of transition uncertainty simulated above. Similarly, the MDP simulations have shown that trait or state increases in the subjective cost of harm strengthen the effect of transition uncertainty. Another target concerns the mitigation of the downward effects of transition uncertainty on shifting the balance between goal-directed and habitual behavior. Theoretically this can involve either interfering with habit learning (computationally, this corresponds with increasing the precision of the E matrix - prior beliefs over policies), or increasing the precision of goal-directed policies.

The current gold-standard treatments for OCD include cognitive-behavioral therapy (CBT; Foa and Kozak, 1996), and SSRIs (Pizarro et al., 2014). Treatment augmentation with atypical antipsychotics (Veale et al., 2014) has become increasingly common. It is tempting to speculate that different treatments might influence different targets. For example, SSRIs might play a role in regulating the perceived threat resulting

from high transition uncertainty. Correspondingly, SSRIs have been shown to be most effective in OCD subtypes in which harm-avoidance plays a major role (Landeros-Weisenberger et al., 2010). Furthermore, most studies have shown that SSRIs do not improve neuropsychological functioning related to perseveration or non-perseverative errors (Mataix-Cols, Alonso, Pifarre, Menchon, & Vallejo, 2002; Nielen & Den Boer, 2003; Roh et al., 2005). In contrast, dopamine has been linked to the encoding of the precision over policies, and the balance between goal-directed and habitual policies (Parr & Friston, 2018). This raises the hypothesis that dopaminergic medication may alleviate perseveration but not the core transition uncertainty impairment. Indeed, preliminary findings suggest that dopamine antagonist administration reduces perseveration in OCD (Ersche et al., 2011), but can also increase errors that are due to an instable response style (i.e., difficulty in maintaining a response set in the Wisconsin card sorting test; De Geus, Denys, & Westenberg, 2007).

CBT involves two cardinal components: exposure, often aimed at helping patients tolerate the risk and uncertainty associated with their obsessions, and response prevention, aiming directly at reducing compulsive rituals (Foa, Steketee, Grayson, Turner, & Latimer, 1984). Computationally, exposure may serve the role of regulating perceived threat, as well as the emotional effects of transition uncertainty (i.e., anxiety-provoking transitions). Response prevention is traditionally considered to be necessary for exposure to work effectively (Foa & Kozak, 1996b), but it can also serve a role in decreasing the prior probability of habitual compulsions (similarly to habit-reversal treatment in OC-related disorders; Gillan & Robbins, 2014). If indeed a partial dissociation exists in the computational parameters affected by these two components, computational phenotyping might help guide therapeutic focus.

Finally, an interesting yet difficult question is whether difficulties in transition uncertainty can be remedied directly. Preliminary evidence suggests that errors indicating excessive exploration or oversensitivity to feedback were not improved by existing treatments, including dopaminergic agents (Ersche et al., 2011) or CBT (Freyer et al., 2011), and a more systematic research of this question is required. One possible direction involves the use of cognitive training programs aimed at reducing the biased estimation of transition uncertainty, along the lines that led to the development of cognitive remediation programs. This first requires the development and psychometric validation

of tasks that directly measure transition uncertainty in different contexts.

More generally, we believe that the model proposed here provides a novel perspective regarding the subjective experience of OCD patients. For example, we show how deficits in goal-directed control may be closely linked to self-reported indecisiveness and doubts. Furthermore, we propose a unitary account for explaining the diversity in the types of obsessions in OCD (e.g., intrusive thoughts vs. NJREs), as well as for understanding patients' different motivations to engage in compulsive behavior. We believe that this can enrich clinicians' understanding of their patients and provide a framework for new computational studies of OCD.

Appendix A – Glossary

State – a configuration of the environment at a given time

Hidden/latent state – a state that is not directly observable, but must be inferred from different sources of information – sensory observations and predictions regarding state transitions (see below).

State transitions – a (probabilistic) description of how states evolve over time. State transitions can be affected either by actions (action-dependent state transitions) or not (action-independent state transitions).

Transition uncertainty – uncertainty about the state transition probabilities, (i.e. being unsure whether things will change). When transition uncertainty is high, past states and/or actions are not predictive of present or future states.

Precision – inverse variance/uncertainty. High uncertainty is equivalent to low precision.

State-space model – a formal mathematical model defining how states relate to each other and to observations. The core problem of simple state space models is to infer hidden states by using these relations.

Bayesian state-space model – a state-space model where hidden states, state transitions, and the relations between states and observations are represented as probability distributions (characterized by their mean and variance/uncertainty), and where inference follows Bayes' theorem.

Particle filter – an algorithm that can solve Bayesian state space models (under minimal assumptions). Distributions of states, state transitions and observations are represented by 'particles', with each particle representing a sample (one possible value) from the respective distribution.

Bayesian brain – the idea that the brain uses Bayesian inference (or approximations of it) to integrate predictions (e.g., predictions regarding state transitions) and sensory observations, by weighting each according to its uncertainty. This process allows the brain to infer a probability distribution over hidden states.

Prediction errors / PEs – a quantifiable difference between the brain's prediction (e.g., perceptual prediction) and its actual input.

Predictive coding – a neurobiological framework describing how the brain implements Bayesian inference over states that change continuously (as in perception). It does this by passing messages (predictions and prediction errors) up and down a hierarchy of neuronal populations.

Active inference – an extension of Bayesian inference schemes (including predictive coding) to action, such that prediction errors can either be resolved by changing the model (perception) or the sensory input (action). Briefly, the brain uses Bayesian inference on expected (rather than actual) observations to select policies (see below). Movement itself can also be described as Bayesian inference on expected proprioceptive observations (e.g., how the posture of my hand should change to realize the sensations I expect, such as gripping a cup). Actions serve both to fulfill goals (utilitarian goals) and to minimize uncertainty (epistemic goals).

Policy – a sequence of actions

Markov decision process (MDP) – a mathematical framework for modeling decision making (i.e. policy selection). Active inference (see above) can be used to solve MDPs (i.e., to infer the optimal policy, while taking into account the relations between policies, states, and observations).

References

- Aardema, F., O'Connor, K. P., Pélissier, M.-C., & Lavoie, M. E. (2009). The Quantification of Doubt in Obsessive-Compulsive Disorder. *International Journal of Cognitive Therapy*, 2(2), 188–205. <https://doi.org/10.1521/ijct.2009.2.2.188>
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. J. (2013). The Computational Anatomy of Psychosis. *Frontiers in Psychiatry*, 4. <https://doi.org/10.3389/fpsy.2013.00047>
- Ahmari, S. E., Risbrough, V. B., Geyer, M. a, & Simpson, H. B. (2012). Impaired sensorimotor gating in unmedicated adults with obsessive-compulsive disorder. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology*, 37(5), 1216–23. <https://doi.org/10.1038/npp.2011.308>
- Ahmari, S. E., Spellman, T., Douglass, N. L., Kheirbek, M. A., Simpson, H. B., Deisseroth, K., ... Hen, R. (2013). Repeated cortico-striatal stimulation generates persistent OCD-like behavior. *Science*, 340(6137), 1234–1239.
- Association, A. P. (2013). *Diagnostic and Statistical Manual of Mental Disorders*. <https://doi.org/10.1176/appi.books.9780890425596>
- Atzil, S., Gao, W., Fradkin, I., & Barrett, L. F. (2018). Growing a social brain. *Nature Human Behaviour*, 1.
- Banca, P., Vestergaard, M. D., Rankov, V., Baek, K., Mitchell, S., Lapa, T., ... Voon, V. (2014). Evidence Accumulation in Obsessive-Compulsive Disorder: The Role of Uncertainty and Monetary Reward on Perceptual Decision-Making Thresholds. *Neuropsychopharmacology*, 40(5), 1192–1202. <https://doi.org/10.1038/npp.2014.303>
- Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, 100(2), 490–509.
- Ben-Sasson, A., Dickstein, N., Lazarovich, L., & Ayalon, N. (2017). Not Just Right Experiences: Association with Obsessive Compulsive Symptoms and Sensory Over-Responsivity. *Occupational Therapy in Mental Health*, 33(3), 217–234. <https://doi.org/10.1080/0164212X.2017.1303418>
- Ben-Sasson, A., & Podoly, T. Y. (2017). Sensory over responsivity and obsessive compulsive symptoms: A cluster analysis. *Comprehensive Psychiatry*, 73, 151–159. <https://doi.org/10.1016/j.comppsy.2016.10.013>
- Blakemore, S. J., Wolpert, D., & Frith, C. (2000). Why can't you tickle yourself? *Neuroreport*, 11(11), 11–16.
- Bogacz, R. (2017). A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*, 76, 198–211. <https://doi.org/10.1016/j.jmp.2015.11.003>
- Boyer, P., & Liénard, P. (2006). Why ritualized behavior? Precaution Systems and action parsing in developmental, pathological and cultural rituals. *Behavioral and Brain Sciences*, 29(6), 595–613. <https://doi.org/10.1017/S0140525X06009332>
- Brem, S., Grünblatt, E., Drechsler, R., Riederer, P., & Walitza, S. (2014). The neurobiological link between OCD and ADHD. *ADHD Attention Deficit and Hyperactivity Disorders*, 6(3), 175–202. <https://doi.org/10.1007/s12402-014-0146-x>
- Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81, 55–79. <https://doi.org/10.1016/j.jmp.2017.09.004>
- Chamberlain, S. R., & Menzies, L. (2009). Endophenotypes of obsessive-compulsive disorder: Rationale, evidence and future potential. *Expert Review of Neurotherapeutics*, 9(8), 1133–1146. <https://doi.org/10.1586/ern.09.36>
- Clair, A.-H., N'diaye, K., Baroukh, T., Pochon, J.-B., Morgieue, M., Hantouche, E., ... Mallet, L. (2013). Excessive checking for non-anxiogenic stimuli in obsessive-compulsive disorder. *European Psychiatry*, 28(8), 507–513.
- Coles, M. E., & Ravid, A. (2016). Clinical presentation of not-just right experiences (NJREs) in individuals with OCD: Characteristics and response to treatment. *Behaviour Research and Therapy*, 87, 182–187. <https://doi.org/10.1016/j.brat.2016.09.013>
- Corlett, P. R., Horga, G., Fletcher, P. C., Alderson-Day, B., Schmack, K., & Powers, A. R. (2018). Hallucinations and Strong Priors. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2018.12.001>
- Cougle, J. R., Salkovskis, P. M., & Wahl, K. (2007). Perception of memory ability and confidence in recollections in obsessive-compulsive checking. *Journal of Anxiety Disorders*, 21(1), 118–130. <https://doi.org/10.1016/j.janxdis.2006.03.015>

- Cruys, S. V. D., Evers, K., Hallen, R. V. D., Eylen, L. V., Boets, B., & Wagemans, J. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological Review*, *121*(4). <https://doi.org/10.1037/a0037665>
- Dar, R. (2004). Elucidating the mechanism of uncertainty and doubt in obsessive-compulsive checkers. *Journal of Behavior Therapy and Experimental Psychiatry*, *35*(2), 153–163. <https://doi.org/10.1016/j.jbtep.2004.04.006>
- Dar, R., Kahn, D. T., & Carmeli, R. (2012). The relationship between sensory processing, childhood rituals and obsessive-compulsive symptoms. *Journal of Behavior Therapy and Experimental Psychiatry*, *43*(1), 679–684. <https://doi.org/10.1016/j.jbtep.2011.09.008>
- Dar, R., Rish, S., Hermesh, H., Taub, M., & Fux, M. (2000). Realism of confidence in obsessive-compulsive checkers. *Journal of Abnormal Psychology*, *109*(4), 673–678. <https://doi.org/10.1037/0021-843X.109.4.673>
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711. <https://doi.org/10.1038/nn1560>
- De Geus, F., Denys, D., & Westenberg, H. G. M. (2007). Effects of quetiapine on cognitive functioning in obsessive-compulsive disorder. *Int Clin Psychopharmacol International Clinical Psychopharmacology*, *22*(22), 77–8477.
- de Leeuw, A. S., Oranje, B., van Meegen, H. J. G. M., Kemner, C., & Westenberg, H. G. M. (2010). Sensory gating and sensorimotor gating in medication-free obsessive-compulsive disorder patients. *International Clinical Psychopharmacology*, *25*(4), 232–240. <https://doi.org/10.1097/YIC.0b013e328338c4f0>
- Deckersbach, T., Savage, C. R., Curran, T., Bohne, A., Wilhelm, S., Baer, L., ... Rauch, S. L. (2002). A study of parallel implicit and explicit information processing in patients with obsessive-compulsive disorder. *American Journal of Psychiatry*, *159*(10), 1780–1782. <https://doi.org/10.1176/appi.ajp.159.10.1780>
- Dorfman, H. M., & Gershman, S. J. (2019). Controllability Governs the Balance Between Pavlovian and Instrumental Action Selection. *BioRxiv*, 596577.
- Edwards, M. J., Adams, R. A., Brown, H., Parees, I., & Friston, K. J. (2012). A Bayesian account of ‘hysteria.’ *Brain*, *135*(11), 3495–3512.
- Endrass, T., & Ullsperger, M. (2014). Specificity of performance monitoring changes in obsessive-compulsive disorder. *Neuroscience and Biobehavioral Reviews*. <https://doi.org/10.1016/j.neubiorev.2014.03.024>
- Ersche, K. D., Roiser, J. P., Abbott, S., Craig, K. J., Mller, U., Suckling, J., ... Bullmore, E. T. (2011). Response perseveration in stimulant dependence is associated with striatal dysfunction and can be ameliorated by a D 2/3 receptor agonist. *Biological Psychiatry*, *70*(8), 754–762. <https://doi.org/10.1016/j.biopsych.2011.06.033>
- Ezrati, O., Sherman, E., & Dar, R. (2018). High obsessive-compulsive individuals may have attenuated access to internal cues associated with active movement: Evidence from a head repositioning study. *Journal of Behavior Therapy and Experimental Psychiatry*, *60*(October 2017), 1–4. <https://doi.org/10.1016/j.jbtep.2018.02.004>
- FitzGerald, T. H. B., Dolan, R. J., & Friston, K. (2015). Dopamine, reward learning, and active inference. *Frontiers in Computational Neuroscience*, *9*(November), 1–16. <https://doi.org/10.3389/fncom.2015.00136>
- FitzGerald, T. H. B., Dolan, R. J., & Friston, K. J. (2014). Model averaging, optimal inference, and habit formation. *Frontiers in Human Neuroscience*, *8*(June), 1–11. <https://doi.org/10.3389/fnhum.2014.00457>
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: A Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews. Neuroscience*, *10*(1), 48–58. <https://doi.org/10.1038/nrn2536>
- Foa, E. B., & Kozak, M. J. (1995). DSM-IV field trial: Obsessive-compulsive disorder. *The American Journal of Psychiatry*, *152*(1), 90.
- Foa, E. B., & Kozak, M. J. (1996a). *Psychological treatment for obsessive-compulsive disorder*.
- Foa, E. B., & Kozak, M. J. (1996b). Psychological treatment for obsessive-compulsive disorder. In *Long-term treatments of anxiety disorders* (pp. 285–309). Arlington, VA, US: American Psychiatric Association.
- Foa, E. B., Steketee, G., Grayson, J. B., Turner, R. M., & Latimer, P. R. (1984). Deliberate exposure and blocking of obsessive-compulsive rituals: Immediate and long-term effects. *Behavior Therapy*, *15*(5), 450–472.
- Fradkin, I., Eitam, B., Strauss, A. Y., & Huppert, J. D. (2019a). How can an overlapping mechanism lead to distinct pathology? The case of psychosis and obsessive compulsive disorder. *Clinical Psychological Science*. <https://doi.org/10.1177/2167702619830393>
- Fradkin, I., Eitam, B., Strauss, A. Y., & Huppert, J. D. (2019b). Thoughts as Unexpected Intruders: Context, Obsessive-Compulsive Symptoms, and the Sense of Agency Over Thoughts. *Clinical Psychological Science*, *7*(1), 162–180. <https://doi.org/10.1177/2167702618797102>
- Fradkin, I., & Huppert, J. D. (2018). When our train of thought goes off track: The different facets of out-of-context thoughts in obsessive compulsive disorder. *Journal of Obsessive-Compulsive and Related Disorders*, *18*, 31–39. <https://doi.org/10.1016/J.JOCD.2018.06.001>
- Fradkin, I., Strauss, A. Y., Pereg, M., & Huppert, J. D. (2018). Rigidly Applied Rules? Revisiting Inflexibility in Obsessive Compulsive Disorder Using Multilevel Meta-Analysis. *Clinical Psychological Science*, *6*(4). <https://doi.org/10.1177/2167702618756069>
- Freyer, T., Klöppel, S., Tüscher, O., Kordon, A., Zurowski, B., Kuelz, A.-K., ... Voderholzer, U. (2011). Frontostriatal activation in patients with obsessive-compulsive disorder before and after cognitive behavioral therapy. *Psychological Medicine*, *41*(1), 207–216.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *360*(1456), 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews. Neuroscience*, *11*(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Friston, K. (2011). What is optimal about motor control? *Neuron*, *72*(3), 488–498. <https://doi.org/10.1016/j.neuron.2011.10.018>
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O’Doherty, J., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience and Biobehavioral Reviews*, *68*, 862–879. <https://doi.org/10.1016/j.neubiorev.2016.06.022>
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, *6*(4), 187–224. <https://doi.org/10.1080/17588928.2015.1020053>

- Frost, R. O., & Shows, D. L. (1993). The nature and measurement of compulsive indecisiveness. *Behaviour Research and Therapy*, 31(7), 683-692. [https://doi.org/10.1016/0005-7967\(93\)90121-A](https://doi.org/10.1016/0005-7967(93)90121-A)
- Gallagher, S. (2000). Self-Reference and Schizophrenia: A Cognitive Model of Immunity to Error through Misidentification. In D. Zahavi (Ed.), *Exploring the Self: Philosophical and Psychopathological Perspectives on Self-experience*, (p. (203-239)). <https://doi.org/10.1075/aicr.23.14gal>
- Gehring, W. J., Himle, J., & Nisenson, L. G. (2000). Action-Monitoring dysfunction in Obsessive Compulsive Disorder. *Psychological Science*, 11(1), 1-6.
- Gentsch, A., Schütz-Bosbach, S., Endrass, T., & Kathmann, N. (2012). Dysfunctional forward model mechanisms and aberrant sense of agency in obsessive-compulsive disorder. *Biological Psychiatry*, 71(7), 652-9. <https://doi.org/10.1016/j.biopsych.2011.12.022>
- Giele, C. L., Van Den Hout, M. A., Engelhard, I. M., Dek, E. C. P., & Hofmeijer, F. K. (2011). Obsessive-compulsive-like reasoning makes an unlikely catastrophe more credible. *Journal of Behavior Therapy and Experimental Psychiatry*, 42(3), 293-297. <https://doi.org/10.1016/j.jbtep.2010.12.012>
- Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *eLife*, 5(MARCH2016), 1-24. <https://doi.org/10.7554/eLife.11305>
- Gillan, C. M., Morein-Zamir, S., Urcelay, G. P., Sule, A., Voon, V., Apergis-Schoute, A. M., ... Robbins, T. W. (2014). Enhanced avoidance habits in obsessive-compulsive disorder. *Biological Psychiatry*, 75(8), 631-8. <https://doi.org/10.1016/j.biopsych.2013.02.002>
- Gillan, C. M., Pappmeyer, M., Morein-Zamir, S., Sahakian, B. J., Fineberg, N. A., Robbins, T. W., & De Wit, S. (2011). Disruption in the balance between goal-directed behavior and habit learning in obsessive-compulsive disorder. *American Journal of Psychiatry*, 168(7), 718-726. <https://doi.org/10.1176/appi.ajp.2011.10071062>
- Gillan, C. M., & Robbins, T. W. (2014). Goal-directed learning and obsessive-compulsive disorder. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 369, 20130475. <https://doi.org/10.1098/rstb.2013.0475>
- Graybiel, A. M., & Grafton, S. T. (2015). The striatum: Where skills and habits meet. *Cold Spring Harbor Perspectives in Biology*, 7(8), a021691.
- Grupe, D. W., & Nitschke, J. B. (2013). Uncertainty and anticipation in anxiety: An integrated neurobiological and psychological perspective. *Nature Reviews Neuroscience*, 14(7), 488.
- Hampton, A. N., Bossaerts, P., & O'doherty, J. P. (2006). The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience*, 26(32), 8360-8367.
- Hashimoto, T., Shimizu, E., Koike, K., Orita, Y., Suzuki, T., Kanahara, N., ... Iyo, M. (2008). Deficits in auditory P50 inhibition in obsessive-compulsive disorder. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, 32(1), 288-296. <https://doi.org/10.1016/j.pnpbp.2007.08.021>
- Hauser, T. U., Fiore, V. G., Moutoussis, M., & Dolan, R. J. (2016). Computational psychiatry of ADHD: Neural gain impairments across Marrian levels of analysis. *Trends in Neurosciences*, 39(2), 63-73.
- Hauser, T. U., Iannaccone, R., Dolan, R. J., Ball, J., Hättenschwiler, J., Drechsler, R., ... Brem, S. (2017). Increased fronto-striatal reward prediction errors moderate decision making in obsessive-compulsive disorder. *Psychological Medicine*, 47(07), 1246-1258. <https://doi.org/10.1017/S0033291716003305>
- Hermans, D., Engelen, U., Grouwels, L., Joos, E., Lemmens, J., & Pieters, G. (2008). Cognitive confidence in obsessive-compulsive disorder: Distrusting perception, attention and memory. *Behaviour Research and Therapy*, 46(1), 98-113. <https://doi.org/10.1016/j.brat.2007.11.001>
- Hermans, D., Martens, K., De Cort, K., Pieters, G., & Eelen, P. (2003). Reality monitoring and metacognitive beliefs related to cognitive confidence in obsessive-compulsive disorder. *Behaviour Research and Therapy*, 41(4), 383-401.
- Hoenig, K., Hochrein, A., Quednow, B. B., Maier, W., & Wagner, M. (2005). Impaired prepulse inhibition of acoustic startle in obsessive-compulsive disorder. *Biological Psychiatry*, 57(10), 1153-1158. <https://doi.org/10.1016/j.biopsych.2005.01.040>
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109(4), 679-709. <https://doi.org/10.1037//0033-295X.109.4.679>
- Huppert, J. D., & Zlotnick, E. (2012). Core fears, values, and obsessive-compulsive disorder: A preliminary clinical-theoretical outlook. *Psicoterapia Cognitiva e Comportamentale*, 18(1), 91-102.
- Jaafari, N., Frasca, M., Rigalleau, F., Rachid, F., Gil, R., Olié, J.-P., ... Rotgé, J.-Y. (2013). Forgetting what you have checked: A link between working memory impairment and checking behaviors in obsessive-compulsive disorder. *European Psychiatry*, 28(2), 87-93.
- Julien, D., O'Connor, K. P., & Aardema, F. (2009). Intrusions related to obsessive-compulsive disorder: A question of content or context? *Journal of Clinical Psychology*, 65(7), 709-722. <https://doi.org/10.1002/jclp.20578>
- Kim, M.-S., Kang, S.-S., Youn, T., Kang, D.-H., Kim, J.-J., & Kwon, J. S. (2003). Neuropsychological correlates of P300 abnormalities in patients with schizophrenia and obsessive-compulsive disorder. *Psychiatry Research: Neuroimaging*, 123(123), 109-123. [https://doi.org/10.1016/S0925-4927\(03\)00045-3](https://doi.org/10.1016/S0925-4927(03)00045-3)
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12), 712-719. <https://doi.org/10.1016/j.tins.2004.10.007>
- Landeros-Weisenberger, A., Bloch, M. H., Kelmendi, B., Wegner, R., Nudel, J., Dombrowski, P., ... Leckman, J. F. (2010). Dimensional predictors of response to SRI pharmacotherapy in obsessive-compulsive disorder. *Journal of Affective Disorders*, 121(1-2), 175-179.
- Lawson, R. P., Mathys, C., & Rees, G. (2017). Adults with autism overestimate the volatility of the sensory environment. *Nature Neuroscience*, 20(9), 1293-1299. <https://doi.org/10.1038/nn.4615>
- Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in Human Neuroscience*, 8(May), 1-10. <https://doi.org/10.3389/fnhum.2014.00302>
- Lazarov, A., Dar, R., Liberman, N., & Oded, Y. (2012a). Obsessive-compulsive tendencies and undermined confidence are related to reliance on proxies for internal states in a false feedback paradigm. *Journal of Behavior Therapy and Experimental Psychiatry*, 43(1), 556-64. <https://doi.org/10.1016/j.jbtep.2011.07.007>
- Lazarov, A., Dar, R., Liberman, N., & Oded, Y. (2012b). Obsessive-compulsive tendencies may be associated with attenuated access to internal states: Evidence from a biofeedback-aided muscle tensing task. *Consciousness and Cognition*, 21(3), 1401-1409. <https://doi.org/10.1016/j.concog.2012.07.002>

- Lazarov, A., Liberman, N., Hermesh, H., & Dar, R. (2014). Seeking proxies for internal states in obsessive-compulsive disorder. *Journal of Abnormal Psychology, 123*(4), 695–704. <https://doi.org/10.1037/abn0000004>
- Lee, H. J., Kwon, S. M., Kwon, J. S., & Telch, M. J. (2005). Testing the autogenous-reactive model of obsessions. *Depression and Anxiety, 21*(3), 118–129. <https://doi.org/10.1002/da.20063>
- Leplow, B., Murphy, R., & Nutzinger, D. O. (2002). *Specificity of Conditional Associative-Learning Deficits in Obsessive-Compulsive Disorder (OCD) and Non-OCD Anxiety Disorders*. 24(6).
- Levy, N. (2018). Obsessive – compulsive disorder as a disorder of attention. *Mind & Language, 33*(1), 3–16. <https://doi.org/10.1111/mila.12172>
- Lewin, A. B., Wu, M. S., Murphy, T. K., & Storch, E. A. (2015). Sensory Over-Responsivity in Pediatric Obsessive Compulsive Disorder. *Journal of Psychopathology and Behavioral Assessment, 37*(1), 134–143. <https://doi.org/10.1007/s10862-014-9442-1>
- Lieder, I., Adam, V., Frenkel, O., Jaffe-Dax, S., Sahani, M., & Ahissar, M. (2019). Perceptual bias reveals slow-updating in autism and fast-forgetting in dyslexia. *Nature Neuroscience, 1*.
- Malloy, P., Rasmussen, S., Braden, W., & Haier, R. J. (1989). Topographic evoked potential mapping in obsessive-compulsive disorder: Evidence of frontal lobe dysfunction. *Psychiatry Research, 28*(1), 63–71. [https://doi.org/10.1016/0165-1781\(89\)90198-4](https://doi.org/10.1016/0165-1781(89)90198-4)
- Martin, J. R., & Pacherie, E. (2013). Out of nowhere: Thought insertion, ownership and context-integration. *Consciousness and Cognition, 22*(1), 111–122. <https://doi.org/10.1016/j.concog.2012.11.012>
- Mataix-Cols, D., Alonso, P., Pifarre, J., Menchon, J. M., & Vallejo, J. (2002). Neuropsychological performance in medicated vs. Unmedicated patients with obsessive-compulsive disorder. *Psychiatry Research, 109*, 255–264.
- Mavrogiorgou, P., Juckel, G., Frodl, T., Hauke, W., Zaudig, M., & Dammann, G. (2002). P300 subcomponents in obsessive-compulsive disorder. *Journal of Psychiatric Research, 36*, 399–406.
- Menzies, L., Chamberlain, S. R., Laird, A. R., Thelen, S. M., Sahakian, B. J., & Bullmore, E. T. (2008). Integrating evidence from neuroimaging and neuropsychological studies of obsessive-compulsive disorder: The orbitofronto-striatal model revisited. *Neuroscience & Biobehavioral Reviews, 32*(3), 525–549.
- Metcalfe, J., & Greene, M. J. (2007). Metacognition of agency. *Journal of Experimental Psychology. General, 136*(2), 184–199. <https://doi.org/10.1037/0096-3445.136.2.184>
- Milad, M. R., & Rauch, S. L. (2012). Obsessive-compulsive disorder: Beyond segregated cortico-striatal pathways. *Trends in Cognitive Sciences, 16*(1), 43–51. <https://doi.org/10.1016/j.tics.2011.11.003>
- Mirza, M. B., Adams, R. A., Parr, T., & Friston, K. (2019). Impulsivity and Active Inference. *Journal of Cognitive Neuroscience, 31*(2), 202–220. https://doi.org/10.1162/jocn_a_01352
- Moore, J., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition, 17*(1), 136–144. <https://doi.org/10.1016/j.concog.2006.12.004>
- Moore, P. J. (2015). A hierarchical narrative framework for OCD. *ArXiv Preprint ArXiv, 1503.00999*, 1–15.
- Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship between feedback-related EEG signals and learning. *BioRxiv*, 581744.
- Nielen, M. M. A., & Den Boer, J. A. (2003). Neuropsychological performance of OCD patients before and after treatment with fluoxetine: Evidence for persistent cognitive deficits. *Psychological Medicine, 33*(5), 917–925.
- Norman, L. J., Carlisi, C. O., Christakou, A., Murphy, C. M., Chantiluke, K., Giampietro, V., ... Rubia, K. (2018). Frontostriatal dysfunction during decision making in attention-deficit/hyperactivity disorder and obsessive-compulsive disorder. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging, 3*(8), 694–703.
- Oades, R. D., Zerbin, D., Dittmann-Balcar, A., & Eggers, C. (1996). Auditory event-related potential (ERP) and difference-wave topography in schizophrenic patients with/without active hallucinations and delusions: A comparison with young obsessive-compulsive disorder (OCD) and healthy subjects. *International Journal of Psychophysiology, 22*(3), 185–214. [https://doi.org/10.1016/0167-8760\(96\)00026-8](https://doi.org/10.1016/0167-8760(96)00026-8)
- O’Connell, K. S., McGregor, N. W., Lochner, C., Emsley, R., & Warnich, L. (2018). The genetic architecture of schizophrenia, bipolar disorder, obsessive-compulsive disorder and autism spectrum disorder. *Molecular and Cellular Neuroscience, 88*(January), 300–307. <https://doi.org/10.1016/j.mcn.2018.02.010>
- Palmer, C. J., Lawson, R. P., & Hohwy, J. (2017). Bayesian Approaches to Autism: Towards Volatility, Action, and Behavior. *Psychological Bulletin, 143*(5), 521–542. <https://doi.org/10.1037/bul0000097>
- Parr, T., & Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface, 14*(136), 20170376. <https://doi.org/10.1098/rsif.2017.0376>
- Parr, T., & Friston, K. J. (2018). The anatomy of inference: Generative models and brain structure. *Frontiers in Computational Neuroscience, 12*, 90.
- Parr, T., Rikhye, R. V., Halassa, M. M., & Friston, K. J. (2019). Prefrontal Computation as Active Inference. *Cerebral Cortex*.
- Pellicano, E., & Burr, D. (2012). When the world becomes “too real”: A Bayesian explanation of autistic perception. *Trends in Cognitive Sciences, 16*(10), 504–510. <https://doi.org/10.1016/j.tics.2012.08.009>
- Peters, A., McEwen, B. S., & Friston, K. (2017). Uncertainty and stress: Why it causes diseases and how it is mastered by the brain. *Progress in Neurobiology, 156*, 164–188. <https://doi.org/10.1016/j.pneurobio.2017.05.004>
- Piantadosi, S. C., Chamberlain, B. L., Glausier, J. R., Lewis, D. A., & Ahmari, S. E. (2019). Lower excitatory synaptic gene expression in orbitofrontal cortex and striatum in an initial study of subjects with obsessive compulsive disorder. *Molecular Psychiatry, 1*–13. <https://doi.org/10.1038/s41380-019-0431-3>
- Pitman, R. K. (1987a). A cybernetic model of obsessive-compulsive psychopathology. *Comprehensive Psychiatry, 28*(4), 334–343. [https://doi.org/10.1016/0010-440X\(87\)90070-8](https://doi.org/10.1016/0010-440X(87)90070-8)
- Pitman, R. K. (1987b). Pierre Janet on obsessive-compulsive disorder (1903): Review and commentary. *Archives of General Psychiatry, 44*(3), 226–232.
- Pizarro, M., Fontenelle, L. F., Paravidino, D. C., Yücel, M., Miguel, E. C., & de Menezes, G. B. (2014). An updated review of antidepressants with marked serotonergic effects in obsessive-compulsive disorder. *Expert Opinion on Pharmacotherapy, 15*(10), 1391–1401.
- Pushkarskaya, H., Tolin, D., Ruderman, L., Henick, D., Kelly, J. M., Pittenger, C., & Levy, I. (2017). Value-based decision making under uncertainty in hoarding and obsessive-compulsive disorders. *Psychiatry Research, 258*, 305–315. <https://doi.org/10.1016/j.psychres.2017.08.058>

- Pushkarskaya, H., Tolin, D., Ruderman, L., Kirshenbaum, A., Kelly, J. M., Pittenger, C., & Levy, I. (2015). Decision-making under uncertainty in obsessive-compulsive disorder. *Journal of Psychiatric Research*. <https://doi.org/10.1016/j.jpsychires.2015.08.011>
- Quattrocki, E., & Friston, K. (2014). Autism, oxytocin and interoception. *Neuroscience and Biobehavioral Reviews*, *47*, 410–430. <https://doi.org/10.1016/j.neubiorev.2014.09.012>
- Rae, C. L., Critchley, H. D., & Seth, A. K. (2019). A Bayesian account of the sensory-motor interactions underlying symptoms of Tourette syndrome. *Frontiers in Psychiatry*, *10*.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. <https://doi.org/10.1038/4580>
- Rauch, S. L., Savage, C. R., Alpert, N. M., Dougherty, D., Kendrick, a, Curran, T., ... Jenike, M. a. (1997). Probing striatal function in obsessive-compulsive disorder: A PET study of implicit sequence learning. *The Journal of Neuropsychiatry and Clinical Neurosciences*, *9*(4), 568–573. <https://doi.org/10.1176/jnp.9.4.568>
- Rauch, Scott L., Wedig, M. M., Wright, C. I., Martis, B., McMullin, K. G., Shin, L. M., ... Wilhelm, S. (2007). Functional Magnetic Resonance Imaging Study of Regional Brain Activation During Implicit Sequence Learning in Obsessive-Compulsive Disorder. *Biological Psychiatry*, *61*(3), 330–336. <https://doi.org/10.1016/j.biopsych.2005.12.012>
- Ray Li, C., Chen, M., Yang, Y., Chang, H., Liu, C., Shen, S., & Chen, C. (2000). Perceptual alternation in obsessive compulsive disorder—Implications for a role of the cortico-striatal circuitry in mediating awareness. *Behavioural Brain Research*, *111*(1–2), 61–69. [http://dx.doi.org/10.1016/S0166-4328\(00\)00140-6](http://dx.doi.org/10.1016/S0166-4328(00)00140-6)
- Reuven-Magril, O., Dar, R., & Liberman, N. (2008). Illusion of control and behavioral control attempts in obsessive-compulsive disorder. *Journal of Abnormal Psychology*, *117*(2), 334–41. <https://doi.org/10.1037/0021-843X.117.2.334>
- Robbins, T. W., Gillan, C. M., Smith, D. G., de Wit, S., & Ersche, K. D. (2012). Neurocognitive endophenotypes of impulsivity and compulsivity: Towards dimensional psychiatry. *Trends in Cognitive Sciences*, *16*(1), 81–91. <https://doi.org/10.1016/j.tics.2011.11.009>
- Roh, K. S., Shin, M. S., KIM, M.-S., HA, T.-H., SHIN, Y.-W., Lee, K. J., & Kwon, J. S. (2005). Persistent cognitive dysfunction in patients with obsessive-compulsive disorder: A naturalistic study. *Psychiatry and Clinical Neurosciences*, *59*(5), 539–545.
- Rossi, S., Bartalini, S., Olivelli, M., Mantovani, A., Di Muro, A., Goracci, A., ... Passero, S. (2005). Hypofunctioning of sensory gating mechanisms in patients with obsessive-compulsive disorder. *Biological Psychiatry*, *57*(1), 16–20. <https://doi.org/10.1016/j.biopsych.2004.09.023>
- Rotge, J. Y., Clair, A. H., Jaafari, N., Hantouche, E. G., Pelissolo, A., Goillandeau, M., ... Burbaud, P. (2008). A challenging task for assessment of checking behaviors in obsessive-compulsive disorder. *Acta Psychiatrica Scandinavica*, *117*(6), 465–473.
- Salkovskis, P. M. (1999). Understanding and treating obsessive—Compulsive disorder. *Behaviour Research and Therapy*, *37*, S29–S52.
- Sambrook, T. D., & Goslin, J. (2015). A neural reward prediction error revealed by a meta-analysis of ERPs using great grand averages. *Psychological Bulletin*, *141*(1), 213–235. <https://doi.org/10.1037/bul0000006>
- Sanborn, A. N. (2017). Types of approximation for probabilistic cognition: Sampling and variational. *Brain and Cognition*, *112*, 98–101. <https://doi.org/10.1016/j.bandc.2015.06.008>
- Sanz, M., Molina, V., Martin-Loeches, M., Calcedo, a, & Rubia, F. J. (2001). Auditory P300 event related potential and serotonin reuptake inhibitor treatment in obsessive-compulsive disorder patients. *Psychiatry Research*, *101*(1), 75–81. [https://doi.org/10.1016/s0165-1781\(00\)00250-x](https://doi.org/10.1016/s0165-1781(00)00250-x)
- Sarig, S., Dar, R., & Liberman, N. (2012). Obsessive-compulsive tendencies are related to indecisiveness and reliance on feedback in a neutral color judgment task. *Journal of Behavior Therapy and Experimental Psychiatry*, *43*(1), 692–697. <https://doi.org/10.1016/j.jbtep.2011.09.012>
- Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human orbitofrontal cortex represents a cognitive map of state space. *Neuron*, *91*(6), 1402–1412.
- Seli, P., Risko, E. F., Purdon, C., & Smilek, D. (2016). Intrusive thoughts: Linking spontaneous mind wandering and OCD symptomatology. *Psychological Research*. <https://doi.org/10.1007/s00426-016-0756-3>
- Sheppard, D. M., Bradshaw, J. L., Purcell, R., & Pantelis, C. (1999). Tourette's and comorbid syndromes: Obsessive compulsive and attention deficit hyperactivity disorder. a common etiology? *Clinical Psychology Review*, *19*(5), 531–552. [https://doi.org/10.1016/S0272-7358\(98\)00059-2](https://doi.org/10.1016/S0272-7358(98)00059-2)
- Shipp, S., Adams, R. A., & Friston, K. J. (2013). Reflections on agranular architecture: Predictive coding in the motor cortex. *Trends in Neurosciences*, *36*(12), 706–716. <https://doi.org/10.1016/j.tins.2013.09.004>
- Speekenbrink, M. (2016). A tutorial on particle filters. *Journal of Mathematical Psychology*, *73*, 140–152. <https://doi.org/10.1016/j.jmp.2016.05.006>
- Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., ... Corlett, P. R. (2018). The Predictive Coding Account of Psychosis. *Biological Psychiatry*, 1–10. <https://doi.org/10.1016/j.biopsych.2018.05.015>
- Sterzer, P., Mishara, A. L., Voss, M., & Heinz, A. (2016). Thought Insertion as a Self-Disturbance: An Integration of Predictive Coding and Phenomenological Approaches. *Frontiers in Human Neuroscience*, *10*(October), 1–12. <https://doi.org/10.3389/fnhum.2016.00502>
- Summerfeldt, L. J. (2004). Understanding and treating incompleteness in obsessive-compulsive disorder. *Journal of Clinical Psychology*, *60*(11), 1155–68. <https://doi.org/10.1002/jclp.20080>
- Swerdlow, N. R., Benbow, C. H., Zisook, S., Geyer, M. A., & Braff, D. L. (1993). A preliminary assessment of sensorimotor gating in patients with obsessive compulsive disorder. *Biological Psychiatry*, *33*(4), 298–301. [https://doi.org/10.1016/0006-3223\(93\)90300-3](https://doi.org/10.1016/0006-3223(93)90300-3)
- Szechtman, H., & Woody, E. (2004). Obsessive-compulsive disorder as a disturbance of security motivation. *Psychological Review*, *111*(1), 111–27. <https://doi.org/10.1037/0033-295X.111.1.111>
- Tekcan, A. I., Topçuoğlu, V., & Kaya, B. (2007). Memory and metamemory for semantic information in obsessive-compulsive disorder. *Behaviour Research and Therapy*, *45*(9), 2164–2172. <https://doi.org/10.1016/j.brat.2006.10.002>
- Toffolo, M. B. J., van den Hout, M. A., Engelhard, I. M., Hooge, I. T. C., & Cath, D. C. (2016). Patients With Obsessive-Compulsive Disorder Check Excessively in Response to Mild Uncertainty. *Behavior Therapy*, *47*(4), 550–559. <https://doi.org/10.1016/j.beth.2016.04.002>
- Toffolo, M. B. J., van den Hout, M. A., Hooge, I. T. C., Engelhard, I. M., & Cath, D. C. (2013). Mild Uncertainty Promotes Checking Behavior in Subclinical Obsessive-Compulsive Disorder. *Clinical Psychological Science*, *1*(2), 103–109. <https://doi.org/10.1177/2167702612472487>

- Towey, J. P., Tenke, C. E., Bruder, G. E., Leite, P., Friedman, D., Liebowitz, M., & Hollander, E. (1994). Brain event-related potential correlates of overfocused attention in obsessive-compulsive disorder. *Psychophysiology*, *31*(6), 535–543. <https://doi.org/10.1111/j.1469-8986.1994.tb02346.x>
- Tuna, Ş., Tekcan, A. İ., & Topçuoğlu, V. (2005). Memory and metamemory in obsessive-compulsive disorder. *Behaviour Research and Therapy*, *43*(1), 15–27. <https://doi.org/10.1016/j.brat.2003.11.001>
- Vaghi, M. M., Luyckx, F., Sule, A., Fineberg, N. A., Robbins, T. W., & De Martino, B. (2017). Compulsivity Reveals a Novel Dissociation between Action and Confidence. *Neuron*, *96*(2), 348–354.e4. <https://doi.org/10.1016/j.neuron.2017.09.006>
- van den Hout, M. A., van Dis, E. A. M., van Woudenberg, C., & van de Groep, I. H. (2018). OCD-like checking in the lab: A meta-analysis and improvement of an experimental paradigm. *Journal of Obsessive-Compulsive and Related Disorders*, (November), 1–11. <https://doi.org/10.1016/j.jocrd.2017.11.006>
- Veale, D., Miles, S., Smallcombe, N., Ghezai, H., Goldacre, B., & Hodsoll, J. (2014). Atypical antipsychotic augmentation in SSRI treatment refractory obsessive-compulsive disorder: A systematic review and meta-analysis. *BMC Psychiatry*, *14*(1), 317.
- Voon, V., Derbyshire, K., Rück, C., Irvine, M. A., Worbe, Y., Enander, J., ... Bullmore, E. T. (2015). Disorders of compulsivity: A common bias towards learning habits. *Molecular Psychiatry*, *20*(3), 345–352. <https://doi.org/10.1038/mp.2014.44>
- Weilhammer, V., Stuke, H., Hesselmann, G., Sterzer, P., & Schmack, K. (2017). A predictive coding account of bistable perception—A model-based fMRI study. *PLoS Computational Biology*, *13*(5), 1–21. <https://doi.org/10.1371/journal.pcbi.1005536>
- Wikenheiser, A. M., & Schoenbaum, G. (2016). Over the river, through the woods: Cognitive maps in the hippocampus and orbitofrontal cortex. *Nature Reviews. Neuroscience*, *17*(8), 513–523. <https://doi.org/10.1038/nrn.2016.56>
- Wu, M. S., Lewin, A. B., Murphy, T. K., & Storch, E. A. (2014). Misophonia: Incidence, Phenomenology, and Clinical Correlates in an Undergraduate Student Sample. *Journal of Clinical Psychology*, *70*(10), 994–1007. <https://doi.org/10.1002/jclp.22098>
- Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The Neural Basis of Error Detection: Conflict Monitoring and the Error-Related Negativity. *Psychological Review*, *111*(4), 931–959. <https://doi.org/10.1037/0033-295X.111.4.931>
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*(4), 681–692. <https://doi.org/10.1016/j.neuron.2005.04.026>
- Zhang, Z., Wang, M., Miao, X., Li, Y., Hitchman, G., & Yuan, Z. (2017). Individuals with high obsessive-compulsive tendencies or undermined confidence rely more on external proxies to access their internal states. *Journal of Behavior Therapy and Experimental Psychiatry*, *54*, 263–269. <https://doi.org/10.1016/j.jbtep.2016.09.003>
- Zorowitz, S., Momennejad, I., & Daw, N. D. (2019). Anxiety, avoidance, and sequential evaluation. *BioRxiv*, 724492.

Supplemental material

Hand-washing, particle filter simulations

Particle filters are a set of sampling algorithm used for Bayesian inference for state-space models, which include both observation uncertainty and transition uncertainty (Speekenbrink, 2016). In contrast to the Kalman filter, which provides a analytic Bayesian solution to state space models, particle filters are not limited to Gaussian variables. The reasons for which we chose to use particle filter, rather than a Kalman filter in these simulations are reviewed below.

State-space models describe the evolution of latent states $x_{1:T}$ as a first-order Markovian process, where each state depends probabilistically only on the previous state:

$$p(x_t) = p(x_t | x_{t-1}, \hat{\sigma}_p^2) \quad (S1)$$

, where $\hat{\sigma}_p^2$ represents a transition uncertainty parameter, which can be variance, or other uncertainty parameters in the case of other distributions. The observation at time t (o_t) depends probabilistically only on the latent state at time t (x_t):

$$p(o_t) = p(o_t | x_t, \hat{\sigma}_o^2) \quad (S2)$$

, where $\hat{\sigma}_o^2$ represents an observation uncertainty parameters, which again can be the variance or other uncertainty parameters.

A generic particle filter represents estimated latent states with a finite set of particles generating a Monte-Carlo distribution. At each time step, particles are weighted in accordance with Equation S2. This procedure, known as importance sampling (Speekenbrink, 2016) generates an approximate posterior distribution, integrating the prior estimate of the latent state (particle distribution before weighting), and the likelihood distribution (Equation S2). Most particle filters then use a resampling procedure, where particles are resampled in proportion to their weights, resulting in a particle distribution where the relative number (rather than the weights) of the particles represent the relative posterior probability of each particle. This resampling step was shown to be critical to prevent the 'particle degeneration problem'. As a final step, the particle set at time t is propagated in correspondence with Equation S1, to account for transition uncertainty (i.e. process noise).

We used a particle filter (rather than a Kalman filter) for two reasons. First, although we assumed that the evolution of latent states is Gaussian, we also assumed that the observations are generated from a Gaussian

distribution truncated at zero, to represent the idea that any evidence for hands' dirtiness cannot reach negative values (i.e. hands can be perfectly clean at most). Second, to model an affective bias which can be described as the belief that the latent state generating observable information is usually more negative than the observation, we used a skewed normal distribution, with a negative skew. In our case, this corresponds with the belief that one's hands are usually dirtier than they look/feel.

The actual sensory variance, denoted by σ_o^2 was set to 50, whereas actual process noise (i.e. transition uncertainty) was set to 0.1. To focus on action-dependent transitions, process noise was set to 0 when no action (i.e. hand-washing) was applied. The values of observation noise and process noise are important only in relation to one each other. Sensory noise was set to a considerably higher value than process noise to represent the assumption that sensory feedback regarding hand dirtiness is in most cases highly non-informative (except for relatively infrequent cases where one can actually see dirt on one's hands). Whereas hand washing changes hand dirtiness, we assumed that it does so in a relatively predictable fashion; that is: we assumed that hand washing is a more informative cue than sensory information regarding hand dirtiness.

Finally, as a simple response model, the probability with which agents were assumed to wash their hands was determined in accordance with the proportion of particles surpassing a specific criterion, which represents a subjective criterion for "how clean does one expect/need one's hands to be". Hand washing was assumed to decrease the actual hand dirtiness by a (arbitrary) factor of 0.8. A multiplicative rather than additive factor was used, to account for the idea that hand washing can never yield perfectly clean hands, and that the effects of hand washing decrease for cleaner hands. The only modeling assumption critical for the results of the simulation is that of representing hand washing as a sequential Bayesian inference, whereas some of the more peripheral assumptions described above could be relaxed with no marked effect on the results of the simulation.

We now describe the particle filter algorithm used for these simulations. In this algorithm, x_t represents the level of hand dirtiness at time t , estimated by a set of P particles: $\hat{x}_t^{(p)}$. At each time t , the weight of each particle is denoted by $w_t^{(p)}$. p_{wash_t} corresponds with the probability of washing one's hands at time t . The model included 3 free parameters. First, b corresponded with the ratio of the estimated state uncertainty ($\hat{\sigma}_p^2$) to the actual state uncertainty (σ_p^2), and it was set to 1 for

simulated control subjects, and to 100 to simulated OCD patients (see Figure 3A vs 3B, respectively, for results). c corresponded with the cleanliness criterion, and it was set to 5 in all but the 'perfectionism' simulation, where it was set to 1 (see Figure 3A for results).

Finally, α corresponds with the skew parameter of the skewed normal distribution, and it was set to 0 (corresponding to a Gaussian distribution) for all but the 'affective bias' simulation (see Figure 3B for results). 1000 particles were used for the simulations. In the algorithm below, asterisks correspond with changes in the real world (rather than in the Bayesian estimate thereof):

Particle filter algorithm

1. *(Actual state) $x_0 = 100$ (initial level of dirtiness)
2. (Initialize) sample P particles from $\hat{x}_0^{(p)} \sim N(x_0, 100)$
3. For $t=1 \dots T$
 - a. *(Action) $wash_t \sim \text{Bernoulli}[P(\hat{x}_{t-1}^{(p)} > c)]$
 - b. *(Actual state) $\begin{cases} \text{if } wash_t = 1 & \text{sample } x_t \sim N(x_{t-1} * 0.8, \sigma_p^2) \\ \text{if } wash_t = 0 & \text{set } x_t = x_{t-1} \end{cases}$
 - c. *(Observation) sample a random observation from $o_t \sim N(x_t, \sigma_o^2)$
 - d. (Propagate) $\begin{cases} \text{if } wash_t = 1 & \text{sample } \hat{x}_t^{(p)} \sim N(\hat{x}_{t-1}^{(p)} * 0.8, \hat{\sigma}_p^2) \\ \text{if } wash_t = 0 & \text{set } \hat{x}_t^{(p)} = \hat{x}_{t-1}^{(p)} \end{cases}$
with $\hat{\sigma}_p^2 = \sigma_p^2 * b$
 - e. (Weight) Weight particles in accordance with $w_t^{(p)} \sim SN(o_t, \hat{\sigma}_o^2, \alpha)$
with $\sum w = 1$
 - f. (Resample) Resample particles $\hat{x}_t^{(p)}$ in accordance with $w_t^{(p)}$

The free energy framework

The free energy framework has been suggested as a unified theory of cognition and action, and it can be viewed as proposing a specific mathematical framework for conceptualizing and quantifying predictive processing in the brain. Interested readers can find detailed theoretical reviews (Friston, 2010), and mathematical tutorials (Bogacz, 2017; Buckley, Kim, McGregor, & Seth, 2017) elsewhere. Here we briefly describe the very basic mathematical principles relating the free energy framework to the Bayesian brain, and predictive coding.

In most realistic cases Bayesian inference requires non-analytic approaches. The particle filter presented above is an example of one type of solutions which use sampling techniques to approximate Bayesian inference. Whereas sampling provides asymptotically correct solutions with infinite sample size, finite samples provide an approximation of the exact solution. This weak point of sampling techniques has been suggested to make it appropriate to explain the often non-optimal cognition (Sanborn, 2017).

Another set of solutions provide tools for Bayesian inference in complex scenarios by approximating the true posterior distribution by a simpler distribution. This class of solutions, called variational Bayes, converts the Bayesian inference problem to an optimization problem, focusing on finding the distribution that best approximates the true posterior. Specifically, assuming the exact Bayesian posterior is:

$$p(x|o) = \frac{p(x, o)}{p(o)} = \frac{p(o|x)p(x)}{\int p(o|x)p(x)dx}. \quad (S3)$$

In variational inference, $p(x|o)$ is approximated by another, simpler distribution $q(x)$, with the goal of minimizing the dissimilarity between these two distributions. The most common measure of dissimilarity is the Kullback-Leibler divergence, defined as:

$$KL(q(x), p(x|o)) = \int q(x) \ln \frac{q(x)}{p(x|o)} dx \quad (S4)$$

Note that when $q(x)$ equals $p(x|o)$, their ratio is equal to 1, and its logarithm, as well as the entire expression is equal to zero.

Critically, solving the KL divergence is still impossible, because it requires one to know the posterior distribution, $p(x|o)$. However, one can substitute this term by its definition above (Equation S3), which gives:

$$\begin{aligned} KL(q(x), p(x|o)) &= \int q(x) \ln \frac{q(x)}{\frac{p(o, x)}{p(o)}} dx \\ &= \int q(x) \ln \frac{q(x)p(o)}{p(o, x)} dx \\ &= \int q(x) \left[\ln \frac{q(x)}{p(o, x)} + \ln p(o) \right] dx \\ &= \int q(x) \ln \frac{q(x)}{p(o, x)} dx + \int q(x) \ln p(o) dx \\ &= \int q(x) \ln \frac{q(x)}{p(o, x)} dx + \ln p(o) \end{aligned} \quad (S5)$$

, where the transitions between the second and third lines were based on logarithm product rule; the transition from the third to the fourth line was based on integral rules; and the transition from the fourth to the last line was based on the fact that $\ln p(o)$ is a constant and thus can be pulled out the integral, whereas the integral of $q(x)$ is 1, because it is a probability distribution.

The integral in the last line of Equation S5 is called free energy (F), such that Equation S5 can be expressed as:

$$KL(q(x), p(x|o)) = F + \ln p(o). \quad (S6)$$

Now, recall that our objective here is to minimize the KL divergence. Because $\ln p(o)$ is a constant with respect to $q(x)$, instead of minimizing the KL we can simply minimize the free energy. Thus, in variational Bayes, approximate Bayesian inference is done by minimizing free energy. This explains why according to the free energy framework, the (Bayesian) brain acts to minimize free energy.

Note further that we can rearrange the terms in Equation S6, to:

$$KL(q(x), p(x|o)) - F = \ln p(o). \quad (S7)$$

Because the KL divergence is, by definition, always positive, the negative free energy is also the lower bound for $\ln p(o)$. That means, that:

$$F \geq -\ln p(o). \quad (S8)$$

Equation S8 also serves a critical role in the free energy framework, because the term on its right side is called surprisal in information theory, and it quantifies the extent to which observations are novel to the agent (i.e. observations with low probability are most surprising). This means, that by minimizing free energy one

affectively minimizes surprisal. Indeed the key theoretical assumption of the free energy framework is that all living organisms strive to minimize surprisal, which ultimately stands for staying in the homeostatic boundaries allowing for their existence (e.g., for a fish, being out of the water is a highly surprising observation).

In terms of cognition, $q(x)$ can be thought of as beliefs regarding the true state of the world x . Minimizing variational free energy thus minimizes the difference between one's beliefs and observed outcomes interpreted in light of a probabilistic generative model with assumptions regarding how observations and latent states covary:

$$\begin{aligned} F &= \int q(x) [\ln q(x) - \ln p(x, o)] dx & (S9) \\ &= E_q [\ln q(x) \\ &\quad - \ln p(x, o)] \end{aligned}$$

In active inference, the selection of policies (or actions) is also assumed to be guided by free energy minimization. However, here agents are assumed to strive to minimize their expected free energy (rather than the current free energy), because this process is assumed to occur during planning. Mirroring Equation S9, expected free energy, given a specific policy π can be defined as:

$$\begin{aligned} G &= E_q [\ln q(x|\pi) - \ln p(x, o|\pi)] \\ &= E_q [\ln q(x|\pi) - \ln p(x|o, \pi) - \ln p(o|\pi)] \\ &= E_q [\ln q(x|\pi) - \ln p(x|o, \pi)] - E_q [\ln p(o|\pi)] \end{aligned} \quad (S10)$$

Note that the expectation here is with respect to a 'posterior predictive' distribution $q(o, x|\pi) = p(o|x)q(x|\pi)$, so that we take expected future observations into account. Assuming that the agent's preferences (prior expectation) regarding outcomes do not depend on the selected policy, the expected free energy becomes:

$$\begin{aligned} G(\pi) &= \underbrace{E_q [\ln q(x|\pi) - \ln p(x|o, \pi)]}_{\text{epistemic value}} - \underbrace{E_q [\ln p(o)]}_{\text{extrinsic value}} \\ &\approx -E_q [KL(q(x|o, \pi), q(x|\pi))] - E_q [\ln p(o)] \quad (S11) \\ &= E_q [H[p(o|x)]] - H[q(o|\pi)] - E_q [\ln p(o)] \end{aligned}$$

This formula weights the epistemic value, and the extrinsic value of a policy. The second line shows that, if we approximate $p(x|o, \pi)q(o|\pi) \approx q(x, o|\pi)$, we can rewrite the epistemic value as an expected KL-Divergence (or information gain). The third line rewrites

epistemic value involving two terms. First, the expected free energy of a policy decreases as the state transitions become more uncertain, as they lead to a greater uncertainty about future outcomes, or a greater entropy of the predictive distribution $q(o|\pi)$. Thus, for example, when a certain policy is believed with high certainty to bring in a certain state, $H[q(o|\pi)]$ is low, signaling that following that policy provides less knowledge, because there is limited uncertainty to resolve. Second, the expected free energy decreases further, as the observations are more informative, that is – as they disambiguate one's prediction regarding the consequences of a policy. The ambiguity associated with a given policy may be expressed as the entropy of the likelihood (first term on line 3 of Equation S11), as this expresses the fidelity with which states map to outcomes. In sum, policies related with states characterized by relatively high state uncertainty and relatively low observational noise have a higher epistemic value. Finally, the extrinsic value of policies is formalized in active inference in accordance with the prior distribution over observations. That is, desired outcomes are defined as observations with high prior probability, such that policies expected to result in desired outcomes have a lower expected free energy.

Generative model for the active inference simulations

The MDP scheme used in our simulations closely resembled the one specified in (Parr & Friston, 2017), with the addition of habit learning, represented as the prior probability over policies ($E(\pi)$), which is learned over time. Thus, policy selection is a softmax function (σ) of these prior beliefs, and the expected free energy of each policy (see above), with the inverse temperature parameter γ :

$$p(\pi) = \sigma[\ln E(\pi) - \gamma \cdot G(\pi)]. \quad (S12)$$

As described above, the expected free energy of a policy is a function of beliefs regarding state transitions for that policy, beliefs regarding the state-outcome mapping, prior preferences regarding outcomes, and prior belief regarding the initial state, defined respectively as:

$$\begin{aligned} p(x_{t+1}|x_t, \pi) &= \text{Cat}(B(u)), & u &= \pi(t) \\ p(o_t|x_t) &= \text{Cat}(A) \\ p(o_t) &= \text{Cat}(C) \\ p(x_1) &= \text{Cat}(D) \end{aligned} \quad (S13)$$

In the current scheme, the outcome and state space are each defined in terms of two factors. That is, the

state space is factorized into hidden states defining position (e.g., where did the agent 'click'), and hidden states defining the context (i.e. the correct cue). Whereas the transition matrix (B) for position states depends on the action u , the transition matrix for context states is defined irrespective of action. Specifically, the probability of each context to persist on the next time step is parameterized by the parameter b , which is at the focus of the current simulations. It is important to note that b here refers to the estimated transition uncertainty (i.e. from the agent's perspective). The actual probability for state transitions was 0 (with the exception of one, predefined transition at the reversal trial).

The outcome space is factorized into two factors: exteroceptive outcomes, that are informative about the agent's position, and affective (which could be defined as interoceptive) outcomes which could be a reward, a punishment or a neutral affective outcome. Each hidden state (in terms of both position and context) is mapped to an outcome. The exteroceptive outcome (position) is a deterministic (identity) function of the position state, whereas the affective (interoceptive) outcome is a probabilistic function of the joint position-context state. Here we assume that the probability for a reward, given that the agent has chosen the correct cue (e.g., choosing the left cue in the context that the left cue is correct), is defined by a , whereas the probability of reward for each of incorrect cues is $1-a$.

The D vectors (i.e. initial states) were parameterized with Dirichlet priors. The agent beliefs each trial starts from the initial/baseline position (position states factor). In the first trial, the agent has a relatively non-informative (symmetric Dirichlet, with the concentration parameter $\alpha = 4$) prior regarding context (context states factor), and this belief is updated on each trial, based on the generative model and observed outcomes. Policies were defined only in terms of position states, because the agent had no direct influence on context states. Correspondingly, habits were learned only for position related policies. Habits were also parameterized in terms of a symmetric Dirichlet prior, where the concentration parameter defines the amount of prior certainty in a given policy (see Figure 4). Finally, the C vector (i.e. preferred outcomes) was set to include preferences for affective outcomes, and no preferences for position outcomes (i.e. the agent prefers to get rewards, but has no preference regarding its position per se).

Specific generative model for task 1.

In task 1 there were 4 hidden position states, and 3 hidden context states. The transition matrix for position, for each policy, was defined such that each action

(in this task policies were of length-1) leads directly to the respective position. For example, for policy 'L':

$$B_{position}('L') = \begin{matrix} & N & L & R & M \\ \begin{matrix} L \\ R \\ M \end{matrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix}$$

Here, the transitions occur from state represented by each column to those represented by rows, with states being: neutral/initial position (N), left position (L), right position (R), and middle position (M). The transition matrix for context, independent of action, was defined as (Here L, R and M refer to context and not to the agent's position) :

$$B_{context} = \begin{matrix} & L & R & M \\ \begin{matrix} L \\ R \\ M \end{matrix} & \begin{bmatrix} b & 0.5 * (1 - b) & 0.5 * (1 - b) \\ 0.5 * (1 - b) & b & 0.5 * (1 - b) \\ 0.5 * (1 - b) & 0.5 * (1 - b) & b \end{bmatrix} \end{matrix}$$

The likelihood matrix for position ($A_{exteroceptive}$) for each context was a simple identity matrix (implying a deterministic relationship between position states and observed positions). The likelihood matrices for affective outcomes prescribe the probability for a reward given the position and context states. So for example, when the left cue is correct, the likelihood matrix mapping position states (columns) to affective outcomes (rows) is (note that the 'L', in the left term, here refers to the context in which the left cue is correct – not to the policy 'choose left cue', as in the transition matrix above):

$$A_{affective}('L') = \begin{matrix} & Neutral & Reward & Punishment & N \\ \begin{matrix} L \\ R \\ M \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & a & 1 - a & 1 - a \\ 0 & 1 - a & a & a \end{bmatrix} \end{matrix}$$

Specific generative model for task 2.

In task 2 there were 4 hidden position states, and 2 hidden context states. The transition matrix for position, for each policy, was defined such that each action (in this task policies were of length-2) leads directly to the respective position, with the exception of the utilitarian cue positions (left/right) being absorbing states, such that once entered (in a given trial) the agent stays in this position until the end of the policy (see Friston et al., 2016). Thus, for example, for policy 'L':

$$B_{position}('L') = \begin{matrix} & N & L & R & C \\ \begin{matrix} L \\ R \\ C \end{matrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \end{matrix}$$

The transition matrix for context, independent of action, was defined as:

$$B_{context} = \begin{matrix} L \\ R \end{matrix} \begin{bmatrix} b & 1-b \\ 1-b & b \end{bmatrix} \begin{matrix} L \\ R \end{matrix}$$

In this task, the exteroceptive likelihood matrix included information regarding position, as well as information signaling context when the agent was at the 'checking' (C) position state. So, when the left cue is correct, the exteroceptive likelihood matrix is:

$$A_{Exteroceptive}(L') = \begin{matrix} N \\ L \\ R \\ context = L \\ context = R \end{matrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{matrix} N \\ L \\ R \\ C \end{matrix}$$

, and when the right cue is correct it becomes:

$$A_{Exteroceptive}(R') = \begin{matrix} N \\ L \\ R \\ context = L \\ context = R \end{matrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{matrix} N \\ L \\ R \\ C \end{matrix}$$

The likelihood array for affective outcomes describes the probability for a reward given the position

and context states. Here the affective outcome of both the initial and the checking positions is neutral (defined as 0 in the C matrix). So for example, when the left cue is correct, the likelihood matrix mapping position states (columns) to affective outcomes (rows) is:

$$A_{affective}(L') = \begin{matrix} Neutral \\ Reward \\ Punishment \end{matrix} \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & a & 1-a & 0 \\ 0 & 1-a & a & 0 \end{bmatrix} \begin{matrix} N \\ L \\ R \\ C \end{matrix}$$

Note also, that in the second task a was set to 1 (only deterministic state-outcome contingencies), because otherwise the outcome of checking behavior should also have been defined as probabilistic, adding a currently unnecessary level of complexity.

References

- Bogacz, R. (2017). A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*, *76*, 198–211. <https://doi.org/10.1016/j.jmp.2015.11.003>
- Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, *81*, 55–79. <https://doi.org/10.1016/j.jmp.2017.09.004>
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews. Neuroscience*, *11*(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O’Doherty, J., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience and Biobehavioral Reviews*, *68*, 862–879. <https://doi.org/10.1016/j.neubiorev.2016.06.022>
- Parr, T., & Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface*, *14*(136), 20170376. <https://doi.org/10.1098/rsif.2017.0376>
- Sanborn, A. N. (2017). Types of approximation for probabilistic cognition: Sampling and variational. *Brain and Cognition*, *112*, 98–101. <https://doi.org/10.1016/j.bandc.2015.06.008>
- Speekenbrink, M. (2016). A tutorial on particle filters. *Journal of Mathematical Psychology*, *73*, 140–152. <https://doi.org/10.1016/j.jmp.2016.05>

