Perception and misperception of bodily symptoms from an Active Inference perspective: Modelling the case of panic disorder.

Domenico Maisto
Institute for High Performance Computing and Networking, National Research Council, Naples, Italy

Laura Barca
Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy

Omer Van den Bergh
Health Psychology, University of Leuven, Belgium

Giovanni Pezzulo
Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy

Corresponding author
Giovanni Pezzulo
ISTC-CNR
Via San Martino della Battaglia 44
00185 Rome, Italy
Email: giovanni.pezzulo@istc.cnr.it
Phone: +39 6 44595206

Abstract

We advance a novel computational model that characterizes formally the ways we perceive or misperceive bodily symptoms, in the context of panic attacks. The computational model is grounded within the formal framework of Active Inference, which considers top-down prediction and attention dynamics as key to perceptual inference and action selection. In a series of simulations, we use the computational model to reproduce key facets of adaptive and maladaptive symptom perception: the ways we infer our bodily state by integrating prior information and somatic afferents; the ways we decide whether or not to attend to somatic channels; the ways we use the symptom inference to make decisions about taking or not taking a medicine; and the ways all the above processes can go awry, determining symptoms misperception and ensuing maladaptive behaviors, such as hypervigilance or excessive medicine use. While recent existing theoretical treatments of psychopathological conditions focus on prediction-based perception (predictive coding), our computational model goes beyond them, in at least two ways. First, it includes action and attention selection dynamics that are disregarded in previous conceptualizations but are crucial to fully understand the phenomenology of bodily symptoms perception and misperception. Second, it is a fully implemented model that generates specific (and personalized) quantitative predictions, thus going beyond previous qualitative frameworks.

Keywords: panic disorder, active inference, predictive coding, maladaptive inference, computational psychiatry

Introduction

There is increasing consensus in neuroscience that the perception (or misperception) of both external events, such as visual scenes, and internal events, such as bodily symptoms, is not a purely bottom-up process, but an inferential process guided by top-down dynamics, which include predictions derived from prior knowledge and top-down attention dynamics (Clark, 2016; Engel et al., 2001; K. J. Friston, 2005; Hohwy, 2013; Rao & Ballard, 1999). In computational psychiatry, this implies that understanding phenomena like symptom perception (or misperception) requires going beyond bottom-up and afferent processes, to also consider top-down, predictive and attention processes. In other words, what is needed is a novel conceptual framework that puts predictive dynamics at its core and affords an "inferential leap" (Van den Bergh et al., 2017) in our understanding of how we perceive and feel our bodily symptoms.

One prediction-based framework that is rapidly gaining prominence in neuroscience is Active Inference. This framework describes the brain as a “predictive machine”, which forms an internal (generative) model of the statistical regularities of its environment, and uses the model for perception and adaptive control (K. J. Friston, 2010; K. J. Friston et al., 2017; Pezzulo et al., 2015; Pezzulo, Rigoli, et al., 2018). An Active Inference agent continuously strives to minimize the discrepancy between its prior expectations (including homeostatic priors, such as “expecting not to be hungry”) and its sensory evidence (including interoceptive sensations, such as “feeling hungry”). To minimize this discrepancy (or more formally, variational free energy), an Active Inference agent uses synergistically perceptual processes (inferring the latent causes of the sensorium, e.g., infer whether one is hungry), action processes (selecting action plans that bring about predicted / desired sensory outcomes, which confirm to expectation, e.g., elicit sensations of satiation by eating) and learning processes (finessing internal models, e.g., learn that some physical activities make one hungry).

Active Inference is increasingly applied to elucidate a wide variety of phenomena such as action selection, learning, motor control, homeostatic regulation, interoceptive processing, and uncertainty reduction or epistemic foraging (Donnarumma et al., 2017; K. Friston et al., 2012, 2015; K. J. Friston, FitzGerald, Rigoli, Schwartenbeck, O’Doherty, et al., 2016; K. J. Friston et al., 2017; Pezzulo, Rigoli, et al., 2018). Furthermore, Active Inference and the closely related theory of *predictive coding* have been recently adopted to explain *failures* of the above (and other) processes and the resulting psychopathologies, including somatic symptom disorder, schizophrenia, psychosis, hysteria, and beyond (Adams et al., 2013; Barrett et al., 2016; Corlett et al., 2019; Edwards et al., 2012; K. J. Friston, Stephan, et al., 2014; Paulus et al., 2019; Paulus, 2019; Pezzulo et al., 2019; A. R. Powers et al., 2017; Sterzer et al., 2018; Van den Bergh et al., 2017). However, it is important to clarify that predictive coding models only consider the perceptual domain (K. J. Friston, 2005; Rao & Ballard, 1999), whereas Active Inference extends it to consider both perception and action, under a unique imperative (variational free energy minimization). The scope of Active Inference theories is broader and includes not just misperceptions or false beliefs but also maladaptive behavior stemming from psychopathological conditions. However, most existing explanations of psychopathology use the constructs of Active Inference at a rather abstract level, without also advancing specific formal (computational) models.

In this article, we go one step beyond these previous attempts and use Active Inference to develop a formal model of symptom perception and misperception in the context of panic disorder. In line with recent pleas for the development of computational

psychiatry and psychosomatics (K. J. Friston, Stephan, et al., 2014; Petzschner, 2017; Petzschner et al., 2017), we present a computational example illustrating how an Active Inference framework can model processes of symptom perception and behavioral action in panic disorder. Our proposal builds on previous conceptualizations of interoceptive processing within the Active Inference framework, which we introduce below.

Interoceptive processing and symptom perception from an active inference perspective

Interoception refers to the ways by which afferent information from visceral and bodily input is processed in the brain (Craig, 2015). Interoception is important for homeostatic regulation and adaptive behavior (Barrett & Simmons, 2015; Pezzulo et al., 2015) but it is also thought to be critically involved in feelings and emotions (Barrett, 2017; Seth, 2013), in the experience of agency and selfhood (Seth & Tsakiris, 2018) and in related psychopathological disorders (Quadt et al., 2018; Smith et al., 2020). Most of these interactions between interoception and psychological function occur outside awareness, but some information may be consciously perceived. The extent to which people are able to accurately perceive their internal state remains a matter of debate. On the one hand, individual differences seem to exist in this ability which may correlate with psychological function such as variations in emotion processing, in the experience of selfhood and in sensitivity for body illusions (Berntson et al., 2018). On the other hand, the most widely used test to assess interoceptive accuracy in such studies, the heartbeat counting task, appears methodologically severely flawed and, if anything, it shows that the large majority of people are poor heartbeat perceivers (Zamariola et al., 2018).

It is important to note that the debate on interoceptive accuracy is largely predicated on the assumption that afferent sensory information, originating in the body, fully determines interoception (Van den Bergh et al., 2018). However, according to the recent theories of *interoceptive inference* (Seth, 2013) and *Embodied Predictive Interoception Coding model* (EPIC) (Barrett & Simmons, 2015), interoception emerges from brain predictions about the state of the body that are interacting with actual sensory input. In this view, interoception is a matter of Active Inference (Seth & Friston, 2016).

This Active Inference account of interoception can be extended to symptom perception. Obviously, the perception of bodily symptoms as indicators of physiological dysfunction relies on interoception, but it is more than that. It can be considered a form of *categorical perception* during which patterns of interoceptive stimulation are grouped into meaningful categories (e.g., cardiac, gastrointestinal) with behavioural relevance. In an Active Inference framework, categorical perception of interoceptive variables is a Bayesian inferential process, akin to categorical perception of exteroceptive variables (e.g., object categories). This process is based on the two typical ingredients of Bayesian inference: prior beliefs (predictions) about categories, given an individual's learning history and current context (Lynn & Barrett, 2014; Rigoli et al., 2017); and sensory evidence (here, interoceptive or somatic input). Crucially, both prior belief and sensory evidence contribute to Bayesian inference, but their relative "weights" depend on their reliability or precision (formally, the inverse of the variance of a distribution, e.g., Gaussian distribution). This implies that, depending on the relative precision of (the neural representations of) both the categorical predictions and the actual interoceptive input, the eventual symptom percept will be closer to the predictions or to the somatic input (Van den Bergh et al., 2017). In conditions of highly precise (but incorrect) categorical priors and low precise interoceptive input (e.g., noisy somatic channels), perceived symptoms may reflect prior beliefs rather than actual

input and hence provide a biased perception of bodily state (Henningsen et al., 2018; Van den Bergh et al., 2018).

An additional important aspect of Active Inference is that it describes perception as an active process – because agents are free to decide what information and what sensory channels to attend or not to attend (this includes both *overt* and *covert* attention allocation, but in this context, we only focus on *covert* attention). The decision of what information should be attended to considers both the agent's prior beliefs and (beliefs about) the *precision* of information sources (e.g., somatic channels). Intuitively, an Active Inference agent only needs to attend precise information that has some *information value*, i.e., has the potential to change his belief state. Information that can resolve the agent's uncertainty about motivationally salient events (e.g., about whether or not one has a dysfunction) and comes from reliable sensory channels has high information gain and *should* be (covertly) attended to; whereas information that is already fully predicted (e.g., is associated with a very strong prior) or comes from noisy sensory channels has little or no information value and hence *should not* be attended to. These attention dynamics are important for symptom perception, too. This is because the same conditions that may promote symptom misperception – highly precise (but incorrect) categorical priors and imprecise interoceptive input – may also determine a progressive *loss of somatic attention*, which further impairs the correct symptom perception. On the contrary, an imprecise internal model that leaves a person uncertain about (present or future) motivationally salient events may promote an excessive level of somatic attention, associated with hypervigilance.

In sum, Active Inference suggests that the imbalance of (the precision of) prior and sensory evidence can cause symptoms misperception, both as a direct effect of (precision-weighted) Bayesian inference, and as an indirect effect of (uncertainty-resolving) attention dynamics. Of note, Active Inference suggests that in most cases, what appears to be *maladaptive inference* is not the result of wrong inferential processes, but of correct inferential processes given odd premises (e.g., excessively precise priors). In this article, we apply this Active Inference perspective on symptoms misperception in the context of a specific psychopathological condition - panic disorder - and later we embody our proposals in a fully implemented computational model.

An Active Inference perspective on panic disorder

Panic disorder is characterized by sudden bouts of massive autonomic activity dominated by cardiorespiratory sensations and accompanied by severe anxiety and a feeling of impending doom (e.g. fear of dying). Although occasional panic attacks are rather prevalent in the general population, panic disorder develops when individuals start to persistently worry about having another attack and to substantially change behavior in order to predict, control or avoid potential new attacks. Fear conditioning through which low-level cardiorespiratory sensations become interoceptive conditioned cues is considered an important process in the maintenance and exacerbation of panic disorder (Bouton et al., 2001), resulting in patients becoming hypervigilant for interoceptive cues that may predict a full-blown panic attack. Safety learning easily develops taking the form of excessive dependence on the presence of others (e.g. a spouse), on being in specific situations (e.g. at home) and/or on taking medication (e.g. alprazolam).

The above brief description of panic disorder is a severe simplification, but it may be sufficient to lay the ground for a computational model explaining the perception and misperception of bodily symptoms in these patients from an Active Inference perspective.

For illustrative purposes, we focus on only a few specific characteristics: a person with panic disorder who wants to decide whether or not to take a medicine (e.g. alprazolam) to stop a possible panic attack. Clearly, the optimal plan is to take the medicine only if he is about to have a possible panic attack and not continuously as dependence might ensue. In the simulations below, we define the patient's goal as a generic "feeling normal" sensation, which can be achieved in two ways: if he is about to have a panic attack and takes the medicine, or if he is not about to have a panic attack and does not take the medicine.

To make an optimal decision about taking or not taking the medicine, the patient has to first resolve his uncertainty about whether he is about to have a panic attack (context 1: "panic attack") or not (context 2: "no panic attack"). He makes this inference by jointly considering two factors. The first factor is his prior belief about experiencing panic attacks more frequently in some situations (e.g., in public places) but not others (e.g., at home). The second factor is his current bodily sensations of two kinds, one of which (heart pounding) is more diagnostic then the other (breathlessness) about upcoming panic attacks. Importantly, in the simulations that we describe below, the person can decide whether or not to attend to these slightly aversive bodily sensations before making the decision, in order to reduce his uncertainty. Our simulations will illustrate how Active Inference agents balance *epistemic* (uncertainty-minimizing) aspects of the decision, such as whether or not to attend somatic stimuli, with *economic* (reward-maximizing) aspects of the decision, such as whether or not to take the medicine. To illustrate the peculiarities of the Active Inference model, we will compare it with a standard decision-making model that only considers reward maximization; and show that the uncertainty-minimization imperative implicit in Active Inference is important to ensure adaptive action in ambiguous contexts. Finally, our simulations will show the ways Active Inference can go awry and cause symptoms misperception, hypervigilance and excessive safety behaviour – therefore explaining the clinical phenomenology or panic disorders.

Specification of the agent's generative model

In the Active Inference framework, an agent makes decisions (e.g., whether or not to take a medicine) based on its *generative model*, which generally includes 4 components, called **A**, **B**, **C**, **D** matrices (see the Appendix for a formal description). The first component (likelihood function, **A**) is a probabilistic description of what *observations* (e.g., somatic sensations, such as heart pounding, breathlessness and "feeling strange") should be expected under different states of affairs, which cannot be directly observed and are hence called *hidden states* (e.g., having or not having a panic attack). The second component (transition function, **B**) is a probabilistic description of how *hidden states* change over time, as an effect of the agent's actions (e.g., what happens if one takes a medicine when he has, or not has, a panic attack). The third component is a (prior, **C**) probability distribution over observations, which in Active Inference encodes the agent's prior preferences or goals (i.e., that he likes "feeling normal" but dislikes "feeling strange", heart pounding and breathlessness). The fourth condition is the agent's prior belief (**D**) about its current hidden state, before receiving any sensory observation.

Once an Active Inference agent is endowed with a specific generative model (fully specified by the **A**, **B**, **C**, **D** matrices) and is exposed to a certain set of environmental stimuli (e.g., it feels pounding), its decision dynamics follow automatically from a standard (variational) inferential scheme that approximates full Bayesian inference: free energy minimization. A full description of the free energy minimization procedure is provided in the

Appendix. Here, it suffices to say that an Active Inference agent uses the generative model for both *perception* and *action selection*. The former corresponds to inferring its current hidden state (e.g., whether it is having a pain attack) based on its observations (e.g., given that it feels breathlessness) and prior beliefs. The latter corresponds to selecting an action or a sequence of actions (called policy) that simultaneously achieves the agent's goals ("feeling normal") and lowers his uncertainty about the hidden state.

In the context of the current paper, perception corresponds to the inference of panic attack symptoms, whereas action selection corresponds to two kinds of decisions: *taking* or *not taking* a medicine, and *attending* or *not attending* to somatic channels (the latter can be considered a sort of attention regulation). We assume that both symptom perception (or misperception) and adaptive (or maladaptive) action selection emerge from Active Inference, under a particular generative model (and **A**, **B**, **C**, **D** matrices), which we introduce below, see Figure 1.

The first component of the generative model is a likelihood (**A**) matrix that maps probabilistically ten *hidden states* to forty *observations* (and is stochastic). The ten *hidden states* are obtained as the tensorial product between two hidden *contexts*, which correspond to the two hypotheses of being about to have (context 1: "panic attack"), or not to have (context 2: "no panic attack"), a panic attack; by five *task situations*: an "initial state", two states that correspond to "having attended to heart pounding" and "having attended to breathlessness" and two states that correspond to "having taken the medicine" or "not having taken the medicine", respectively. The forty observations are obtained as the tensorial products between five *exteroceptive sensations* and eight *interoceptive sensations*. The five exteroceptive sensations correspond one-to-one to the five task situations explained above (not shown in Figure 1 for the sake of simplicity). As the agent receives a different exteroceptive sensation in each task situation, it always knows unambiguously in which task situation it is. However, exteroceptive sensations do not disambiguate between contexts: an agent in the "initial state" receives the same exteroceptive sensation, independent of the context (panic or not panic) it is in. To infer in which context it is, the agent has to consider its eight interoceptive sensations. These include two noninformative initial sensations (*a* and *b*), two bodily sensations that correspond to the presence or absence of heart pounding (*c* and *d*), two bodily sensations that correspond to the presence or absence of breathlessness (*e* and *f*) and two final sensations: a more positive sensation called "feeling normal" (*g*) and a more negative sensation called "feeling strange", e.g., lightheadedness (*h*). As mentioned above, the probabilistic mapping between the ten hidden states and the eight observations corresponds to the likelihood (**A**) matrix and it is set as stochastic in our simulations. Note that the mapping from hypotheses to the heart pounding sensations (*c* and *d*) is "sharper" (i.e., has lower entropy) than the mapping from hypotheses to breathlessness sensations (*e* and *f*), reflecting the fact that the presence of the former is more diagnostic for panic attack than the latter. In other words, if I am having a panic attack I'm very likely to sense heart pounding when attending to heart pounding but less likely to sense breathlessness, when attending to breathlessness.

Note also that the agent "feels normal" (i.e., observes *g*) in two conditions: when he is about to have a panic attack and takes the medicine; and when he is not about to have a panic attack and does not take the medicine. Conversely, the agent "feels strange" (i.e., observes *h*) when he is about to have a panic attack and does not take the medicine; and when he is not about to have a panic attack and takes the medicine (e.g., feels a bit dry in the mouth).

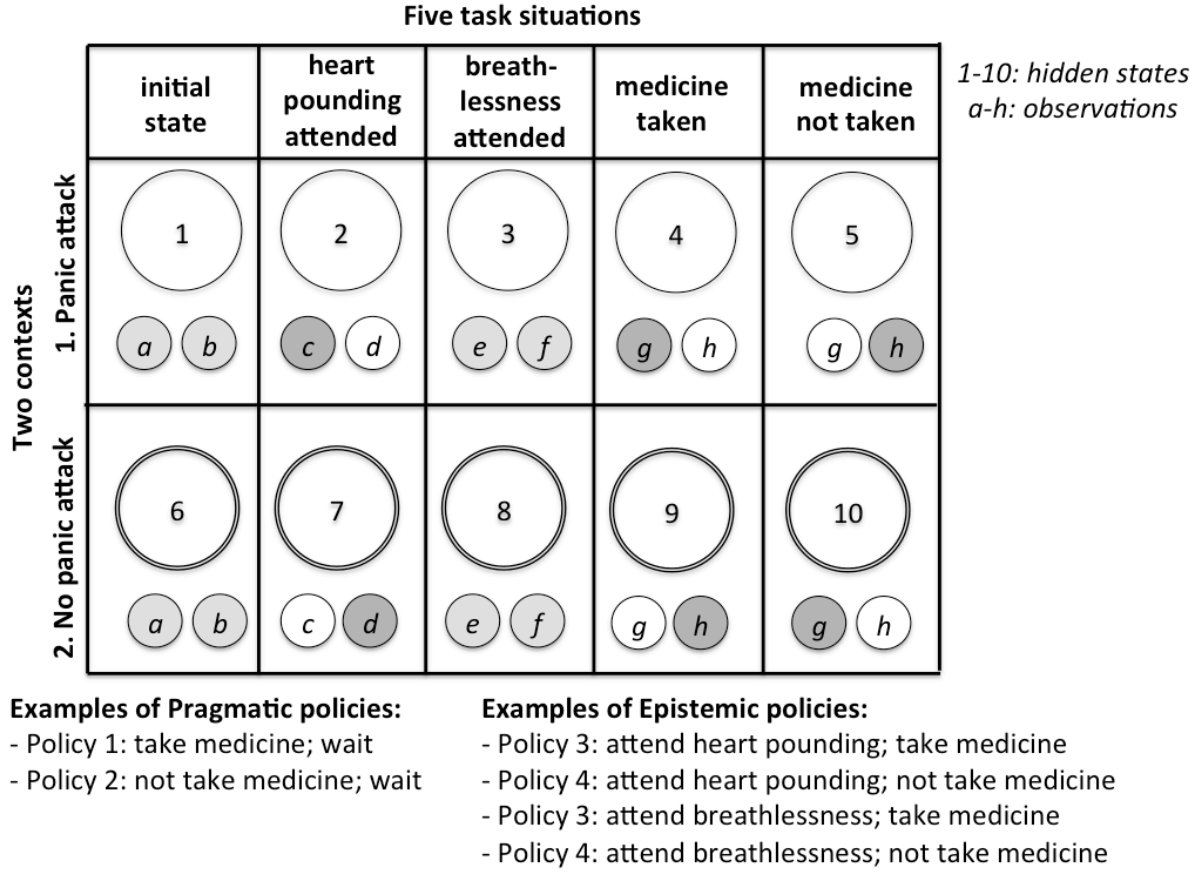

Figure 1. The agent's generative model. The agent's generative model includes ten hidden states (1-10): five task situations (initial state, having attended to heart pounding or breathlessness, having or not having taken the medicine) by two contexts (panic and not panic). Furthermore, the generative model includes five exteroceptive sensations (not shown in the figure for the sake of simplicity) that map one-to-one to the five task situations; and eight interoceptive sensations (a-h): two noninformative initial sensations, the presence or absence of heart pounding, the presence or absence of breathlessness, a positive sensation ("feeling normal") and a negative sensation (feeling "strange"). See the main text for explanation.

Furthermore, the generative model includes five actions, which correspond to moving to any of the five task situations (i.e., going to start, attend heart pounding, attend breathlessness, take medicine, not take medicine). Note that in active inference, actions (or more formally, *control states*) are hidden variables that need to be inferred. The probabilistic mapping between actions and subsequent hidden states is encoded in the transition (**B**) matrix. In our simulations, this mapping depends on the agent's hidden context. Importantly, transitions between task situations are deterministic, which implies that when an agent has selected a desired task situation (e.g., "heart pounding attended"), it can confidently reach it, by selecting the appropriate action (e.g., "attend heart pounding"), whichever its initial state. However, there are exceptions: we set the terminal states 4, 5, 9 and 10 to be "absorbing states", i.e., once reached, they cannot be left. This means that executing any action (e.g., go to start or attend to heart pounding) in states 4, 5, 9, or 10 will imply a self-transition to the same states 4, 5, 9, or 10.

Transitions between contexts are stochastic; but the probability to change context is small (0.1) and constant for all actions. For example, an agent that starts from (any state of) the "panic attack" context and selects the action to "attend heart pounding" will make a high-probability (0.9) transition to hidden state 2 (context: "panic attack", task situation: "heart pounding attended") and a low-probability (0.1) transition to hidden state 7 (context: "no panic attack", task situation: "heart pounding attended"). Conversely, an agent that starts from (any state of) the "no panic attack" context and selects the action to "attend heart pounding" will make a high-probability (0.9) transition to hidden state 7 (context: "no panic attack", task situation: "heart pounding attended") and a low-probability (0.1) transition to hidden state 2 (context: " panic attack", task situation: "heart pounding attended"). This implies that the agent has no specific actions to change its hidden context (e.g., from "no panic" to "panic"), but it can rarely experience contextual changes, by chance.

By starting from this action set, the agent can consider, and select amongst, several 2-step *policies* (i.e., sequences of 2 actions). In our simulation, we will consider the choice between two kinds of policies: *pragmatic policies* to go directly to one of the four terminal states 4, 5, 9, or 10 and remain there (e.g., Policy 1: take medicine and then wait); and *epistemic policies* to attend to bodily sensations (by going to states 2, 3, 7 or 8) before reaching a terminal state (e.g., Policy 3: attend to heart pounding and then take medicine).

The latter, epistemic policies are important as they visit hidden states pertaining to heart pounding and breathlessness, which are the only states that provide reliable information about the agent's context (panic or no panic). These can be regarded as "attentional states" in the following sense. We assume that at the beginning of the trial, the agent does not attend to its interoceptive channels - and hence it receives uninformative sensations (a,b) from which it cannot infer whether or not it is having a panic attack. To receive informative sensations, the agent has to attend to its interoceptive channels (heart pounding or breathlessness). When attending to heart pounding, there is a precise likelihood mapping to the presence or absence of corresponding sensations, depending upon whether one is about to have a panic attack or not. In other words, the sensations of heart pounding rest upon an interaction between two hidden states - the context (impending panic attack or not) and attentional set (attending to heart pounding or attending to breathlessness). Note that the subject has to be in one of the five task situations and in one of the two contexts. Technically, these constitute hidden factors with five and two levels, respectively. Importantly, a subject can move between the five task situations, choosing whether to attend to heart pounding, breathlessness or forego any particular attentional set and take medicine (or not). However, the subject cannot intentionally move between the two contexts (i.e., chose whether or not to have a panic attack). Contextual changes can occur stochastically in our simulations, but their (small) probability does not depend on the selected action - hence the agent has no control over its context (panic or not panic).

Finally, the generative model includes two kinds of priors. The former is a prior over observations, which in active inference plays the role of the person's prior preferences or goals. This is encoded in the model as the (**C**) vector, which assigns positive valence to feeling normal, negative valence to feeling strange, and slight negative valence to heart pounding and breathlessness (which are felt as aversive by patients). The second is a prior belief about the person’s current state, encoded in the model as the (**D**) matrix; and in our simulations, it reflects the knowledge that it starts from the initial state but does not know

its context (i.e., the probabilities of being in "panic attack" and "no panic attack" contexts are the same).

Note that this latter (prior) belief is subjective – it is part of the person's *generative model* – and distinct from the true state the person is in, encoded in (**S**), which is instead part of environmental dynamics (also called the *generative process* in Active Inference). This separation between the generative model (encoding what the persons knows and prefers) and the generative process (encoding the "true" environmental state, including the hidden context the agent is in) is crucial in our simulations. This is because, initially, the agent is not sure about his context (having or not having a panic attack) and has to resolve this uncertainty to decide what to do.

The perception and misperception of bodily symptoms: five simulations

After specifying the agent's generative model, we illustrate the agent's behavior (prescribed by its generative model) in a series of five simulations. These simulations illustrate the functioning of Active Inference and how maladaptive parametrizations of the generative model give rise to various manifestations of panic disorders.

First simulation: correct inference of (not having) a panic attack

The first simulation illustrates a case of correct inference that one is not actually experiencing a panic attack. The simulation uses the generative model illustrated in Figure 1. The agent always begins from the start state in the "no panic attack" context (i.e., the initial hidden state encoded in **S** is state 6 in Figure 1) and samples a sensory observation. Importantly, while the agent knows it is in the start state, it does not know in which context it is (i.e., having or not having a panic attack). This corresponds to the fact that the **D** values for the two contexts are equal (0.5) and the observation it samples in the start state is uninformative. This simulation illustrates that before being able to select the optimal action (which in this case is not taking the medication), it has to firstly infer in which context it is ("no panic attack") – and this epistemic behavior is automatically derived in Active Inference. The parameters used in this and the following simulations are shown in Table A1.

The agent can select amongst multiple 2-step policies (i.e., sequences of 2 actions), such as those illustrated in Figure 1, e.g., attending to heart pounding and then taking medicine; attending to heart pounding and then not taking medicine; taking medicine directly; not taking medicine directly. This choice guides transitions between hidden states (as described in the matrix **B)** and the sampling of observations (as described in the matrix **A**), resulting in a "feeling normal" sensation if the agent has not taken medicine and a "feeling strange" sensation if the agent has taken medicine. Note that while entire 2-step policies are evaluated at the first time step, only their first action is selected and executed. After selecting an action, the agent does a (deterministic) transition to a new hidden state, collects an observation, re-does the policy selection, and then selects the second action.

The simulation results of the Active Inference agent are illustrated in Figure 2A. In this and the subsequent figures, the results are the average of 128 simulations, with the same initial conditions (which is useful to explore the diversity of the possible solutions, as some aspects of the simulation are stochastic). The top panel shows the proportion of times the agent is in the five task situations (ordinate) during the tree time steps of the simulation (abscissa). At the first time step, the agent always starts from the task situation 1 (initial situation). It can be appreciated that, in most cases, the agent selects the epistemic policy to first attend to the most diagnostic source of information (heart pounding) to resolve his

uncertainty; and then not to take the medicine, hence often achieving the desired "feeling normal" sensation. This epistemic behavior is necessary because initially, the agent is only certain about its task state (see first column of top panel), but not its context (see first column of center panel). In most cases, the agent selects an epistemic action to go to the "heart pounding attended" state (second column of top panel), where it receives an informative bodily sensation (i.e., not having heart pounding), with high probability. In turn, this bodily sensation increases the agent's belief that he is in the "no panic attack" context (second column of center panel). At this point, it can confidently select a pragmatic action, to go to the "medicine not taken" state (third column of top panel) and hence "feeling normal" most of the time (not shown).

The center panel shows which context (panic or not panic) the agent infers being in during the simulation. The agent is initially fully uncertain about context but its uncertainty drastically reduces after the first choice (i.e., at the second time step), because it frequently selects the epistemic action to attend heart pounding - and hence gathers informative interoceptive sensations. Furthermore, the agent's contextual uncertainty reduces again at the third time step, as it usually gathers a "feeling normal" sensation at the terminal state that is informative about context.

The bottom panel shows the average pragmatic (utility maximization) and epistemic (uncertainty reducing) values of the states actually visited by the agent during the three time steps. It can be appreciated that the agent's first choice (after the first time step) brings greater epistemic value and resolves most of the agent's uncertainty; whereas the agent's second choice (after the second time step) has greater utility.

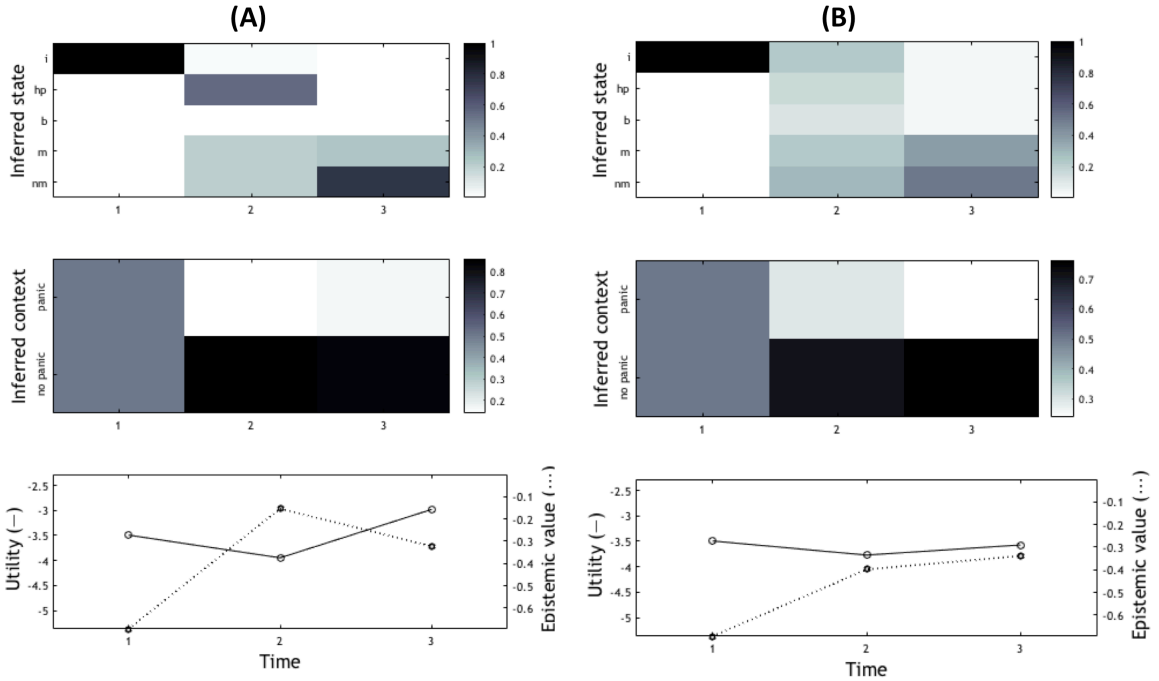

Figure 2. Results of the first simulation. (A) Results of the Active Inference model. The top panel shows in ordinate in which of the five hidden task situations the agent is in, from top to bottom (i: initial state; hp: heart pounding attended; b: breathlessness attended; m: medicine taken; nm: medicine not taken) and in abscissa the three time steps of the simulation (left: first; center: second; right: third time step). Colors denote probabilities, with darker colors corresponding to higher probabilities. The center panel shows the hidden context (having or not having a panic attack) inferred by the agent. Note that the agent is

uncertain about its context during the first time step; it becomes more certain at the second time step and it fully resolves its uncertainty at the third time step. The bottom panel shows the average pragmatic (utility maximization) and epistemic (uncertainty reducing) values of the states actually visited by the agent during the three time steps. These are the two terms of Expected Free Energy illustrated in the Appendix. (B) Results of the utility maximization model that lacks the epistemic component. See the main text for explanation.

This first simulation illustrates two main things. First, while deciding under uncertainty, it is optimal to firstly execute an *epistemic action* (to attend to a bodily sensation) to disambiguate the context, and then make the second, *pragmatic action* (whether or not to take medicine) with high confidence. Second, the selection of a specific epistemic action depends on the informativeness of the information sources; because in our simulations heart pounding is more diagnostic than breathlessness (i.e., its column in the **A** matrix has lower entropy), it is attended to more often.

An important peculiarity of Active Inference compared to traditional models of decision-making as utility maximization in (neuro)economics (Glimcher & Rustichini, 2004; Loewenstein et al., 2007) is that it automatically balances pragmatic (utility maximization) and epistemic (uncertainty minimization) components of choice - which are two equally important parts of the *expected free energy* equation used for planning (see the Appendix). The importance of jointly considering pragmatic and epistemic components of choice can be appreciated by comparing the performance of the Active Inference agent with a classical utility maximization agent. The latter can derived from the Active Inference agent by simply removing the "epistemic value" component of the free energy equation, see (K. Friston et al., 2015). The simulation results of the utility maximization agent are shown in Figure 2B. While the Active Inference agent tends to select the state having highest epistemic value (e.g., attend to heart pounding) at the first time step, the utility maximization agent does not consider the epistemic value of attending to heart pounding, but only its slightly negative value - and hence rarely selects it. As a consequence, the utility maximization agent remains more uncertain about its context (panic or not panic; see the second panel of Figure 2B) and makes less informed choices - selecting the wrong action (take medicine) almost half of the times. This comparison illustrates that making decisions in an ambiguous context (e.g., without knowing in advance whether or not it is having a panic attack) requires going beyond classical utility maximization schemes, to also include uncertainty minimization. Indeed, epistemic actions that resolve uncertainty are important prerequisites to increase utility afterwards. While it is certainly possible to add external incentives (e.g., novelty or curiosity bonuses) to classical utility maximizing agents, Active Inference realizes the integration of pragmatic and epistemic imperatives in a principled way - as they both stem from the same free energy minimization imperative.

Second simulation: the role of prior information

In the first simulation, we considered a situation in which the agent was completely uncertain (or had flat prior belief) about the possibility of having a panic attack. However, an agent can also start with a more informative prior about it, either correct or incorrect. For example, one can consider that it is a priori more (or less) likely to have a panic attack because one is for example in a public place (or at home). In this second simulation, we vary such prior information (encoded in the agent's **D** vector), to see how its initial belief influences his plans. Figure 3 shows the results of the simulations, with five different levels

of initial belief about being in a "panic attack" context; one column for each simulation. Note that like in the first simulation the agent's "true" but unknown context (encoded in the **S** vector) is always "no panic attack".

The first column shows the results of a simulation with an initially very high belief (0.9) about having a panic attack. This high prior may occur, for example, when a person is in a place where he had a panic attack in the past (and shows some form of conditioning). In the case of a high prior about having a panic attack, different from our first simulation, it becomes optimal for the agent to select a pragmatic policy, to take the medicine directly. Note that in this case, the prior belief was completely misleading (as the agent's true context is "no panic attack"); and hence the performance of the agent is very poor (bottom panel). The fifth column shows a similar situation, with an initially very high belief (0.9) about not having a panic attack. Even in this case, the agent selects directly a pragmatic policy – not to take the medicine – that yields the desired results, as the prior belief matches the "true" context. These two simulations show that high levels of prior belief favor the selection of pragmatic policies, which can be more or less effective depending on how appropriate the prior belief is. Because most contextual uncertainty has been already minimized, epistemic actions (to attend to heart pounding or breathlessness) cease to be informative and are not selected. This contrasts with our first simulation (whose results are reproduced in the third column of Figure 3, for the sake of the reader), where epistemic actions are required to reduce uncertainty before the agent can confidently select whether or not to take the medicine.

The second and fourth columns of Figure 4 show intermediate situations, in which the agent is endowed with a moderate (0.6) prior that he is about to have, or not to have, a panic attack. These situations are akin to the first simulation and illustrate the importance of resolving uncertainty (by selecting an epistemic action) before a pragmatic choice can be made confidently.

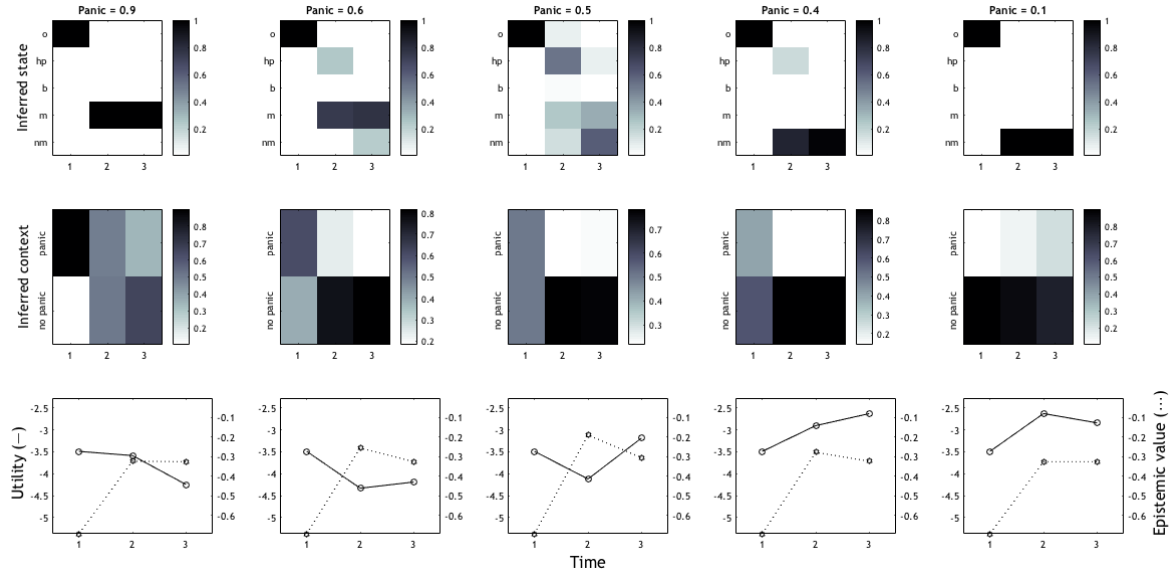

*Figure 3. Results of the second simulation. The five columns show five parameterizations of the agent's prior belief about its hidden context, from 0.9 (first column) to 0.1 (last column) probability of being in "no panic attack" context, as encoded in the agent's **D** vector. The meaning of symbols is the same as the first simulation.*

In sum, this second simulation illustrates that an agent's prior belief influences both hidden state estimation and policy selection. The former occurs because high (or very precise) priors may easily dominate over evidence during hidden state estimation. In keeping, several active inference treatments of psychopathology highlight the potential maladaptive roles of excessively high (or excessively precise) priors in perceptual inference and hidden state estimation. A standard assumption of Bayesian (and predictive coding) models of perception is that what an agent actually perceives corresponds to the hypothesis having the highest probability. In our example, an agent that assigns a high probability to the "panic attack" context would actually perceive that he is in a panic-related somatic state, which (if one assumes a hierarchical agent model) in a cascade would entail the perception of heart pounding while they are actually not there, and possibly other panic symptoms. If an excessively high (or precise) prior dominates the inference, it can automatically determine the (mis)perception of bodily symptoms.

However, our simulation goes beyond predictive coding and mere perceptual processing. In Active Inference (as opposed to predictive coding) the inference of the most likely context does not only determine what is perceived but also action and policy selection, in two ways: by determining the final choice (e.g., about the medicine) and the selection of epistemic versus pragmatic policies. Hence, not only an agent endowed with a strong prior about having a panic attack would misperceive its bodily symptoms, but it would also favor policies that take the medicine and disfavor epistemic policies. The reason for the latter is that strong and/or excessively precise priors would be hardly changed by novel information - and hence by definition, this novel information acquires low information gain, rendering epistemic policies unnecessary. But in turn, disfavoring epistemic actions may render priors largely impermeable to novel evidence – as observed in psychopathological conditions (Van den Bergh et al., 2017).

Third simulation: uninformative bodily states (and the suppression of epistemic actions)

In this third simulation, we consider the case of an agent endowed with non-informative bodily sensations. While in the previous simulations we considered heart pounding sensation to be highly informative and diagnostic about having or not having a panic attack, here we render these streams less informative, by increasing the entropy (i.e., decreasing the precision) of their associated columns in the **A** matrix. Specifically, the only parts of the **A** matrix that we changed (become more entropic) compared to the previous simulation is the mapping between hidden contexts (having or not having a panic attack) to bodily sensations (heart pounding and breathlessness). This implies that whatever the bodily sensation (e.g., perceiving or not perceiving heart pounding), the hidden context would remain uncertain. We consider the case of an agent endowed with five levels of prior belief (see the five columns of Figure 4), from 0.9 to 0.1, in favor of the hypothesis of having a panic attack.

The simulation results show that with non-informative bodily sensations, epistemic actions are largely ignored (they are only selected infrequently in the case of uncertain prior beliefs, third column of Figure 4). This situation can be compared with the results of the second simulation, when epistemic actions were much more frequent, especially with uncertain prior beliefs. Of course, the absence of informative bodily sensations also prevents the agent to make appropriate decisions, thus lowering its performance.

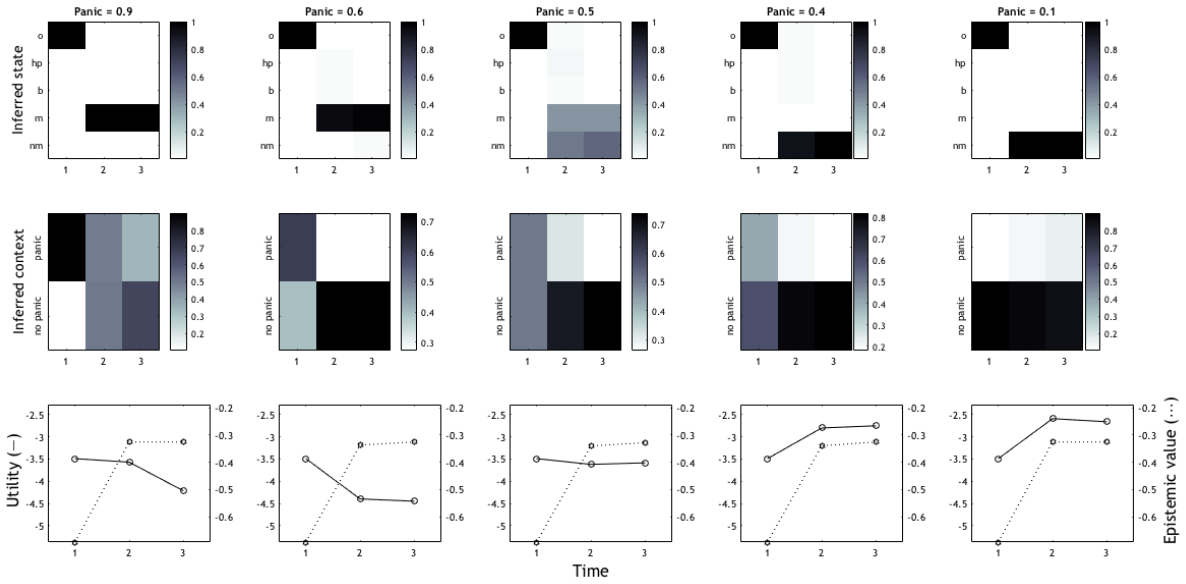

*Figure 4. Results of the third simulation. An agent is endowed with a matrix **A** in which the mapping between the two contexts (having or not having a panic attack) and the two bodily sensations (heart pounding and breathlessness) has high entropy. The mapping between the two contexts and the two final feelings (feeling normal or strange) has instead low entropy, as in the previous simulations. The five columns show simulations with different prior beliefs about having a panic attack, from 0.9 (first column) to 0.1 (last column).*

The simulation illustrated here shows that precise sensory (here, bodily or somatic) sensations are key to accurate inference. In Bayesian treatments of perception (and predictive coding), prior information and evidence are integrated in an optimal fashion and weighted according to their relative precision; hence imprecise sensory information only influences perceptual inference to a minor extent.

Active Inference adds an active (attention modulation) component to this idea. Since an Active Inference agent is free to select whether or not to attend to specific sources of information, it will tend to ignore imprecise sources of information (as indexed by a high entropy of the **A** matrix), as they have little information gain. In other words, patients who perceive their interoceptive channels to be imprecise (either because they really are, due to some physiological dysfunction; or because the patient's model includes incorrect precision parameters) will systematically fail to attend to informative signals from their bodies. In psychopathological conditions, this (rational) inattention prevents a patient's internal models to be correctly updated in the light of novel evidence and hence can contribute to maintain incorrect priors, and lower the awareness of one's bodily condition (Barrett et al., 2016; Seth & Tsakiris, 2018; Smith et al., 2020).

Fourth simulation: imprecise transition model and hypervigilance

The previous simulations illustrated the importance of having a flexible prior (**D**) and an appropriate likelihood model (**A** matrix), with a precise mapping between hidden states and observations. This is because imprecise mappings prevent a person from correctly updating their beliefs and selecting informative (epistemic) actions. Here, we illustrate the importance of having a sufficiently precise transition model (**B** matrix), to maintain knowledge about the correct context over time. In all the above simulations, we assumed that the transitions between hidden states were deterministic (given the agent's choice of

action), and the agent maintained a perfect knowledge of his previous context (or did not consider that the context has changed). This implies that the agent automatically carries over his previous knowledge about his current context (e.g., having a panic attack), across consecutive time steps, and once he becomes sufficiently confident about it, he does not need to check it anymore. However, if the mapping between consecutive hidden states was imprecise – meaning that the agent has no perfect (working) memory of his previous estimates or, equivalently, considers that the context has changed in the meantime – he would experience a loss of information over consecutive trials (Parr & Friston, 2017a). Therefore, even if the agent acquired sufficient confidence that he was not having a panic attack, he would need to continuously confirm this belief, by continuing executing epistemic actions that perceive bodily sensations – hence indulging in a form of hypervigilance.

This situation is illustrated in the simulations shown in Figure 5, which compares the agent's behavior endowed with a **B** matrix that has low entropy, analogous to the first simulation (Figure 5A), and with a **B** matrix that has high entropy (Figure 5B). Note that to better appreciate the retention or loss of information over time as a function of the entropy of the **B** matrix, the simulations shown in Figure 5 have five steps, not three as in the previous simulations. When an agent is endowed with a **B** matrix that has low entropy, he resolves his uncertainty (and reaches one of the two terminal states, taking or not taking the medicine) after one epistemic action, in the majority of cases. Rather, when the agent is endowed with a **B** matrix that has high entropy, his uncertainty decreases less steeply (as he considers the possibility that the context will change after each action) and hence he continues executing epistemic actions also during successive time steps.

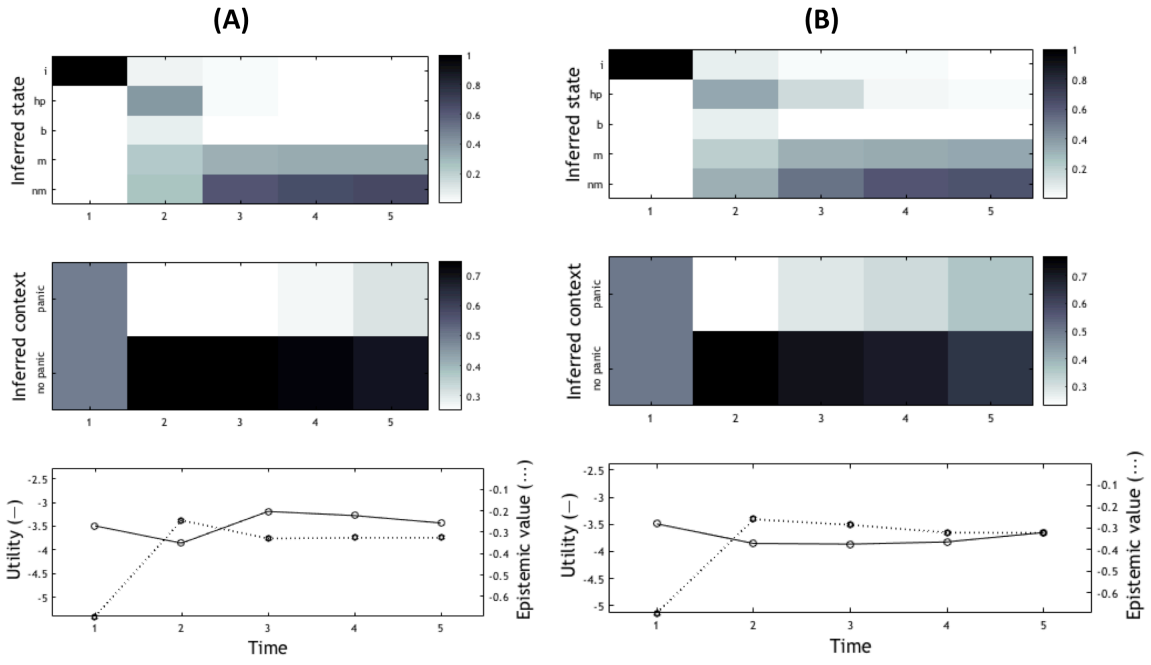

*Figure 5. Results of the fourth simulation, (A) when the agent is endowed with a **B** matrix that has low entropy, as in the first simulation; and (B) when it is endowed with a **B** matrix that has high entropy. Note that in this last case, the generative process that governs the agent's "true" transitions is the same as in the previous simulations (i.e., it has low entropy); it is only the agent's generative model (**B** matrix) that has high entropy.*

The simulation illustrated here shows that an imprecise mapping between successive hidden states (i.e., a **B** matrix having high entropy in the context-to-context transitions) endows the agent with a limited working memory about the context he is in. This implies that the agent can never resolve his contextual uncertainty and needs to continuously monitor his bodily state, hence showing a form of hypervigilance, which is common in some psychopathological conditions, including panic disorder. The Active Inference model shows that hypervigilance may emerge as an adaptive solution to the loss of information within high-entropy generative models.

While we described this loss of information as a failure of working memory, it could be equally considered as an expectation that the current context will not remain stable for long - and hence the future context is uncertain. In Active Inference, contextual uncertainty or the expectation of a contextual change raises attention, by increasing the salience of diagnostic cues (Parr & Friston, 2017a). In keeping, the anticipation of an imminent panic attack and the impossibility to fully resolve contextual uncertainty may keep a person in constant alertness, to detect (for example) initial sensations of heart pounding. As this increased alertness requires extra mental resources — or an "allostatic load" — it may create stress and brain malfunctions (Peters et al., 2017). Furthermore, in clinical conditions, the constant search for safety can "escalate", producing excessive safety behavior even when unnecessary. We address this point in our next simulation.

Fifth simulation: an escalation to excessive clinging to safety

A common clinical issue in panic disorder is an excessive search for safety, which escalates over time. For example, patients may develop clinging and psychological dependence to their alprazolam: they strongly prefer having it with them all the time as a safety signal or also consume it when unnecessary.

The next simulation shows how in Active Inference, the escalation of excessive safety behavior may stem from statistical learning in aversive environmental conditions. This is illustrated in Figure 6, where an Active Inference agent experiences frequent episodes of panic attack in a specific place or situation (e.g., at work) for 45 trials; and it evolves a conditioned (appetitive) response to having alprazolam — which does not disappear, even if panic attacks cease afterwards. Figure 6A shows a single, representative Active Inference agent; whereas Figure 6B shows the average behavior of 128 agents initialized with the same parameters, but which experience slightly different learning trajectories (since the generative process is stochastic).

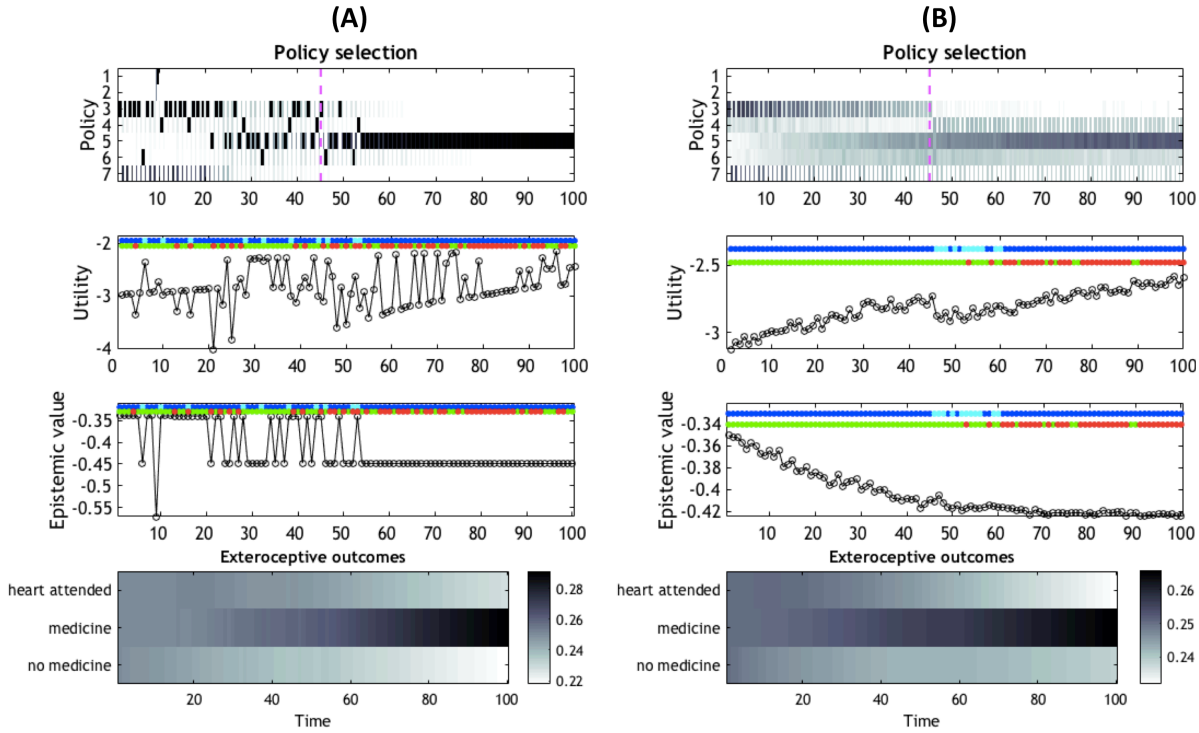

Figure 6. Results of the fifth simulation, in which Active Inference agents learn in an aversive situation, in which panic attacks are frequent (occur 90% of the times) for 45 consecutive trials. After the first 45 trials, panic attacks cease for both agents. The contextual change is marked by the vertical dotted bar. (A) A single representative Active Inference agent. (B) Average behavior of 128 Active Inference agents initialized with the same parameters. See the main text for explanation.

This simulation uses the same generative model as in the first stimulation, with two differences. First, we removed the uninformative *breathlessness* sensation and the hidden state of *having attended to breathlessness*, which (as shown in the first simulation) are irrelevant. Second, and crucially, the agent can learn preferences over its four exteroceptive sensations, one for each its hidden states, and which we label “initial state reached” (not shown in the figures), “heart attended", "medicine with me" and " medicine not with me". While in previous simulations these exteroceptive sensations had neutral value (i.e., their value in the **C** matrix was fixed to zero and could not change), in this simulation they can acquire value as an effect of learning (i.e., their value in the **C** matrix can change over time). Conceptually, this form of valence learning is realized by endowing the model with prior beliefs about the parameters of **C** and treating learning as another form of Bayesian inference: the passage from prior to posterior beliefs about the parameters of **C**. We follow the standard approach in statistical theory to choose a Dirichlet distribution (i.e., a multivariate probability distribution that is the conjugate prior of the categorical distribution used for **C**) as prior. The posterior is another Dirichlet distribution with updated pseudo-counts (i.e., whose parameter values are increased with counts of all the new observations), see the Appendix and (Bishop, 2006).

The effect of this learning process is an increase of the value of "medicine with me" in the **C** matrix when the agent executes policies to avoid panic attacks by taking the medicine— and receives corresponding exteroceptive sensations of "medicine with me". This is a form of *evaluative conditioning*, which refers to a change in the valence of a

(conditioned) stimulus CS that results from pairing it with another, positive or negative (unconditioned) stimulus US (Baeyens et al., 1992; Hofmann et al., 2010). Here, the initially neutral stimulus of having “medicine with me” acquires valence in virtue of the fact of being part of a goal-directed policy that successfully avoids panic attacks and hence achieves a valued outcome (“feeling normal”). In other words, having experienced the positive effects of alprazolam in an aversive situation (i.e., the avoidance of a panic attack) may cause substantial evaluative shifts in the value of the medicine; and subsequently, just having the medicine around (even without taking it) may induce safety feelings.

The first panel of Figure 6A shows the policies selected by the agent over time, when it starts from the same uncertain contextual belief as in the first simulation (i.e., the **D** values for having or not having panic attack are both 0.5). The key thing to appreciate is the crucial difference between the *epistemic policy* of firstly attending to heart pounding and then taking the medicine (policy 3), which is selected more frequently in the first part of the simulation, for about 20 trials; and the *pragmatic policy* to take the medicine immediately and then waiting (policy 5), which is selected afterwards — and persists despite after 45 trials panic attacks cease. Note that in this simulation, we could interpret the action to take medicine equivalently as “bring the medicine”, to consider that one can evolve a form of conditioning to an exteroceptive sensation (e.g., the presence of the medicine) and experience safety feelings (or the reduction of anticipatory fear) without necessarily consuming it. The agent selects the other five policies less frequently (policy 1: staying in the initial state and then taking the medicine; policy 2: staying in the initial state and then not taking the medicine; policy 4: attending to heart pounding and then not taking the medicine; policy 6: taking the medicine immediately and then wait; policy 7: attending to heart pounding twice).

The second and third panels show the exteroceptive and interoceptive outcomes observed by the agent, respectively (top: blue is "medicine with me", cyan is "medicine not with me"; bottom: green is "feel good", red is "feel strange") and its performance in terms of utility and epistemic value, which are calculated at the end of each trial rather than at each time step, as in the previous simulations. The fourth panel shows how the learned prior preferences over exteroceptive sensations ("heart attended", "having the medicine with me" and "not having the medicine with me") change over time (black is low probability, white is high probability). The prior preference for heart pounding increases transiently during the first trials, and then decreases. The prior preference for having the medicine increases over time (and symmetrically, the prior preference for not having the medicine decreases) during all the simulation.

This preference learning mechanism is crucial to establish a form of conditioning, which imbues an initially neutral sensation (the presence of the medicine) with value. In turn, this produces an escalation of safety behavior. While during the first few trials the agent shows the same epistemic behavior illustrated in the first simulation (Figure 2A), afterwards it starts selecting more rigidly a policy to take (or have) the medicine. This happens for two fundamental reasons. First, after preference learning, the value of having the medicine is sufficient to favor the direct selection of a policy to take (or have) it. While in this context having the medicine may seem instrumental to relieve from panic, this is not the case: the policy is selected because the exteroceptive sensation of having the medicine has acquired some value per se, not because it is instrumental for the final outcome of "feeling normal". As the value of the medicine has become independent from the value of the initial goal of relieving from panic, the policy to have the medicine would not fully meet

the criteria for goal-directedness in associative learning theories (Dickinson & Balleine, 1994) – or at least one should consider that the agent's goal has shifted from "feeling normal" to "having the medicine with me".

Second, and equally importantly, the direct selection of a policy to take the medicine abolishes epistemic actions – which implies that the agent remains largely uncertain about its context (panic or no panic) and cannot confidently select instrumental actions.

The combined effect of these two problems (i.e., the fact that having the medicine acquires conditioned value and the abolition of epistemic behavior) becomes evident if one considers that the remaining of the simulation, when the context changes (after trial 45), the agent continues to select the policy to take a medicine, despite its panic attacks cease. Given that it does not execute epistemic actions, the agent cannot fully resolve its uncertainty about having (or not having) a panic attack. Furthermore, it suffers from a "sampling bias": because it keeps selecting always the same policy, it does not have the chance to discover that not taking the medicine is beneficial; and this does not allow the agent to recover efficiently from the initial maladaptive learning.

In sum, this simulation shows that an escalation of excessive safety behavior can arise from statistical learning and conditioning in aversive conditions. This is because the learning process converges prematurely to a poor solution - a phenomenon sometimes called "bad bootstrap" in machine learning, see (Tschantz et al., 2020) - and the model "overfits" to the current situation (frequent panic attacks) and fails to generalize to subsequent situations (absence of panic attacks).

This situation can be contrasted with that of another Active Inference agent, which learns in a significantly less aversive situation, where panic attacks occur 40% of the times. Similar to Figures 6A and 6B, also Figures 7A and 7B show the behavior of a single agent and the average behavior of 128 identical agents, respectively. In this situation, Active Inference agents do not "overfit" but are sensitive to the contextual change. This is evident by considering that after the contextual change, they tend to select the epistemic policy to attend to heart pounding and then take the medicine (policy 4) or the pragmatic policy not to take the medicine (policy 6) - contrary to the agents shown in Figure 6, which tended to select the pragmatic policy to take the medicine (policy 5) instead.

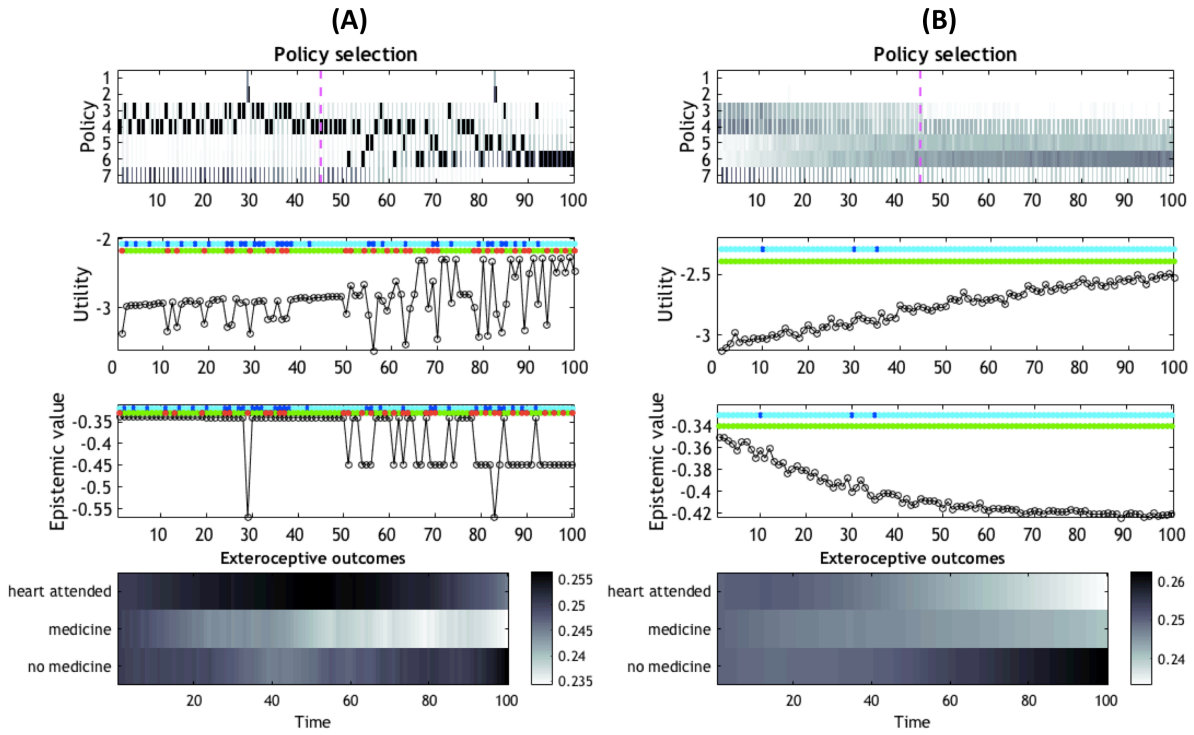

Figure 7. Results of the fifth simulation. This Figure is analogous to Figure 6 but shows the behavior of (A) a single Active Inference agent and (B) the average behavior of 128 Active Inference agents that learn in less aversive situations, where panic attacks occur 40% of the times. See the main text for explanation.

Note that the agent shown in Figure 7A develops not just a preference for "no medicine" but also a preference for "heart pounding" - which may potentially result in forms of hypervigilance. This preference is however is not prevalent in the larger population of 128 agents (Figure 7B), where the preference for "no medicine" largely dominates. This example illustrates that single (initially identical) agents can accidentally develop preferences and behavioral policies that are slightly different from their "group", depending on small differences in their statistical history (despite their broad behavioral patterns would not tend to diverge too much). Hence, idiosyncratic deficits and individual differences in patient populations may depend on statistical history and not just on a different parametrizations of the patients' generative models - a fact that is very difficult to unveil experimentally.

Discussion and Clinical Relevance

We presented an Active Inference model of the perception and misperception of bodily symptoms in the context of panic disorders. While previous predictive coding accounts of psychopathological conditions focus on perceptual processing, the Active Inference framework goes beyond perception and also considers action and planning – and treats all of the above cognitive operations as complementary ways to fulfill a unique normative imperative: free energy minimization (K. J. Friston, 2010).

In the decision-making domain, Active Inference can be considered as a bounded rational model, because it uses approximate (variational) Bayesian inference and it departs from a purely utility maximization scheme, to consider the importance of two complementary kinds of actions: utilitarian actions that achieve preferred states or goals (called *pragmatic actions*, e.g., take or not take the medicine) and actions that resolve

uncertainty (called *epistemic actions*, e.g., attend to or perceive bodily sensations to resolve uncertainty about having or not having a panic attack). Our simulations – and the comparison between Active Inference and a traditional (neuro)economic model that only optimizes utility – show that adaptive regulation requires correctly balancing and sequencing the two kinds of actions (e.g., resolving uncertainty before deciding whether to take a medicine). Note that while most accounts of perceptual processing (and psychopathology) treat perception as a passive process, here we treated the act of attending to one's own bodily sensations as a form of choice. While it may be assumed that, at some level, the peripheral system continuously informs the brain about the body state, here we wanted to model a more subtle process: the decision to attend to one's own bodily signals, to render them more precise and hence informative.

In our simulation, the inference about the context corresponds to the perception (or misperception) of one's own bodily state. This is in keeping with the standard assumption of Bayesian theories of perception, where the "percept" is the hypothesis having the highest probability. One can also consider that inferring that one is about to have a panic attack also entails the perception (or misperception) of panic-associated bodily sensations, such as heart pounding or breathlessness, even when they are not there. This is because an assumption of the Bayesian framework is that the (generative) architecture of the brain is continuously generating predictions about incoming sensory observations (e.g., heart pounding) that are to be expected under the current hypothesis (e.g., that one is about to have a panic attack). Under normal conditions, when incorrect (top-down) predictions are generated, these would be rapidly falsified by checking with (bottom-up) sensory streams; and in turn, this would permit to revise the incorrect hypothesis that has motivated the prediction in the first place (i.e., the hypothesis that one is having a panic attack). However, this process can go awry in several ways, possibly producing maladaptive inference and action selection, and biased perception of the bodily state.

The latter is more the rule than the exception. An analysis of clinical data on heartbeat perception in panic disorder showed that most perceptions are inaccurate, suggesting that “once a patient with panic disorder perceives a situation as threatening, an `anxiety' schema is activated, and that perception of symptoms is more guided by the schema (that is, by past information) than based upon present physiological status. This hypothetical anxiety schema would include shifting of attentional focus, selective perception, high HR, other arousal symptoms and anxiety” (Van der Does et al., 2000, p. 61). Also excessive safety behaviors (e.g. clinging to benzodiazepines, unnecessarily taking medicine) are well-known in panic disorder (Fujii et al., 2015; Hamm et al., 2014).

Our simulations illustrated various situations that are strikingly consistent with these clinical phenomena. We considered four main causes of suboptimal inference underlying these clinical features. The first three causes correspond to incorrect parametrizations of the generative model; and namely, an incorrect prior belief which can produce rigid and uninformed decisions; an incorrect mapping between hidden states and bodily sensations which can produce the disappearance of epistemic behavior; and an incorrect mapping between hidden states at consecutive time steps which can produce hypervigilance. The fourth cause of suboptimal inference is the over-exposure to an aversive learning environment, which can produce excessive safety behavior. We discuss these four points and their importance for psychopathology in the next four paragraphs.

Incorrect prior belief. Theories of psychopathological conditions and interoceptive processing framed within the predictive coding framework highlighted that when prior

beliefs are excessively strong or are assigned excessively high precision, they can dominate the inference, thus preventing sensory information to be correctly integrated - and possibly generating maladaptive inference, illusions or disorders of the body schema (Adams et al., 2013; Barca & Pezzulo, 2020; Barrett et al., 2016; Edwards et al., 2012; Iodice et al., 2019; Janssens et al., 2009; Pezzulo, 2013; Pezzulo et al., 2015; Pezzulo, Iodice, et al., 2018; Seth, 2013; Stephan et al., 2016; Sterzer et al., 2018; Van den Bergh et al., 2017). The Active Inference approach adopted in this study extends predictive coding to the domains of action selection and planning, and highlights how incorrect priors could produce maladaptive behavior and not just maladaptive perception. Specifically, an excessively high prior (corresponding to the agent's **D** vector) prevents epistemic actions to correctly infer one's state. It promotes, for example, false heartbeat perceptions and favors rigid and uninformed decisions (e.g., taking a medicine even when one is not having a panic attack). Note that there are two subtle differences between previous treatments of priors in predictive coding and our treatment in active inference. First, in active inference, there are two kinds of priors: priors over initial states (in the **D** vector) and over observations (in the **C** vector). The latter plays the role of preferred states and goals; in our simulation, it is the prior of "feeling normal" that motivates action. Second, predictive coding is formulated in continuous time; and priors have a precision that, if the distribution is Gaussian, corresponds to the inverse of the variance of the distribution. The Active Inference system that we used works in discrete time, as a Markov decision process; and the precision of priors cannot be defined in the same way as in continuous time predictive coding. In our simulations, we have assumed that the prior over observations, which corresponds roughly to prior preferences (**C** vector), cannot change; whereas the priors over initial states (**D** vector) are used to infer in which state the agent is, along with observations, according to Bayes rule. Bayes rule combines prior information and observations by assigning the same weight (so to say) to both; which would be analogous to having the same precision for prior and observations in continuous time formulations.

Incorrect mapping between hidden states and observations. This cause of suboptimal behavior corresponds to the precision of the agent's **A** matrix, and in particular the causes of heart pounding and breathlessness sensations. This incorrect mapping reduces epistemic actions when its entropy is too high and increases them when its entropy is too low. These examples show that setting accurate precision levels is essential to adaptive inference. Note that in our simulations, epistemic actions are limited to attending to bodily sensations, but in practice, they encompass any information gathering action, such as for example visually attending or enquiring somebody to obtain relevant information for a decision at hand. This implies that the model presented here can be readily extended to include additional attention dynamics. For example, it is known that panic disorder with agoraphobia often develops as a complication of panic disorder. Patients learn to pay attention not only to bodily sensations, but also to any other external event that is relevant to the inference about panic attack, such as places and situations where previous experiences of panic took place, could potentially happen or would have more dramatic consequences (Starcevic et al., 1993).

Incorrect mapping between hidden contexts at consecutive time steps. This cause of suboptimal behavior corresponds to the precision of the agent's **B** matrix, and in particular the context-to-context mappings. This incorrect mapping implies that the agent's confidence about its current context (e.g. being in a "no panic" context) decreases with time and hence the agent needs to continuously monitor its bodily states (hypervigilance), by

executing an excessive number of epistemic actions. This sort of hypervigilance may seem odd but is, in fact, an optimal behavior mandated by excessive uncertainty in the internal (transition) model. It is typically observed in panic disorder (Schmidt et al., 1997).

Miscalibration of the parameters of the generative model. This last source of suboptimal behavior can arise from over-exposure to aversive learning conditions dominated by frequent negative events (e.g., frequent panic attacks at work). In such conditions, it is adaptive to recalibrate the parameters of the generative model to expect more aversive events (e.g., set thresholds for aversive events to lower levels) and prepare policies to deal with them (e.g., prepare to take a medicine). In our simulations, we have illustrated this recalibration by giving the agent the possibility to learn prior preferences over exteroceptive outcomes (but note that other forms of learning are possible; e.g., learning the prior probability of panic attacks in specific situations, such as at work). In any case, the recalibration of model parameters implies that aversive events may remain expected, causing the agent to need more evidence to change predictions or policies, even when the situation ceases to be aversive (e.g., when the panic attacks reduce significantly). This fits in with overgeneralization of fear learning and impaired safety learning to safety cues, which has been extensively documented in panic disorder (Lissek et al., 2009, 2010).

The model advanced here is in many ways an impoverished description of the clinical reality. Nevertheless, it shows that even a simple model is able to account for quite a number of relevant clinical features. For example, the model describes the conditions in which one may progressively diminish attention to the body, when somatic information entails (or is expected to entail) poor information gain. The model also describes how hypervigilance will develop as an active form of epistemic foraging, given an excessive uncertainty in the transition model, which causes beliefs about panic attack increasingly less precise with time (and of course because establishing whether or not one is about to have a panic attack is highly important for the agent). Eventually, maladaptive patterns of behavior such as poor somatic attention and hypervigilance may also become chronic, as an effect of a progressive transition from goal-directed to more habitual forms of adaptive control (K. J. Friston, FitzGerald, Rigoli, Schwartenbeck, O'Doherty, et al., 2016), as an effect of allostatic load and chronic stress (Peters et al., 2017), or because the presence of highly precise prior beliefs about having a panic attack may render the agent rather impermeable for new conflicting evidence and, hence, for adaptive updating.

Model predictions

This model makes also a number of predictions about the ways different manifestations of panic disorders arise from specific (maladaptive) parametrizations of the generative model. First, it predicts that both excessively high prior beliefs and excessively low precision of interoceptive channels (or their combination) favor rigid and uninformed decisions. This is because misregulated prior or interoceptive precision parameters may prevent interoceptive channels to be attended, and uncertainty about one's condition (e.g., panic or not panic) to be resolved, before a choice to take or not to take a medicine. The possible dysfunctions associated to imprecise interoceptive channels are currently under scrutiny, but further evidence is necessary to assess their roles in psychopathological conditions (Barca & Pezzulo, 2020; Smith et al., 2020).

Second, this model predicts that hypervigilance may stem from an imprecise (high entropy) transition model, which causes a loss of confidence in the current contextual estimate, which results in an irresolvable contextual uncertainty (and in some cases, a

constant expectation of imminent panic attacks). The excessive monitoring of interoceptive states may result from the urge to minimize this irresolvable contextual uncertainty.

Third, the model predicts that an excessive concern for safety may emerge as a consequence of learning in aversive conditions, such as the experience of frequent panic attacks. In our simulation, the excessive concern is situation-specific: it arises in the specific situation (e.g., at work) where frequent panic attacks have been experienced. For this, the model also allows understanding how context specificity may become a feature of panic disorder (i.e., avoiding some places). Note that the model suggests that excessive safety behaviors may emerge because of an "overfit" to the aversive situation, which produces an escalation of safety measures and a suppression of epistemic behavior. In turn, the suppression of epistemic behavior makes a person unable to attend to (interoceptive or exteroceptive) cues that may signal a contextual change. The model therefore suggests that promoting the resurgence of epistemic behavior by increasing the saliency of interoceptive cues may potentially help notice contextual changes and reduce the excessive safety behaviors. In fact, new epistemic behavior in different forms is promoted in exposure therapy resulting in new (inhibitory) learning. Especially exposure to interoceptive cues is one of the most effective therapies for panic disorder (Craske et al., 2014; Pompoli et al., 2018).

It is worth noting that our model is intended as a general framework to explain panic attacks, but it can be *personalized* to describe more specifically the idiosyncratic characteristics of different individuals within the overarching framework offered by the model. For example, in our simulations, we assumed that “feeling normal” is positive, whereas “feeling strange” is negative and the bodily sensations of heart pounding and breathlessness are slightly aversive as they are in reality. Yet the specific values associated with “feeling normal”, “feeling strange” and the bodily sensations are person-specific. Other parameters, such as the precisions of A and B matrices, are person-specific, too. This diversity gives room to include in the model individual difference variables that play a role in panic disorder, such as anxiety sensitivity (McNally, 2002) or intolerance for uncertainty (Carleton et al., 2014). The former can be represented by parameters that quantify the drive to obtain a state of “feeling normal” (or, conversely, the aversiveness of a state of feeling strange); whereas the latter can be represented by excessive entropy of the B matrix, as shown in the fourth simulation. These and additional individual difference variables are key to characterize quantitatively the behavior (and the inference) of different individuals - or to "phenotype" them computationally (P. Schwartenbeck & Friston, 2016).

Note that in the simulations illustrated in this paper, subjects can cause their own outcomes and sensations through taking medicine (or not). In other words, they can engage certain behaviors to induce state transitions in the world (and body) generating sensory outcomes. Crucially, we could have also included autonomic or interoceptive actions (e.g., as mediated by autonomic reflexes) based upon the interoceptive predictions (Tschantz et al., 2021). In other words, when attending to heart pounding in the presence of a panic attack, the subject could actually induce heart pounding in a way that is not dissimilar from the simulations of tachycardia reported in (Allen et al., 2019). This raises the interesting opportunity to simulate the induction of a panic attack simply through aberrant belief updating in which the sensations that provide evidence for a panic attack is actually generated during the process of inferring whether one is having a panic attack. In fact, such vicious circle dynamic has been proposed as an important clinical feature of panic disorder (Bouton et al., 2001).

Conclusions

We presented a fully specified computational model of adaptive and maladaptive symptom perception and behavioral control in the context of panic disorder. This model goes beyond most previous attempts to characterize psychopathological conditions that only focus on perception dynamics (under a predictive coding framework), to also include action control and attention dynamics (under an Active Inference framework). Furthermore, the model goes beyond current theoretical (qualitative) attempts to describe psychopathologies, as it is able to derive *quantitative* predictions about what factors may determine maladaptive inference and behavior and can be parameterized differently for different individuals, to account for their individual differences. Moving from qualitative to quantitative (and personalized) models is a key step to advance the field of computational psychiatry from theory to clinical practice (K. J. Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2016; Petzschner et al., 2017; Stephan et al., 2016) - and this model represents a first attempt in this important research direction.

References

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/fpsyt.2013.00047

Allen, M., Levy, A., Parr, T., & Friston, K. J. (2019). In the Body's Eye: The Computational Anatomy of Interoceptive Inference. *BioRxiv*, 603928.

Ashby, W. R. (1952). *Design for a brain* (Vol. ix). Wiley.

Baeyens, F., Eelen, P., Crombez, G., & van den Bergh, O. (1992). Human evaluative conditioning: Acquisition trials, presentation schedule, evaluative style and contingency awareness. *Behaviour Research and Therapy*, *30*(2), 133–142. https://doi.org/10.1016/0005-7967(92)90136-5

Barca, L., & Pezzulo, G. (2020). Keep your interoceptive streams under control: An active inference perspective on anorexia nervosa. *Cognitive, Affective, & Behavioral Neuroscience*, *20*(2), 427–440. https://doi.org/10.3758/s13415-020-00777-6

Barrett, L. F. (2017). The theory of constructed emotion: An active inference account of interoception and categorization. *Social Cognitive and Affective Neuroscience*, *12*(1), 1–23. https://doi.org/10.1093/scan/nsw154

Barrett, L. F., Quigley, K. S., & Hamilton, P. (2016). An active inference theory of allostasis and interoception in depression. *Phil. Trans. R. Soc. B*, *371*(1708), 20160011.

Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience*, *16*(7), 419–429. https://doi.org/10.1038/nrn3950

Beal, M. J. (2003). *Variational algorithms for approximate Bayesian inference*. University of London. http://www.cse.buffalo.edu/faculty/mbeal/papers/beal03.pdf

Berntson, G. G., Gianaros, P. J., & Tsakiris, M. (2018). Interoception and the autonomic nervous system: Bottom-up meets top-down. *The Interoceptive Mind: From Homeostasis to Awareness*, 1.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

Botvinick, M., & Toussaint, M. (2012). Planning as inference. *Trends in Cognitive Sciences*, *16*(10), 485–488. https://doi.org/10.1016/j.tics.2012.08.006

Bouton, M. E., Mineka, S., & Barlow, D. H. (2001). A modern learning theory perspective on the etiology of panic disorder. *Psychological Review*, *108*(1), 4.

Carleton, R. N., Duranceau, S., Freeston, M. H., Boelen, P. A., McCabe, R. E., & Antony, M. M. (2014). "But it might be a heart attack": Intolerance of uncertainty and panic disorder symptoms. *Journal of Anxiety Disorders*, *28*(5), 463–470. https://doi.org/10.1016/j.janxdis.2014.04.006

Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.

Corlett, P. R., Horga, G., Fletcher, P. C., Alderson-Day, B., Schmack, K., & Powers, A. R. (2019). Hallucinations and Strong Priors. *Trends in Cognitive Sciences*, *23*(2), 114–127. https://doi.org/10.1016/j.tics.2018.12.001

Craig, A. D. (2015). *How do you feel?: An interoceptive moment with your neurobiological self*. Princeton University Press.

Craske, M. G., Treanor, M., Conway, C., Zbozinek, T., & Vervliet, B. (2014). Maximizing Exposure Therapy: An Inhibitory Learning Approach. *Behaviour Research and Therapy*, *58*, 10–23. https://doi.org/10.1016/j.brat.2014.04.006

Cullen, M., Davey, B., Friston, K. J., & Moran, R. J. (2018). Active Inference in OpenAI Gym: A Paradigm for Computational Investigations Into Psychiatric Illness. *Biological Psychiatry.*

Cognitive Neuroscience and Neuroimaging, *3*(9), 809–818. https://doi.org/10.1016/j.bpsc.2018.06.010
Dickinson, A., & Balleine, B. (1994). Motivational control of goal-directed action. *Animal Learning & Behavior*, *22*(1), 1–18. https://doi.org/10.3758/BF03199951
Donnarumma, F., Costantini, M., Ambrosini, E., Friston, K., & Pezzulo, G. (2017). Action perception as hypothesis testing. *Cortex*. https://doi.org/10.1016/j.cortex.2017.01.016
Donnarumma, F., Maisto, D., & Pezzulo, G. (2016). Problem Solving as Probabilistic Inference with Subgoaling: Explaining Human Successes and Pitfalls in the Tower of Hanoi. *PLOS Computational Biology*, *12*(4), e1004864. https://doi.org/10.1371/journal.pcbi.1004864
Edwards, M. J., Adams, R. A., Brown, H., Pareés, I., & Friston, K. J. (2012). A Bayesian account of "hysteria." *Brain: A Journal of Neurology*, *135*(Pt 11), 3495–3512. https://doi.org/10.1093/brain/aws129
Engel, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions: Oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience*, *2*(10), 704–716.
Friston, K. J. (2005). A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci*, *360*(1456), 815–836. https://doi.org/10.1098/rstb.2005.1622
Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nat Rev Neurosci*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787
Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, *68*, 862–879. https://doi.org/10.1016/j.neubiorev.2016.06.022
Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2016). Active Inference: A Process Theory. *Neural Computation*, 1–49. https://doi.org/10.1162/NECO_a_00912
Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active Inference: A Process Theory. *Neural Comput*, *29*(1), 1–49. https://doi.org/10.1162/NECO_a_00912
Friston, K. J., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2014). The anatomy of choice: Dopamine and decision-making. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *369*(1655), 20130481. https://doi.org/10.1098/rstb.2013.0481
Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: The brain as a phantastic organ. *The Lancet Psychiatry*, *1*(2), 148–158. https://doi.org/10.1016/S2215-0366(14)70275-5
Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, *0*(0), 1–28. https://doi.org/10.1080/17588928.2015.1020053
Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: Optimal control without cost functions. *Biological Cybernetics*, *106*(8–9), 523–541.
Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2013). The anatomy of choice: Active inference and agency. *Frontiers in Human Neuroscience*, *7*, 598. https://doi.org/10.3389/fnhum.2013.00598
Fujii, K., Uchida, H., Suzuki, T., & Mimura, M. (2015). Dependence on benzodiazepines in patients with panic disorder: A cross-sectional study. *Psychiatry and Clinical Neurosciences*, *69*(2), 93–99. https://doi.org/10.1111/pcn.12203
Glimcher, P. W., & Rustichini, A. (2004). Neuroeconomics: The consilience of brain and

decision. *Science*, *306*(5695), 447–452. https://doi.org/10.1126/science.1102566
Hamm, A. O., Richter, J., & Pané-Farré, C. A. (2014). When the threat comes from inside the body: A neuroscience based learning perspective of the etiology of panic disorder. *Restorative Neurology and Neuroscience*, *32*(1), 79–93. https://doi.org/10.3233/RNN-139011
Henningsen, P., Gündel, H., Kop, W. J., Löwe, B., Martin, A., Rief, W., Rosmalen, J. G. M., Schröder, A., van der Feltz-Cornelis, C., Van den Bergh, O., & EURONET-SOMA Group. (2018). Persistent Physical Symptoms as Perceptual Dysregulation: A Neuropsychobehavioral Model and Its Clinical Implications. *Psychosomatic Medicine*, *80*(5), 422–431. https://doi.org/10.1097/PSY.0000000000000588
Hofmann, W., De Houwer, J., Perugini, M., Baeyens, F., & Crombez, G. (2010). Evaluative conditioning in humans: A meta-analysis. *Psychological Bulletin*, *136*(3), 390–421. https://doi.org/10.1037/a0018916
Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
Iodice, P., Porciello, G., Bufalari, I., Barca, L., & Pezzulo, G. (2019). An interoceptive illusion of effort induced by false heart-rate feedback. *Proceedings of the National Academy of Sciences*, *116*(28), 13897–13902. https://doi.org/10.1073/pnas.1821032116
Janssens, T., Verleden, G., De Peuter, S., Van Diest, I., & Van den Bergh, O. (2009). Inaccurate perception of asthma symptoms: A cognitive-affective framework and implications for asthma treatment. *Clinical Psychology Review*, *29*(4), 317–327. https://doi.org/10.1016/j.cpr.2009.02.006
Lissek, S., Rabin, S., Heller, R. E., Lukenbaugh, D., Geraci, M., Pine, D. S., & Grillon, C. (2010). Overgeneralization of conditioned fear as a pathogenic marker of panic disorder. *The American Journal of Psychiatry*, *167*(1), 47–55. https://doi.org/10.1176/appi.ajp.2009.09030410
Lissek, S., Rabin, S. J., McDowell, D. J., Dvir, S., Bradford, D. E., Geraci, M., Pine, D. S., & Grillon, C. (2009). Impaired discriminative fear-conditioning resulting from elevated fear responding to learned safety cues among individuals with panic disorder. *Behaviour Research and Therapy*, *47*(2), 111–118. https://doi.org/10.1016/j.brat.2008.10.017
Loewenstein, G., Rick, S., & Cohen, J. D. (2007). Neuroeconomics. *Annual Review of Psychology*, *59*(1), 647–672. https://doi.org/10.1146/annurev.psych.59.103006.093710
Lynn, S. K., & Barrett, L. F. (2014). "Utilizing" Signal Detection Theory. *Psychological Science*, *25*(9), 1663–1673. https://doi.org/10.1177/0956797614541991
Maisto, D., Friston, K., & Pezzulo, G. (2019). Caching mechanisms for habit formation in Active Inference. *Neurocomputing*, *359*, 298–314. https://doi.org/10.1016/j.neucom.2019.05.083
McNally, R. J. (2002). Anxiety sensitivity and panic disorder. *Biological Psychiatry*, *52*(10), 938–946. https://doi.org/10.1016/s0006-3223(02)01475-0
Parr, T., & Friston, K. J. (2017a). Working memory, attention, and salience in active inference. *Scientific Reports*, *7*(1), 14678.
Parr, T., & Friston, K. J. (2017b). Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface*, *14*(136), 20170376. https://doi.org/10.1098/rsif.2017.0376
Paulus, M. P. (2019). Driven by Pain, Not Gain: Computational Approaches to Aversion-Related Decision Making in Psychiatry. *Biological Psychiatry*. https://doi.org/10.1016/j.biopsych.2019.08.025
Paulus, M. P., Feinstein, J. S., & Khalsa, S. S. (2019). An Active Inference Approach to Interoceptive Psychopathology. *Annual Review of Clinical Psychology*, *15*(1), 97–122.

https://doi.org/10.1146/annurev-clinpsy-050718-095617
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
Petzschner, F. H. (2017). Stochastic Dynamic Models for Computational Psychiatry and Computational Neurology. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, *2*(3), 214–215. https://doi.org/10.1016/j.bpsc.2017.03.003
Petzschner, F. H., Weber, L. A. E., Gard, T., & Stephan, K. E. (2017). Computational Psychosomatics and Computational Psychiatry: Toward a Joint Framework for Differential Diagnosis. *Biological Psychiatry*, *82*(6), 421–430. https://doi.org/10.1016/j.biopsych.2017.05.012
Pezzulo, G. (2013). Why do you fear the Bogeyman? An embodied predictive coding model of perceptual inference. *Cognitive, Affective, and Behavioral Neuroscience*.
Pezzulo, G., Donnarumma, F., Iodice, P., Maisto, D., & Stoianov, I. (2017). Model-Based Approaches to Active Perception and Control. *Entropy*, *19*(6), 266. https://doi.org/10.3390/e19060266
Pezzulo, G., Iodice, P., Barca, L., Chausse, P., Monceau, S., & Mermillod, M. (2018). Increased heart rate after exercise facilitates the processing of fearful but not disgusted faces. *Scientific Reports*, *8*(1), 398. https://doi.org/10.1038/s41598-017-18761-5
Pezzulo, G., Maisto, D., Barca, L., & Van den Bergh, O. (2019). Symptom Perception From a Predictive Processing Perspective. *Clinical Psychology in Europe*, e35952,.
Pezzulo, G., Rigoli, F., & Chersi, F. (2013). The Mixed Instrumental Controller: Using Value of Information to combine habitual choice and mental simulation. *Frontiers in Cognition*, *4*, 92. https://doi.org/10.3389/fpsyg.2013.00092
Pezzulo, G., Rigoli, F., & Friston, K. (2018). Hierarchical Active Inference: A Theory of Motivated Control. *Trends in Cognitive Sciences*, *22*(4), 294–306. https://doi.org/10.1016/j.tics.2018.01.009
Pezzulo, G., Rigoli, F., & Friston, K. J. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, *136*, 17–35.
Pompoli, A., Furukawa, T. A., Efthimiou, O., Imai, H., Tajika, A., & Salanti, G. (2018). Dismantling cognitive-behaviour therapy for panic disorder: A systematic review and component network meta-analysis. *Psychological Medicine*, *48*(12), 1945–1953. https://doi.org/10.1017/S0033291717003919
Powers, A. R., Mathys, C., & Corlett, P. R. (2017). Pavlovian conditioning–induced hallucinations result from overweighting of perceptual priors. *Science*, *357*(6351), 596–600. https://doi.org/10.1126/science.aan3458
Powers, W. T. (1973). *Behavior: The Control of Perception*. Aldine.
Quadt, L., Critchley, H. D., & Garfinkel, S. N. (2018). The neurobiology of interoception in health and disease. *Annals of the New York Academy of Sciences*, *1428*(1), 112–128. https://doi.org/10.1111/nyas.13915
Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci*, *2*(1), 79–87. https://doi.org/10.1038/4580
Rigoli, F., Pezzulo, G., Dolan, R., & Friston, K. (2017). A Goal-Directed Bayesian Framework for Categorization. *Frontiers in Psychology*, *8*, 408. https://doi.org/10.3389/fpsyg.2017.00408
Schmidt, N. B., Lerew, D. R., & Trakowski, J. H. (1997). Body vigilance in panic disorder:

Evaluating attention to bodily perturbations. *Journal of Consulting and Clinical Psychology*, *65*(2), 214–220. https://doi.org/10.1037//0022-006x.65.2.214
Schwartenbeck, P., & Friston, K. (2016). Computational Phenotyping in Psychiatry: A Worked Example. *ENeuro*, *3*(4). https://doi.org/10.1523/ENEURO.0049-16.2016
Schwartenbeck, Philipp, FitzGerald, T. H., Mathys, C., Dolan, R., & Friston, K. (2014). The dopaminergic midbrain encodes the expected certainty about desired outcomes. *Cerebral Cortex*, bhu159.
Schwartenbeck, Philipp, Passecker, J., Hauser, T. U., FitzGerald, T. H., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. *ELife*, *8*, e41703. https://doi.org/10.7554/eLife.41703
Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, *17*(11), 565–573. https://doi.org/10.1016/j.tics.2013.09.007
Seth, A. K. (2014). The cybernetic Bayesian brain: From interoceptive inference to sensorimotor contingencies. *MIND Project, Eds. T. Metzinger & J. Windt*.
Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Phil. Trans. R. Soc. B*, *371*(1708), 20160007.
Seth, A. K., & Tsakiris, M. (2018). Being a Beast Machine: The Somatic Basis of Selfhood. *Trends in Cognitive Sciences*, *22*(11), 969–981. https://doi.org/10.1016/j.tics.2018.08.008
Smith, R., Kuplicki, R., Feinstein, J., Forthman, K. L., Stewart, J. L., Paulus, M. P., Investigators, T. 1000, & Khalsa, S. S. (2020). A Bayesian computational model reveals a failure to adapt interoceptive precision estimates across depression, anxiety, eating, and substance use disorders. *MedRxiv*, 2020.06.03.20121343. https://doi.org/10.1101/2020.06.03.20121343
Starcevic, V., Kellner, R., Uhlenhuth, E. H., & Pathak, D. (1993). The phenomenology of panic attacks in panic disorder with and without agoraphobia. *Comprehensive Psychiatry*, *34*(1), 36–41. https://doi.org/10.1016/0010-440x(93)90033-z
Stephan, K. E., Manjaly, Z. M., Mathys, C. D., Weber, L. A. E., Paliwal, S., Gard, T., Tittgemeyer, M., Fleming, S. M., Haker, H., Seth, A. K., & Petzschner, F. H. (2016). Allostatic Self-efficacy: A Metacognitive Theory of Dyshomeostasis-Induced Fatigue and Depression. *Frontiers in Human Neuroscience*, *10*. https://doi.org/10.3389/fnhum.2016.00550
Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., Petrovic, P., Uhlhaas, P., Voss, M., & Corlett, P. R. (2018). The Predictive Coding Account of Psychosis. *Biological Psychiatry*, *84*(9), 634–643. https://doi.org/10.1016/j.biopsych.2018.05.015
Tschantz, A., Barca, L., Maisto, D., Buckley, C. L., Seth, A. K., & Pezzulo, G. (2021). Simulating homeostatic, allostatic and goal-directed forms of interoceptive control using Active Inference. *BioRxiv*, 2021.02.16.431365. https://doi.org/10.1101/2021.02.16.431365
Tschantz, A., Seth, A. K., & Buckley, C. L. (2020). Learning action-oriented models through active inference. *PLOS Computational Biology*, *16*(4), e1007805. https://doi.org/10.1371/journal.pcbi.1007805
Van den Bergh, O., Witthöft, M., Petersen, S., & Brown, R. J. (2017). Symptoms and the body: Taking the inferential leap. *Neuroscience and Biobehavioral Reviews*, *74*(Pt A), 185–203. https://doi.org/10.1016/j.neubiorev.2017.01.015
Van den Bergh, O., Zacharioudakis, N., & Petersen, S. (2018). Interoception, symptom perception and categorization. In *The Interoceptive Mind. From homeostasis to awareness* (pp. 212–226).
Van der Does, A. J. W., Antony, M. M., Ehlers, A., & Barsky, A. J. (2000). Heartbeat perception in panic disorder: A reanalysis. *Behaviour Research and Therapy*, *38*(1), 47–62.

https://doi.org/10.1016/s0005-7967(98)00184-3
Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. The MIT Press.
Zamariola, G., Maurage, P., Luminet, O., & Corneille, O. (2018). Interoceptive accuracy scores from the heartbeat counting task are problematic: Evidence from simple bivariate correlations. *Biological Psychology*, *137*, 12–17.

Appendix: A formal introduction to Active Inference

Active Inference is a formal framework that integrates the cybernetic concepts of feedback and error control (Ashby, 1952; W. T. Powers, 1973; Wiener, 1948) with a Bayesian inferential scheme (Pezzulo et al., 2015; Pezzulo, Rigoli, et al., 2018; Seth, 2014). In Active Inference, perception and action (or policy) selection form a closed-loop process, whose execution can be cast in terms of approximate Bayesian inference (Botvinick & Toussaint, 2012; Donnarumma et al., 2016; Pezzulo et al., 2013, 2017), which is rendered tractable using a variational approximation stemming from the free-energy-minimization principle (K. Friston et al., 2012).

In Active Inference, an agent adopts an *internal generative model* to understand its observations and how they may be generated by external, environmental dynamics (called *generative process* in Active Inference). As shown in Figure A1, the generative model includes hidden *states* s as causes of the observed *outcomes* o. Hidden states move forward in time controlled by a *policy* (sequence of actions) π that depend on a continuous variable: the *precision* γ of the policies.

To select optimal actions, an Active Inference agent needs to evaluate all its policies π, for any possible future state an agent could be in. That means computing the (negative) expected free energy G_π of each policy π. In turn, this requires Active Inference agents to have prior beliefs $P(o_\tau)$ about outcomes o_τ it will experience, and their likelihood $P(o_\tau|s_\tau)$, namely, the conditional distribution of outcomes under the (hidden) states. This allows formulating G_π as:

$$\mathrm{G}_\pi = \underbrace{D_{KL}[Q(\tilde{o}|\pi)||P(\tilde{o})]}_{\text{risk}} - \underbrace{\mathrm{E}_{Q(\tilde{s}|\pi)}\big[\mathrm{H}[P(\tilde{o}|\tilde{s})]\big]}_{\text{expected ambiguity}}$$

where $\tilde{s}$ and $\tilde{o}$ are sequences of states and outcomes, respectively, determined by the application of the policy π.

Note that the value (or quality) of a policy π is formally the sum of two terms. The former (*risk*) term is the Kullback-Leibler divergence between the posteriors $Q(\tilde{o}|\pi)$ and the priors $P(\tilde{o})$ over the outcomes; and is related to the agent's preferences. The second (*expected ambiguity*) is the entropy $\mathrm{H}[P(\tilde{o}|\tilde{s})]$ expected under the posteriors over hidden states $Q(\tilde{s}|\pi)$, which represents the inaccuracy due to a mapping between states and outcomes in relation to the posteriors about the state of the world (K. Friston et al., 2015). One can consider these two terms as the extent to which the policy will allow achieving the agent's goals ("pragmatic value") and the capacity of the policy to reduce uncertainty ("epistemic value") by disambiguating states, respectively. In other words, for a given policy, the pragmatic value measures the difference between predicted and preferred expected outcome in the future while the epistemic value quantifies how much expecting to be in future state diminishes uncertainty on future outcomes. From a machine learning perspective, this would be equivalent to say that G_π includes a "regularisation" term, which balances between exploitive (pragmatic) and exploratory (epistemic) behaviour.

Once every potential policy has been scored (and its associated "quality" value is set), the agent selects the action that minimises the expected divergence between the outcomes predicted at the next times step through the policy and the outcome predicted after each action. Hence, in Active Inference, an action is the result of an inferential process that scores possible futures. This Active Inference scheme has been used to address a variety of cognitive phenomena, including decision-making (K. Friston et al., 2013; K. J.

Friston, Schwartenbeck, et al., 2014), habitual behavior, salience and curiosity-driven planning (K. J. Friston, FitzGerald, Rigoli, Schwartenbeck, O'Doherty, et al., 2016; Maisto et al., 2019; Parr & Friston, 2017b; Philipp Schwartenbeck et al., 2019), and in general to develop a process theory for neural computation (K. J. Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2016; Philipp Schwartenbeck et al., 2014) as well as for investigating psychiatric disorders (Barrett et al., 2016; Cullen et al., 2018; K. J. Friston, Stephan, et al., 2014).

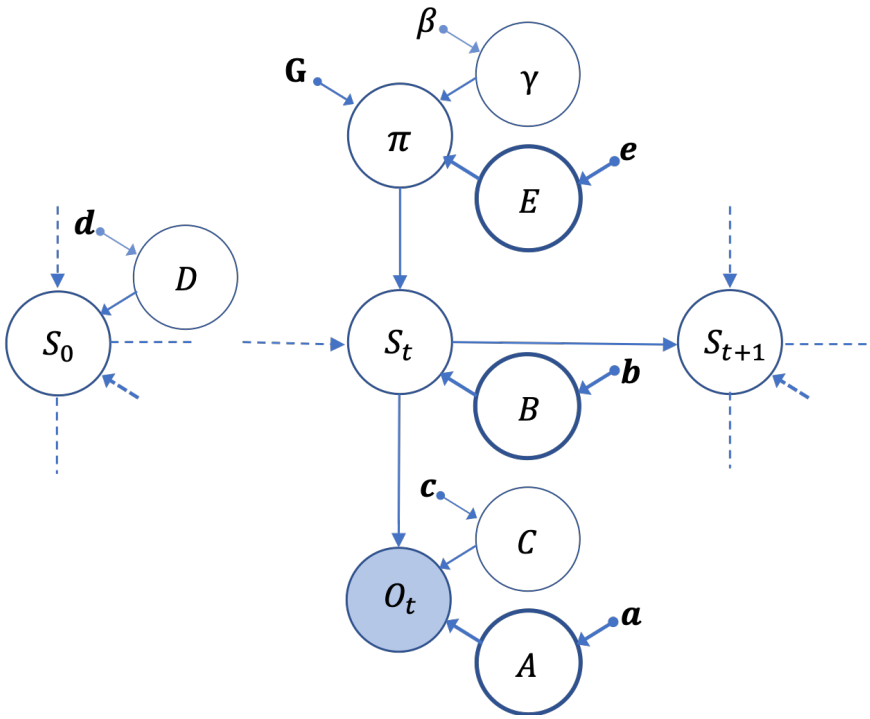

Figure A1. Graphical model for Active Inference. See the main text for explanation.

Formally, the discrete version of Active Inference can be described as a Partially Observed Markov Decision Process (POMDP) represented by a tuple $(S, O, U, \Theta, R, P, Q)$ where:

- S is the set of agent's hidden states s by which the agent infers the environmental state;
- O is the set of observations o;
- U is the set of actions u. A sequence of actions $\tilde{u}$ is called *policy* and denoted as π. Thus, $\pi = [u_t, \dots, u_T] = [\pi(t), \dots, \pi(T)]$;
- $\Theta = (A, B, C, D, E, \gamma)$ is a tuple of additional variables introduced to regulate the S, O, U distributions. Their expected values play a crucial role during the inference and their priors are usually updated through the posteriors, at a slower timescale (corresponding to learning dynamics) at the end of each trial. In all the simulations shown in this paper except the fifth, the distributions of the variables in Θ are static and fixed in advance.
- $R(\tilde{o}, \tilde{s}, \tilde{u})$ is a generative process defined over the outcomes, the states and the actions;
- $P(\tilde{o}, \tilde{s}, \pi, \theta)$ is the generative model. It is a probabilistic distribution describing the agent's observations; namely, it is a *model* of the environment. It depends on the set $\theta = \{\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}, \boldsymbol{e}, \beta\}$ of parameters provided by the model designer to shape the generative model; more specifically $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}, \boldsymbol{e}$ have to be fixed to set the Dirichlet

distributions of the variables A, B, C, D, E, and β to define the gamma distribution of the precision γ.

Following the graphical representation in Fig. 1A, a series of marginalisations permit to characterize the generative model as:

$$P(\tilde{o}, \tilde{s}, \pi, \theta) = P(\pi|\gamma, C, E)P(\theta)\prod_{t=1}^{T} P(o_t|s_t, A)P(s_t|s_{t-1}, \pi, B, D)$$

where:

$P(o_t|s_t, A) = Cat(\mathbf{A})$
$P(s_t|s_{t-1}, \pi, B) = Cat\big(\mathbf{B}(u = \pi(t))\big)$
$P(s_1|s_0, \pi, D) = Cat(\mathbf{D})$
$P(\pi|\gamma, C, E) = \sigma(\ln \mathbf{E} - \gamma \cdot \mathbf{G})$
$\mathbf{A} \sim Dir(\boldsymbol{a})$
$\vdots$
$\mathbf{E} \sim Dir(\boldsymbol{e})$

$$P(\gamma) = \Gamma(1, \beta)$$

The matrix $\mathbf{A}$ encodes the likelihood of observations given a hidden state, while $\mathbf{C}$ represents their prior distribution or (preferred) outcomes. State transitions are specified by $\mathbf{B}$, the prior distribution of the initial state is given by $\mathbf{D}$, while the prior distribution over the policies is coded into $\mathbf{E}$. Furthermore, β is the rate parameters of the gamma density that controls the precision γ. Finally, the parameters $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}, \boldsymbol{e}$ are the concentration hyperparameters of the distributions from which the model generates hidden states, outcomes, and policies.

The quantity $\mathbf{G}$ can be considered as a score vector for the policies and can be viewed as a belief over the variable π given future (expected) states and observations, together with the preferred outcomes (more details below in this Section). $\mathbf{G}$ and $\mathbf{E}$ together define the policy posterior as a a Gibbs distribution (a softmax function σ) with γ as inverse temperature – or precision – parameter.

- $Q_{\boldsymbol{\mu}}(\tilde{s}, \pi, \theta)$ is the approximate posterior distribution over states and parameters, with sufficient statistics $\boldsymbol{\mu} = \big(\boldsymbol{s}_{\pi 0}, \dots, \boldsymbol{s}_{\pi T}, \hat{\boldsymbol{\pi}}, \hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{c}}, \hat{\boldsymbol{d}}, \hat{\beta}\big)$. It makes it possible to invert the generative model using variational methods for approximate inference – an operation that is otherwise intractable (Beal, 2003). By using a *mean-field assumption*, the approximate posterior can be factorized over states, action, and precision, as follows:

$$Q_{\mu}(\tilde{s}, \pi, \theta) = Q(s_1|\pi) \dots Q(s_T|\pi)Q(\pi)Q(A) \dots Q(E)Q(\gamma)$$

where:
$Q(s_t|\pi) = Cat(\boldsymbol{s}_{\pi t})$
$Q(\pi) = Cat(\hat{\pi})$
$Q(A) = Dir(\hat{\boldsymbol{a}})$
$\vdots$
$Q(E) = Dir(\hat{\boldsymbol{e}})$
$Q(\gamma) = \Gamma(1, \hat{\beta})$

A key aspect of Active Inference is the free energy minimization with respect to the sufficient statistics $\boldsymbol{\mu}$. By exploiting some mathematical identities, the variational free energy functional can be expressed in terms of the approximate beliefs Q and the generative model P as:

$$\begin{aligned} F &= \mathbb{E}_{Q_\mu(\tilde{s},\pi,\theta)}\left[\ln Q_\mu(\tilde{s},\pi,\theta) - \ln P(\tilde{s},\tilde{o},\pi,\theta)\right] \\ &= D_{KL}\left[Q_\mu(\tilde{s},\pi,\theta)||P(\tilde{s},\pi,\theta|\tilde{o})\right] - \ln P(\tilde{o}) \\ &\geq -\ln P(\tilde{o}) \end{aligned}$$

In these equations, $\mathbb{E}_Q[\cdot]$ denotes an expected value under Q, $D_{\mathrm{KL}}[\cdot\,||\,\cdot]$ is the Kullback-Leibler divergence, that estimates the difference between two distributions, and $-\ln P(\tilde{o})$ (i.e., the negative logarithm of the *model evidence* $P(\tilde{o})$) is called *surprise*. When the variational distribution $Q_\mu(\tilde{s},\pi,\theta)$ tends to get closer to the posterior $P(\tilde{s},\pi,\theta|\tilde{o})$, free energy decreases. If they match entirely, and their divergence is zero, free energy becomes the same as surprise. Therefore, one could summarize variational inference by saying that minimizing free energy entails approximating the posterior using an approximate distribution.

Note that variational Bayesian methods transform inference (namely, calculating posterior from prior beliefs) into an optimization problem: finding sufficient statistics $\boldsymbol{\mu}$ such that the corresponding free energy is minimum. As explained in (K. J. Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2016), is possible to demonstrate that such condition is verified when the following equations are used to define the sufficient statistics at any time τ:

$$\boldsymbol{s}_{\pi\tau} = \sigma(\ln \mathbf{A} \cdot o_\tau + \ln \mathbf{B}_{\pi\tau-1}\, \boldsymbol{s}_{\pi\tau-1} + \ln \mathbf{B}_{\pi\tau} \cdot \boldsymbol{s}_{\pi\tau+1})$$
$$\widehat{\boldsymbol{\pi}} = \sigma(\ln \mathbf{E} - \mathbf{F} - \gamma \cdot \mathbf{G})$$
$$\hat{\beta} = \beta + (\widehat{\boldsymbol{\pi}} - \widehat{\boldsymbol{\pi}}_0) \cdot \mathbf{G}$$

$$\widehat{\boldsymbol{a}} = \boldsymbol{a} + \sum_\tau \boldsymbol{o}_\tau \otimes \boldsymbol{s}_{\pi\tau}$$
$$\widehat{\boldsymbol{b}}_{\pi\tau} = \boldsymbol{b}_{\pi\tau} + \sum_\tau \boldsymbol{s}_{\pi\tau} \otimes \boldsymbol{s}_{\pi\tau-1}$$
$$\hat{\boldsymbol{c}} = \boldsymbol{c} + \sum_\tau \boldsymbol{o}_\tau$$
$$\widehat{\boldsymbol{d}} = \boldsymbol{d} + \boldsymbol{s}_{\pi 1}$$
$$\hat{\boldsymbol{e}} = \boldsymbol{e} + \widehat{\boldsymbol{\pi}}$$

Here, we use the symbols "$\cdot$" and "$\otimes$" to denote respectively the inner and outer products defined as $\mathbf{A} \cdot \mathbf{B} = \mathbf{A}^T\mathbf{B}$ and $\mathbf{A} \otimes \mathbf{B} = \mathbf{A}\mathbf{B}^T$, where $\mathbf{A}$ and $\mathbf{B}$ are two generic matrices. We use the vector $\boldsymbol{o}_\tau$, with all the elements equal to zero except one, as an "one-hot" encoding to represent a specific outcome at the time τ. In the second equation, the letter $\mathbf{F}$ refers to the vector whose elements F_π represent the free energy under the policy π. In the third equation $\widehat{\boldsymbol{\pi}}_0 = \sigma(\ln \mathbf{E} - \gamma \cdot \mathbf{G})$, and $\gamma = 1/\hat{\beta}$.

Solving the updating equations means finding out the posterior expectations that minimise the variational free energy. It is possible split such a list of equations in two different groups, which have distinct functions and timescales. The first three equations support the Bayesian estimation of the hidden states (inference) and are updated at a faster

timescale, after each observation. The other equations support parameter learning by evidence accumulation and are only updated at a slower timescale, at the end of each trial.

The first of the three inference equations estimates the expected hidden state and corresponds to that part in Active Inference related to perception, whilst the second equation derives from a Boltzmann distribution of the policies' quality values. The expected value of γ is the sensitivity (or inverse temperature parameter) of the distribution: it adjusts the tendency to select a policy with greater or lesser stochasticity. The third equation tunes the value of the expected precision on the base of the values of the policy quality.

The other, learning equations (from the fourth to the last) provide a simple scheme to update the hyperpriors of the hidden variables A, B, C, D, and E of the generative model by substituting the parameters (pseudocounts) of their prior Dirichlet distribution with their posteriors, computed over the whole trial. This mechanism resembles Hebbian plasticity: when an element occurs repeatedly, its probability increases, reflecting an implicit expectation that it will likely happen again in the future.

Finally, it is worth paying attention to $\mathbf{G}$: it is a vector whose elements are the "expected free energies" G_π with respect to the expected future outcomes and states, and they score the quality of the policies. One can formally calculate G_π by integrating the (negative) free energy expected under the policy π, from the current instant t to the final one T:

$$\mathbf{G}_\pi = \sum_{\tau=t}^{T} \mathrm{G}(\pi, \tau)$$

with:

$$\begin{aligned}
\mathrm{G}(\pi,\tau) &= F_\tau(\pi) \\
&= \mathbb{E}_{\tilde{Q}}[\ln Q(s_\tau|\pi) - \ln P(o_\tau, s_\tau|\pi, C)] \\
&= \mathbb{E}_{\tilde{Q}}[\ln \ln Q(s_\tau|\pi) - \ln P(s_\tau|o_\tau, \pi) - \ln P(o_\tau|C)] \\
&\geq \mathbb{E}_{\tilde{Q}}[\ln Q(s_\tau|\pi) - \ln Q(s_\tau|o_\tau, \pi)] - \mathbb{E}_{\tilde{Q}}[\ln P(o_\tau|C)] \\
&= \mathbb{E}_{\tilde{Q}}[\ln Q(o_\tau|\pi) - \ln Q(o_\tau|s_\tau, \pi)] - \mathbb{E}_{\tilde{Q}}[\ln P(o_\tau|C)] \\
&= -D_{\mathrm{KL}}[Q(o_\tau|\pi)||P(o_\tau)] - \mathbb{E}_{\tilde{Q}}\big[H[P(o_\tau|s_\tau)]\big]
\end{aligned}$$

where $\mathbb{E}_{\tilde{Q}}[\cdot]$ is the expected value under the predicted posterior distribution defined as $\tilde{Q} = Q(o_\tau, s_\tau|\pi) \triangleq P(o_\tau|s_\tau)Q(s_\tau|\pi)$ over hidden states and their outcomes under a particular policy and at a certain instant of time. The second identity, carried out by factorizing $P(o_\tau, s_\tau|\pi)$ can be bounded down (third row) by using the variational distribution $Q(s_\tau|o_\tau, \pi)$ to approximate the real posterior distribution $P(s_\tau|o_\tau, \pi)$ over the hidden states. The fourth identity easily derives from Bayes' rule. The final identity is the same as shown at the beginning of this Appendix: it provides an interpretation of the expected free energy as the sum of a term representing the risk of using $Q(o_\tau|\pi)$ in the place of $P(o_\tau|C)$, and another term representing the expected uncertainty due to the complexity of the generative model.

A more effective form of G_π is obtaining by considering the equations of sufficient statistics mentioned above, together with the generative model:

$$\mathrm{G}(\pi,\tau) = \boldsymbol{o}_{\pi\tau} \cdot (\ln \boldsymbol{o}_{\pi\tau} - \ln \mathbf{C}) + \boldsymbol{s}_{\pi\tau} \cdot \mathbf{H}$$

with

$$\begin{aligned}
\boldsymbol{o}_{\pi\tau} &= \mathbb{E}_Q[A] \cdot \boldsymbol{s}_{\pi\tau} \\
\ln \mathbf{C} &= \ln P(o_\tau|C)
\end{aligned}$$

$$\mathbf{H} = -diag\big(\mathbb{E}_Q[A] \cdot \mathbb{E}_Q[\ln A]\big)$$

where $\mathbb{E}_Q[A] = \boldsymbol{a}\times\boldsymbol{a}_0^{-1}$, such that $\mathbf{a}_{0ij} = \sum_i \boldsymbol{a}_{ij}$. In this form of the expected free energy, the vectors $\boldsymbol{o}_{\pi\tau}$, $\ln \mathbf{C}$, and the diagonal matrix $\mathbf{H}$ represent the predicted and the preferred future outcomes and their expected entropy.

The determination of $\mathbf{G}$ is key for the agent's action selection process. Mathematically, action selection can be expressed as the minimization of prediction errors about the outcomes:

$$\begin{aligned} u_t &= \min_u \mathbb{E}_Q\big[D_{KL}[P(o_{t+1}\,|s_{t+1})||R(o_{t+1}\,|s_t, u)]\big] \\ &= \min_u \boldsymbol{o}_{t+1} \cdot \varepsilon_{t+1}^u \end{aligned}$$

where:

$\varepsilon_{t+1}^u = \ln \boldsymbol{o}_{t+1} - \ln \boldsymbol{o}_{t+1}^u$
$\boldsymbol{o}_{t+1} = \mathbf{A}\boldsymbol{s}_{\mathbf{t+1}}$
$\boldsymbol{o}_{t+1}^u = \mathbf{AB}(u)\boldsymbol{s}_t$
$\boldsymbol{s}_t = \sum_\pi \boldsymbol{\pi} \cdot \boldsymbol{s}_{\boldsymbol{\pi t}}$

Thus, in Active Inference, actions are selected to minimize prediction errors about expected outcomes. In our simulations, we used the variant of this process adopted in (K. Friston et al., 2015), where the action is sampled as posterior of the sufficient statistics $\widehat{\boldsymbol{\pi}}$:

$$p(u_t|\widehat{\boldsymbol{\pi}}) = \sum_{\{\hat{\pi}:u_t=\hat{\pi}(t)\}} p(\hat{\pi}|u_t)$$

where we assumed that u_t was uniformly distributed and that $\widehat{\boldsymbol{\pi}} = (\hat{\pi}_1, \dots, \hat{\pi}_{\mathrm{m}}) = \big((u_{11}, \dots, u_{1T}), \dots, (u_{\mathrm{m}1}, \dots, u_{\mathrm{m}T})\big)$.

	First simulation	Second simulation	Third simulation	Fourth simulation	Fifth simulation
Prior for panic context	0.5	From 0.9 to 0.1	From 0.9 to 0.1	0.5	0.5
Values of A matrix	0.9 heart pound 0.5 breathless	0.9 heart pound 0.5 breathless	0.7 heart pound 0.5 breathless	0.9 heart pound 0.5 breathless	0.9 heart pound
Values of B matrix (transitions between contexts)	0.9	0.9	0.9	0.6	0.9
C matrix	2 feel normal -2 feel strange -0.1 heart pound -0.1 breathless	2 feel normal -2 feel strange -0.1 heart pound -0.1 breathless	2 feel normal -2 feel strange -0.1 heart pound -0.1 breathless	2 feel normal -2 feel strange -0.1 heart pound -0.1 breathless	2 feel normal -2 feel strange -0.1 heart pound -0.1 breathless
Learning rate	-	-	-	-	0.0025
Percentage of pain in the first and second blocks	-	-	-	-	0.9 / 0 for the first agent; 0.4 / 0 for the second agent

Table A1. Parameters used in the simulations