

RESEARCH ARTICLE

In the Body's Eye: The computational anatomy of interoceptive inference

Micah Allen^{1,2*}, Andrew Levy³, Thomas Parr³, Karl J. Friston³**1** Centre of Functionally Integrative Neuroscience, Aarhus University Hospital, Denmark, **2** Cambridge Psychiatry, Cambridge University, Cambridge, United Kingdom, **3** Wellcome Centre for Human Neuroimaging, University College London, London, United Kingdom

* micah@cfin.au.dk



Abstract

A growing body of evidence highlights the intricate linkage of exteroceptive perception to the rhythmic activity of the visceral body. In parallel, interoceptive inference theories of affective perception and self-consciousness are on the rise in cognitive science. However, thus far no formal theory has emerged to integrate these twin domains; instead, most extant work is conceptual in nature. Here, we introduce a formal model of cardiac active inference, which explains how ascending cardiac signals entrain exteroceptive sensory perception and uncertainty. Through simulated psychophysics, we reproduce the defensive startle reflex and commonly reported effects linking the cardiac cycle to affective behaviour. We further show that simulated 'interoceptive lesions' blunt affective expectations, induce psychosomatic hallucinations, and exacerbate biases in perceptual uncertainty. Through synthetic heart-rate variability analyses, we illustrate how the balance of arousal-priors and visceral prediction errors produces idiosyncratic patterns of physiological reactivity. Our model thus offers a roadmap for computationally phenotyping disordered brain-body interaction.

OPEN ACCESS

Citation: Allen M, Levy A, Parr T, Friston KJ (2022) In the Body's Eye: The computational anatomy of interoceptive inference. *PLoS Comput Biol* 18(9): e1010490. <https://doi.org/10.1371/journal.pcbi.1010490>

Editor: Sahib S. Khalsa, Laureate Institute for Brain Research, UNITED STATES

Received: February 8, 2021

Accepted: August 13, 2022

Published: September 13, 2022

Copyright: © 2022 Allen et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The underlying MDP scheme here is available as part of the open-access distribution of SPM12 (<https://www.fil.ion.ucl.ac.uk/spm/software/spm12/>). A demonstration of the scheme can be accessed by typing >>DEM into the Matlab command prompt and selecting the Interoception demo from the graphical user interface that appears. Further, all the code required to generate the simulations and figures herein can be found at the following github page: <https://github.com/embodied-computation-group/cardiac-active-inference>. Full installation instructions can be found in the following code

Author summary

Understanding interactions between the brain and the body has become a topic of increased interest in computational neuroscience and psychiatry. A particular question here concerns how visceral, homeostatic rhythms such as the heart beat influence sensory, affective, and cognitive processing. To better understand these and other oscillatory brain-body interactions, we here introduce a novel computational model of interoceptive inference in which a synthetic agent's perceptual beliefs are coupled to the rhythm of the heart. Our model both helps to explain emerging empirical data indicating that perceptual inference depends upon beat-to-beat heart rhythms, and can be used to better quantify intra- and inter-individual differences in heart-brain coupling. Using proof-of-principle simulations, we demonstrate how future empirical works could utilize our model to better understand and stratify disorders of interoception and brain-body interaction.

wrapper, which reproduces all figures and simulations: https://github.com/embodyed-computation-group/cardiac-active-inference/blob/master/paper_wrapper.m.

Funding: MA is supported by a Lundbeckfonden Fellowship (R272-2017-4345). TP is supported by the Rosetrees Trust (Award Number 173346). KJF is a Wellcome Principal Research Fellow (Ref: 088130/Z/09/Z). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Introduction

The enactive view of perception—implied by active vision and inference—suggests an intimate co-dependency between perception and the active sampling of our sensorium. In this work, we take the embodied view to its ultimate conclusion and consider perception as a function of the physical and physiological body we use to ‘measure’ the world. In particular, our focus is on the coupling—or interaction—between interoceptive and exteroceptive perception; namely, how bodily states and states of affairs beyond the body are inferred—and how inference about each domain affects the other. For example, does what we see depend upon our autonomic status and how does visual perceptual synthesis affect sympathetic or parasympathetic outflow? Similarly, if precisions are tied to bodily rhythms, then the metacognitive ability to monitor and integrate these signals is also likely closely related to interoceptive inference. The body is, in essence, an ensemble of fluctuating systems with biorhythms nested at multiple timescales. How then do these physiological fluctuations interact with perceptual and metacognitive inferences in the exteroceptive domains?

There is a rapidly growing body of evidence suggesting that bodily and autonomic states affect perceptual and metacognitive decisions [1–11]. Much of this evidence emphasizes the dynamic aspect of our physiology; usually assessed in terms of how perceptual choice or subjective confidence depends upon the phase of some physiological cycle. Most of the empirical evidence suggests that biorhythms gate or modulate the way that sensory evidence is accumulated during perception [3,12–14]. In the predictive coding literature, this is usually treated as fluctuating, context sensitive, changes in the precision of sensory sampling (e.g., the precision or gain of prediction errors). Clear examples of this include the fast waxing and waning of precision during active visual sampling. For example, saccadic suppression—during saccadic eye movements—alternates with attention to fixated visual information every 250 ms or so. This process of actively sampling the environment via ballistic saccade itself varies with the cardiac cycle [15–17]. At still slower timescales, respiratory [11,18–21] and gastric cycles [22–24] are coupled to neuronal oscillations and behaviour. Given the ample evidence for a close linkage of these rhythms of the body with those of the brain and sensory epithelium, we hypothesized that many of these effects may arise from a process of *interoceptive self-inference* [25,26], that is, inference about fluctuations in sensory precision induced by these brain-body rhythms.

Our focus on the multimodal integration of interoceptive and exteroceptive domains is driven by the overwhelming evidence for interoception as a key modality in hedonics, arousal, emotion and selfhood [22,27–37]. This is generally treated under the rubric of interoceptive inference; namely, active inference in the interoceptive domain. There are several compelling formulations of interoceptive inference from the perspective of neurophysiology, neuroanatomy and, indeed, issues of consciousness in terms of minimal selfhood. For example, the integration of interoceptive precision has been advanced as a key mechanism underpinning both interoceptive awareness and the stability of the body schema [38,39]. More generally, brain-body rhythms and their association with interoception has been proposed as a key influence on embodied perception [18,22]. This newly emerging computational view of brain-body interaction may also carry important new insights for dissociating unique brain-to-body or body-to-brain phenotypes in computational psychiatry and neurology [40–45].

However, a key limitation thus far is that the majority of extant work interlinking predictive processing and interoceptive inference has been of a conceptual nature, underpinned by some notion of active (Bayesian) inference about states of the world (including the body). More recently, active inference models have been utilized to develop formal models of interoceptive awareness in the gastric and cardiac domain [46,47], as well as emotional and affective inference [36,43,48]. In this work, we offer a formal mathematical analysis that casts brain-body

interactions in terms of temporal coupling between interoceptive and exteroceptive inference, in the hope of providing a point of reference for both theoretical and empirical investigations.

In brief, we constructed a minimal active inference architecture to simulate embodied perception and concomitant arousal. Here, we focused on simulating interactions between the cardiac cycle and exteroceptive perception. In principle however, our simulation provides a computational proof-of-principle that can be expanded to understand brain-body coupling at any physiological or behavioral timescale. Using a Markov decision process formulation, we created a synthetic subject who exhibited physiological (cardio-acceleration) responses to arousing stimuli. Our agenda was twofold: first, to provide a sufficiency proof that—in at least one example—the interaction between interoception and exteroception emerges from the normative (formal) principles of active inference. Furthermore, having an *in-silico* subject at hand, means that we can simulate the effects of various disconnections and pathophysiology. For example, we can examine the effect of deafferentation of interoceptive signals on arousal, exteroceptive perception, and the posterior uncertainty associated with these signals—a key input for metacognition in the interoceptive domain.

Second, we constructed our synthetic subject in such a way that the same paradigm could be replicated in real subjects. The motivation for this is that the active inference scheme used below has an associated process theory [49]. In other words, neuronal and behavioral responses associated with inferential processes can be simulated on a trial-by-trial basis. This means that we can use electrophysiological, eye tracking, pupillometry and other physiological proxies to test various hypotheses that can be instantiated in the model [50]. Crucially, this provides a link between neuronal and behavioural responses—as characterised by the latency between stimuli onset and autonomic responses (e.g., heart rate acceleration or variability). In this paper, we will focus on the basic phenomenology and results thereof. In subsequent work, we will use this formalism to model real responses under various experimental manipulations.

In what follows, we briefly describe the generative model and inversion scheme used to simulate cardiac arousal responses. We then demonstrate the results of anatomical (deafferentation) lesions on perceptual and metacognitive signals (i.e., the uncertainty or posterior entropy associated with interoceptive and exteroceptive beliefs), as well as simulated belief updating. Finally, we will examine the effects on synthetic heart-rate variability when changing the precision of various prior beliefs that underlie perceptual inference. We conclude with a discussion of the implications for existing research in this area—and how this research could be informed by a formal approach providing guidelines to discovery.

Methods

Markov decision process

The simulations reported below build upon the notion of active inference. This is a ‘first-principles’ approach to understanding (Bayes) optimal behaviour. Simply put, active inference treats the brain as using an internal (generative) model of the world to explain exteroceptive, proprioceptive, and interoceptive sensory data. By optimizing beliefs about variables in this model (perceptual inference), or by changing their internal or external environment (action), creatures can ensure their sensations and prior beliefs are consistent. The term ‘belief’ here is used in the technical sense of a Bayesian belief, or probability distribution, typically considered to be sub-personal. A Markov decision process (MDP) is a form of probabilistic generative model that describes the sequential dynamics of unobserved (hidden) variables (e.g., the current state of the cardiac cycle) and the sensations they cause (e.g., baroreceptor signals). In this paper, we deal with a special case of a Markov decision process known as a partially observed Markov decision process (or POMDP). The difference between the two is that the latter

include unobserved variables, while other forms of MDP assume all information is directly available to the agent employing that model. The hidden variables of an MDP are hidden states (s_t) and sequences of actions or policies (π). Within this model there are four kinds of precision; namely, sensory precision in the visual (α) and cardiac (β) domains and the precision of state transitions among interoceptive and exteroceptive states (ζ). For clarity, we will refer to the precision of transitions as (inverse) volatility and use precision to refer to sensory (i.e., likelihood) precision. The generative model then embodies the conditional dependencies between these variables, as expressed graphically in Fig 1. While we provide a brief overview here, we refer readers to [49,51] for more technical detail.

It is important to be clear about what is meant by precision, which can be a confusing concept when first encountered. Technically, precision is a feature of a probability distribution, and can be thought of in several complementary ways: the negentropy of that distribution, its inverse temperature, or its inverse (co)variance. All of these quantify the uncertainty in a given distribution, where uncertainty grows as precision declines. This means precision is only meaningful when we ask which distribution it is applied to. In this paper, the precisions we deal with primarily relate to the likelihood distributions (i.e., the probability of sensory data given their latent or hidden causes). In this context, precision becomes the quality or reliability of sensory data (c.f., that can be read as a sensory signal-to-noise ratio).

Hidden states generate observable sensory data with probabilities expressed in a likelihood matrix **A**. The states evolve through time according to a transition probability matrix, **B** and depend only on the state at the previous time, and on the policy, π . Finally, we equip the generative model with preferences (**C**), prior beliefs about initial states (**D**), and prior beliefs about policies. Beliefs about policies have two parts. The first of these is a fixed bias (**E**). This may be thought of as a habit; i.e., ‘what I expect to do’ *a priori*. The second is a belief that the most probable policies are those that have the lowest expected free energy (**G**); i.e., ‘what I expect to do’ after considering the consequences of action. A simple intuition for the latter is to think of the selection between alternative courses of action as we might think of Bayesian hypothesis testing (i.e. model comparison); namely, planning as inference [52,53]. Here, each policy can be thought of as an alternative hypothesis about ‘how I am going to behave’. These are evaluated in terms of prior beliefs (**E**), and the (predicted) evidence future data affords (**G**). Just as free energy is used to approximate the evidence data affords a hypothesis, expected free energy evaluates the expected evidence, under beliefs about how data are actively generated. As expressed in Fig 1, expected free energy can be separated into two parts. ‘Risk’ quantifies how far predicted observations deviate from preferred outcomes. Minimizing this ensures maintenance of homeostasis. ‘Ambiguity’ quantifies the uncertainty in the mapping from states to outcomes. Minimizing this component ensures that salient, uncertainty-resolving data are sought (leading to epistemic, information gathering, behavior).

Synthetic cardiac arousal

Using the MDP scheme detailed above, we set out to simulate a cardiac arousal response to threatening stimuli (e.g., a vicious looking spider), in comparison to non-arousing stimuli (e.g., some flowers). To do this, we had to define ‘arousal’ and its interoceptive correlates. To keep things as simple as possible, we assumed the subject’s generative model included two sorts of hidden states (*interoceptive* and *exteroceptive*—and that she could adopt two modes of engagement with the world (*relaxed* and *aroused*). These sorts of generative models are generally cast as Markov decision processes, whereby transitions among (hidden) states generate observable outcomes in one or more modalities. The modalities considered here were *exteroceptive* (*non-arousing* versus *arousing* visual stimuli) and interoceptive (the cardiac phase;

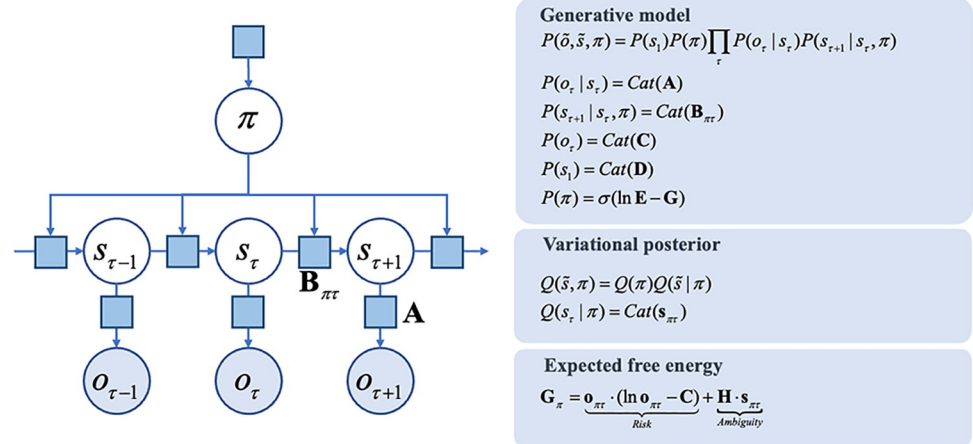


Fig 1. A Markov decision process generative model: the factor graph on the left illustrates the conditional dependencies, and independencies, between the variables in the generative model (see the main text for a description of the variables). The variables are shown in circles (with filled circles showing observable variables). An arrow from one variable to another indicates that the latter depends upon the former. The square nodes each represent probability distributions. The panels on the right give the forms of the distributions (associated with each square node) in the generative model, in addition to defining the expected free energy, and specifying the factorization of the approximate posterior (variational) distributions the agent possesses. In this figure, the tilde symbol (~) means a sequence over time. The σ is a softmax (normalized exponential) distribution. H is a vector of the likelihood entropy (where each element of the vector corresponds to an alternative state). The bold o is a predicted observation, while the italicized o represents realized observations. Generally, bold refers to the parameters of the generative model, while italicized symbols represent the random variables described by the model.

<https://doi.org/10.1371/journal.pcbi.1010490.g001>

diastolic or *systolic*). Having defined the nature of the state space generating outcomes, this model can then be parameterised in a relatively straightforward fashion as outlined above. For any set of **A, B, C, D,** and **E** parameters, one can then simulate active inference using standard marginal message passing schemes [54] to optimize expectations about hidden states of the world—and the action or policy currently being pursued (technically, a policy is a sequence of actions. In what follows, we only consider policies with one action) [49,55].

Crucially, inference about policies rests upon prior beliefs that the policies will minimise expected free energy in the future. This expected free energy has both epistemic and instrumental terms; namely; the ability of any particular course of action to resolve uncertainty about hidden states (known as salience, Bayesian surprise, information gain, etc.) [56–59] and the pragmatic affordance (known as expected value, utility, reward, etc.) as specified by the prior preferences [60].

To capture the fundamentals of multimodal integration of interoceptive and exteroceptive modalities, we assumed the following, reasonably plausible, form for the model. The synthetic subject had to infer which of two policies she was pursuing: a *relaxed* policy or an *aroused* policy. These are defined operationally in terms of transitions among interoceptive states. Here, we model this in terms of two distinct forms of cardiac cycling among *diastolic* and *systolic* bodily states. When *relaxed*, the probability transitions among cardiac states meant that there were two phases of *diastole* and one of *systole*. Conversely, when *aroused*, the first *diastolic* state jumped immediately to *systole*. In brief, this means that being aroused causes cardiac acceleration and increases the average amount of time spent in systole. The outcomes generated by these states were isomorphic; in other words, there was a simple likelihood mapping from states to sensations; such that the subject received a precise or imprecise interoceptive cue about the current cardiac status (i.e., *diastole* or *systole*).

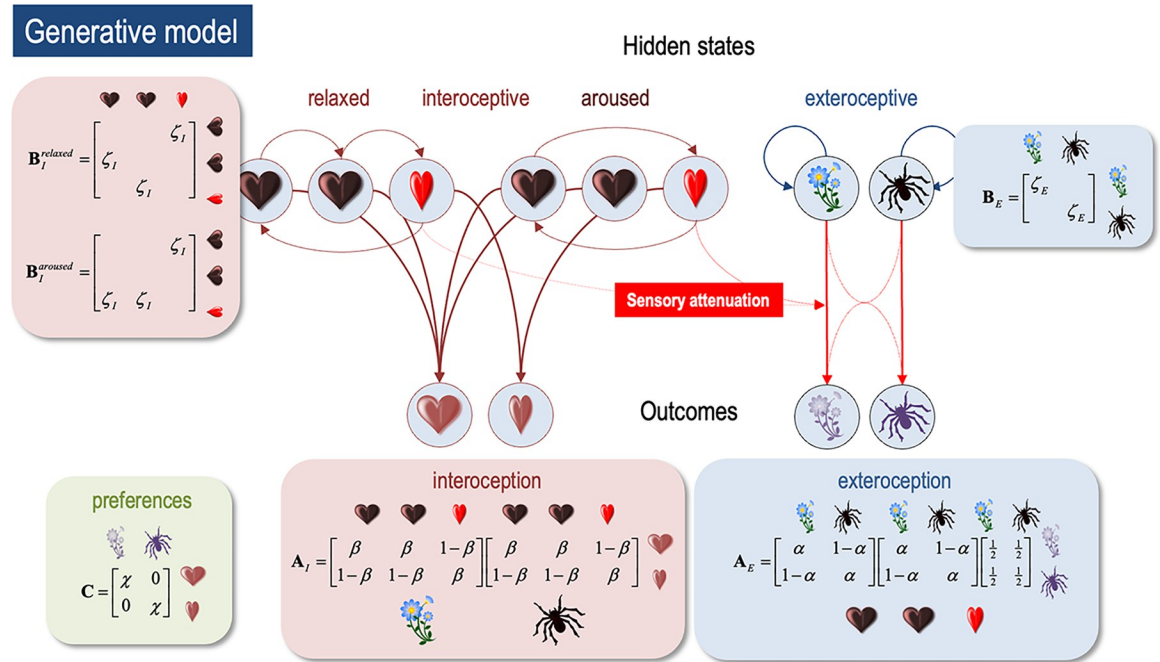


Fig 2. The generative model. This schematic illustrates how hidden states cause each other and sensory outcomes in the interoceptive and exteroceptive domain. The upper row describes the probability transitions among hidden states, while the lower row specifies the outcomes that would be generated by combinations of hidden states that are inferred on the basis of outcomes. The green panel specifies the model’s prior preferences; namely, the sorts of outcomes it expects to encounter. Please see main text for a full explanation. Although this figure portrays interoceptive and exteroceptive outcomes as separate modalities, they were in fact modelled as combinations—so that the prior preferences could be evaluated (this is necessary because the preferred physiological outcome depends upon the visual cue). In this model, the precisions are denoted by Greek letters and control the fidelity of various probabilistic mapping is (i.e., the likelihood or **A** matrices and the transition or **B** matrices).

<https://doi.org/10.1371/journal.pcbi.1010490.g002>

Here, we modelled interoceptive and exteroceptive modalities using a joint factorial approach. Probably the simplest way to understand the joint modelling of exteroceptive and interoceptive modalities is with reference to Fig 2. The important message from this figure is that the same hidden states generate outcomes in both modalities. In practice, this means that combinations of exteroceptive and interoceptive observations are generated by the hidden states (and therefore combinations of multimodal sensations inform inferences about hidden states). The joint modelling is reinforced by the specification of preferences over all combinations of observations.

On the exteroceptive side, we considered two states of the visual world; namely, the subject was confronting an arousing or non-arousing visual object. The corresponding visual modality again had two levels (arousing versus non-arousing picture). Crucially, the fidelity or precision of this mapping depended upon the interoceptive state. When the subject was in systole, this mapping became very imprecise. In other words, all outcomes were equally plausible under each hidden state of the visual world. Conversely, during diastole, there was a relatively precise likelihood mapping. This is the crucial part of our model that links the state of the body to the way that it samples the world. Put simply, precise visual information is only available during certain parts of the cardiac cycle, which itself depends upon the state of arousal (i.e., the policy currently inferred and selected). This can be thought of as a simple approximation of cardiac and other bodily timing effects, expressed as a momentary occlusion or attenuation of sensory input by (for example) afferent inhibitory baroreceptor effects [3,7,8,12,17,32,61–63], or by the

brief flooding of the retina during cardiac contraction. We note however that this linkage of interoceptive and exteroceptive precision is but one operational hypothesis regarding how frequently observed effects on cardiac, respiratory, and gastric cycles could be defined—and should accordingly be treated more as a proof of principle than a definitive empirical model of brain-body coupling. Nevertheless, this minimal model evinces some of the curious and familiar results in empirical studies of visceral-perceptual coupling.

This simple structure produced some remarkable results that speak to the intimate relationship between interoception and exteroception. These phenomena (see below) rest upon the final set of beliefs; namely, preferred outcomes. Here, the subject believed that she would be, on average, in a *systolic* state when confronted with an arousing picture and in a *diastolic* state otherwise. These minimal prior preferences then present the subject with an interesting problem. She has to choose between extending periods of precise evidence accumulation (i.e., a *relaxed* state with more *diastolic* episodes) and sacrificing precise information, via cardio-acceleration, should she infer there is something arousing ‘out there’. However, to infer what is ‘out there’, she has to resolve her uncertainty, through epistemic foraging; i.e., maintaining a *relaxed* state. More precisely, the generative model ensures a preference for a combination of ‘systole’ and the ‘arousing picture’ outcomes and a combination of ‘diastole’ and the ‘non-arousing picture’ outcome. We therefore hypothesized that at the beginning of each trial or exposure to a picture, subjects would be preferentially in a relaxed state until they had accumulated sufficient evidence to confidently infer the visual object was *arousing* or not. If *arousing*, she would then infer herself to be aroused and enter into a period of cardio-acceleration (illustrated in Fig 3). Every simulation started off with a weak prior over hidden states that the picture was not arousing—and a weaker prior in favor of the *relaxed* policy.

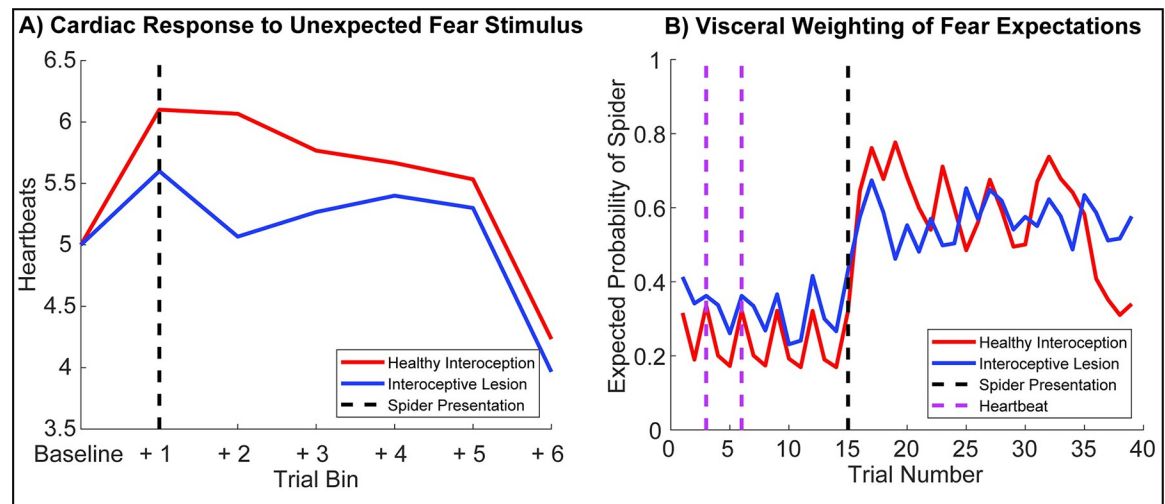


Fig 3. Simulated Physiology and Perceptual Inference. To establish the face validity of our model, we first set out to reproduce some basic psychophysiological phenomenology and establish how these phenomena change under ‘healthy’ (i.e., normative) versus ‘visceral lesion’ parameter settings. To do so, we fed agents a fixed sequence of cardiac and exteroceptive stimuli, such that the first 14 trials constituted a ‘baseline’ period of cardiac quiescence (i.e., a steady heart rate), in the absence of arousing stimuli. On the 15th trial, an unexpected arousing stimulus (a ‘spider’) is presented and a further 85 trials simulated. This simulation was repeated for 60 simulated participants, each with randomized starting values, half of which had ‘lesioned’ interoceptive precision ($\beta = 0.5$, blue lines). Under these conditions our synthetic subjects exhibit a clear ‘startle’ or ‘defense’ reflex (Graham and Clifton, 1966; Sokolov, 1963), characterized by an immediate cardio-acceleration (left panel) and a dramatic shift in the posterior expectation of encountering another threatening stimulus. Interestingly, during the baseline period the posterior expectation of encountering a threat stimulus oscillates with the heartbeat; i.e., the lesioned subjects show both an attenuation of the cardiac response and a blunted belief update. Note that for the right panel, only trials 1–40 are shown. On the left, blue lines show summed heartbeats (time spent in systole) for 15-trial bins; on the right, lines depict the median posterior probability that the agent will see a spider on the next trial. See *Methods* and *Results* for more details.

<https://doi.org/10.1371/journal.pcbi.1010490.g003>

By carefully adjusting the precision of sensory evidence (through adjusting the \mathbf{A} matrix), we could trade-off the evidence accumulation against these imperatives to simulate the elaboration of an arousing response to, and only to, *arousing* stimuli. Furthermore, we anticipated that a failure to implement a selected policy of arousal would both confound inference about the policy being pursued (i.e., an aroused state of mind) and—importantly—confidence about the exteroceptive state of affairs. The latter can be measured quantitatively in terms of the entropy or average uncertainty over hidden exteroceptive states (after taking a Bayesian model average over policies). This leads to the prediction that confidence in perceptual categorisation would not only evolve over time but would depend upon interoceptive inference, with implications for metacognitive inferences at higher levels of the hierarchy not here implemented. We tested this hypothesis *in silico* through various lesion experiments reported in the subsequent sections (Figs 3–5). In what follows, we illustrate the belief updating and arousal responses under ‘normal’ priors (i.e. precisions) based upon the generative model above (summarized graphically in Fig 2).

For the purposes of this paper, the generative model and process are assumed to be identical i.e., that the generative model is a good model of its environment, in the spirit of the good regulator theorem [64]. Conceptually, this assumes an important connection between brain, body, and external environment: namely that a particular kind of body and a particular kind of environment implies a particular kind of brain. Given that the brain regulates the body and environment, this implication is reciprocal. An alternative approach would have been to use a fixed generative process, and to assume alternative models that deviate from this process. This may be more interesting in the context of autonomic pathologies, in which the brain fails to regulate its environment in the way it anticipates.

Simulations

We implemented a minimal model of interoceptive inference: in the sense that one’s state of active engagement with the world may be inferred from its interoceptive and exteroceptive consequences. In this minimal model, the two domains of perception are coupled by—and only by—sensory attenuation: i.e., attenuation of sensory precision in the visual domain during (inferred) systole. Precision refers to the reliability or confidence ascribed to a given probabilistic belief. Within this model there are four kinds of precision; namely, sensory precision in the visual (α) and cardiac (β) domains and the precision of state transitions among interoceptive and exteroceptive states (ζ). For clarity, we will refer to the precision of transitions as (inverse) volatility and use precision to refer to sensory (i.e., likelihood) precision. To see how these precisions manifest in the equations of the generative model, please see Fig 2. In this example, because there are only two states, the corresponding parameters of the generative model control both the expected contingencies and their precision. In other words, when α (or β) decreases to 1/2, sensory signals become imprecise and completely ambiguous. In what follows, we will focus on manipulations of precision under a canonical volatility of $\zeta = 0.9$. In other words, we will assume that our synthetic subject believes state transitions among phases of the cardiac cycle follow each other fairly reliably with a 90% probability. Similarly, if there is a flower ‘out there’, then there is a 90% probability that it will remain there at the next sample. Cardiac and visual stimuli were generated by the same precisions and volatilities as assumed by the subject’s generative model.

We conducted three sets of simulations to illustrate the sorts of behaviours that emerge under this active inference scheme—and to establish the construct validity of the model in relation to empirical phenomena that speak to the influence of interoception on exteroception and *vice versa*. This enabled us to illustrate the basic phenomenology of our agent—in terms of

simulated perceptual inference and cardiac physiology—under some differing levels of sensory precision.

In the first set of simulations (Fig 3), we focused on the physiological and psychological response to arousing stimuli. To do so, we tested the hypothesis that the unexpected presentation of a ‘spider’ would induce an aroused state—as reflected in an increased heart rate—and a greater posterior expectation of encountering an arousing spider stimulus on the next trial. To evaluate this hypothesis, we supplied the subject with a fixed sequence of 15 stimuli—in both the cardiac and visual domains—and examined the posterior beliefs about the next exteroceptive state following a period of relaxed cardiac input. Note that this is possible precisely because our generative model includes beliefs about the future—including the next hidden state and subsequent sensory sample. Here, we used as outcome measures the agent’s evoked cardiac acceleration response (calculated by binning the number of systole events across the experiment) and the agent’s posterior belief that the next stimulus would be threatening. These simulations were repeated 60 times with randomized starting values, such that the first thirty ‘healthy’ agents were compared to an ‘interoceptive lesion’ group for whom interoceptive precision had been attenuated ($\beta = 0.5$). This enabled us to not only establish the interaction of arousal-based expectations and cardiac arousal, but also to demonstrate how these responses change when interoceptive sensory precision is ablated.

In the second set of simulations (Fig 4), our focus moved from perceptual to metacognitive signals. Here, we examined the interaction between exteroceptive and interoceptive sensory precision on the one hand and their coupling to cardiac timing and metacognitive uncertainty (posterior confidence) on the other. Our goal here was to illustrate how both interoceptive and exteroceptive precision interact to influence the uncertainty signals that form the basis of metacognitive inference, and to link these to empirical findings showing that cardiac arousal biases confidence [1,65]. We note however that our model does not constitute a full generative or hierarchical model of metacognitive inference itself: rather, we here examine the influence of first-order interactions in interoceptive and exteroceptive precision on uncertainty, which implies a difference of input to a putative, higher-order metacognitive system [66–70]. Metacognitive, in this context, is used in a technical sense to refer to the updating of beliefs about beliefs. In this setting, this translates into predictions of predictability or precision; where precision is itself an attribute of a (Bayesian) belief about the causes of sensations. There are arguments that this kind of metacognitive belief updating underwrites sensory attenuation, selective attention and perhaps sentience per se [71–73].

For these simulations, we used the uncertainty about inferred exteroceptive and interoceptive states (as quantified by the summed entropy of posterior beliefs for each state) as outcome measures, simulated under a range of cardiac and visual precision settings (Fig 4A). To further illustrate how these effects oscillate with the cardiac rhythm, we separated these measures for each phase of the cardiac cycle (early diastole, late diastole, systole). We then repeated these analyses comparing ‘healthy’ interoceptive inference agents (α & $\beta = 0.9$), to agents for whom either exteroceptive or interoceptive precision was lesioned (α or $\beta = 0.5$, respectively). Note that there is nothing special about these numerical values, which were chosen simply to demonstrate substantive changes in behaviour. Their particular values become more meaningful when active inference models are fit to empirical choice behaviour: e.g., [50]. In virtue of our coupling of exteroceptive sensory precision to the cardiac cycle, we anticipated that metacognitive uncertainty (scored by the negative entropy of posterior beliefs) would depend on the precision of both interoceptive and exteroceptive states, and that this effect would clearly oscillate with the cardiac cycle. Further, we expected in the extreme case of our ‘lesioned’ subjects, these effects would be further exacerbated such that interoceptive and exteroceptive uncertainty would increase dramatically, under their respective lesion conditions.

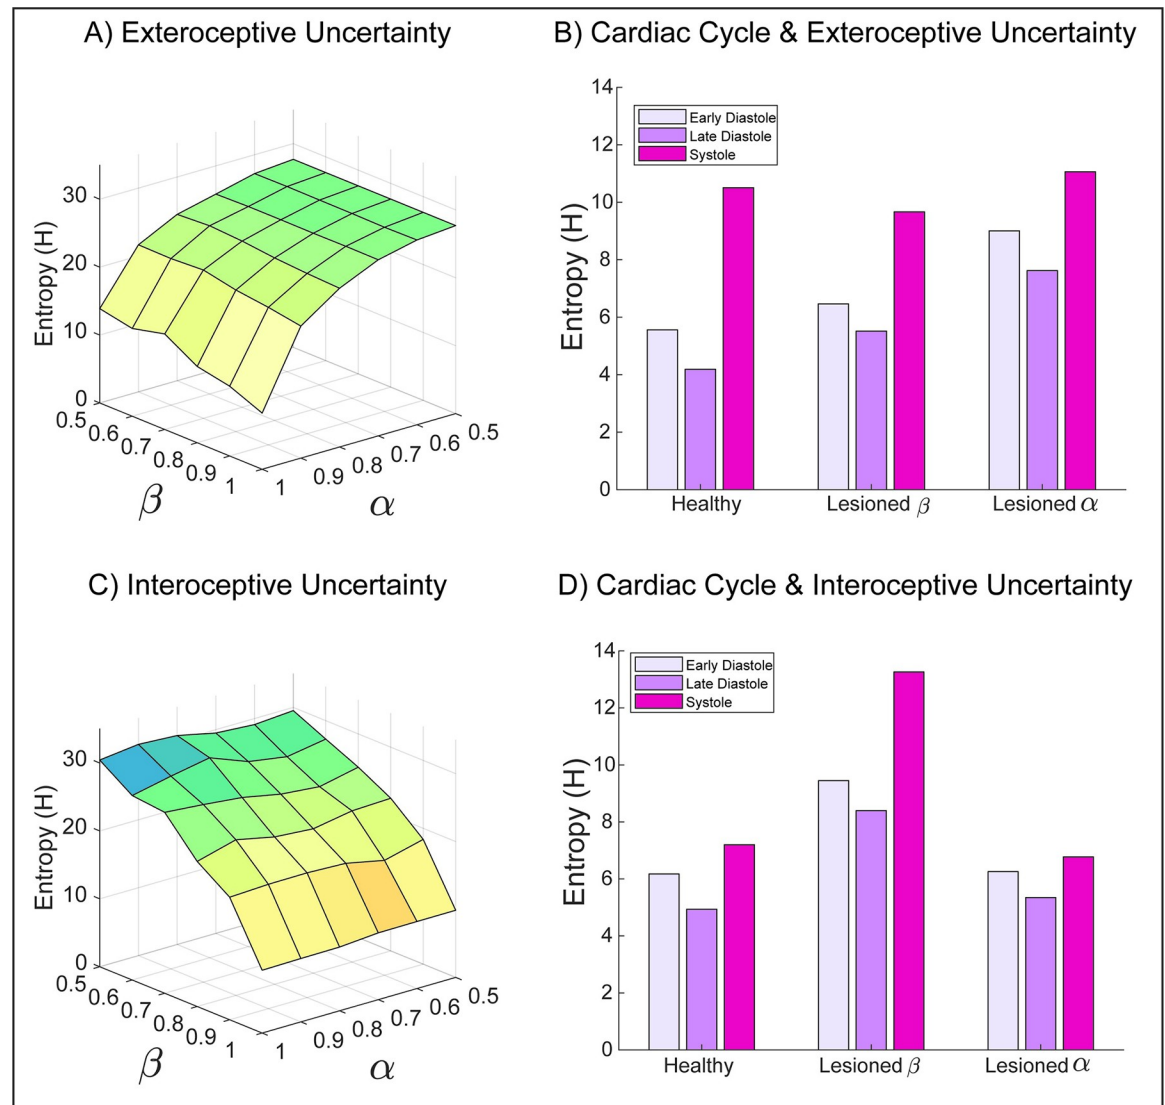


Fig 4. Simulating the influence of interoceptive and exteroceptive precision on uncertainty. To explore how interoceptive inference may shape metacognition, i.e., inference over inferred sensory uncertainty, we measured the summed entropy of beliefs for both exteroceptive (top panels) and interoceptive (bottom panels) states. By simulating the full range of sensory precision values, from lesioned precision (α or $\beta = 0.5$) to extremely high precision (α or $\beta = 1$), the predominant pattern of interactions is revealed. **A)** For exteroceptive inferences (i.e., the agent's belief that a spider or flower is present), the principle entropy gradient is characterized by reductions in exteroceptive precision. This effect is modulated in part by interoceptive precision; for example, the lowest uncertainty is obtained when interoceptive and exteroceptive precision are maximal. **B)** Separating exteroceptive uncertainty by each phase of the cardiac cycle reveals a clear effect of the heartbeat on belief entropy, which is modulated most strongly by lesioning the precision of exteroceptive predictions. Lesioning interoceptive uncertainty does raise the overall level of exteroceptive uncertainty, but to a lesser degree. Note that altering exteroceptive precision only affects the diastolic phases (as precision is already attenuated during systole). Interoceptive lesions preclude precise inferences about the cardiac phase, so reduce the discrepancy in uncertainty between these phases. **C)** Similar to exteroceptive belief, interoceptive metacognition is predominately influenced by interoceptive precision. **D)** The cardiac cycle also modulates the overall uncertainty of interoceptive beliefs; this effect is greatly increased when interoceptive precision is lesioned. Interestingly, exteroceptive lesions primarily reduce the differentiation between cardiac states.

<https://doi.org/10.1371/journal.pcbi.1010490.g004>

Finally, to complement these simulations we modelled the response of first and second order statistics of the physiological responses to changes in sensory precision. These were based upon simulated heart rate (frequency of systole) and the heart rate variability (HRV) assessed over multiple trials or heartbeats (Fig 5). Our objectives here were; 1) to test the

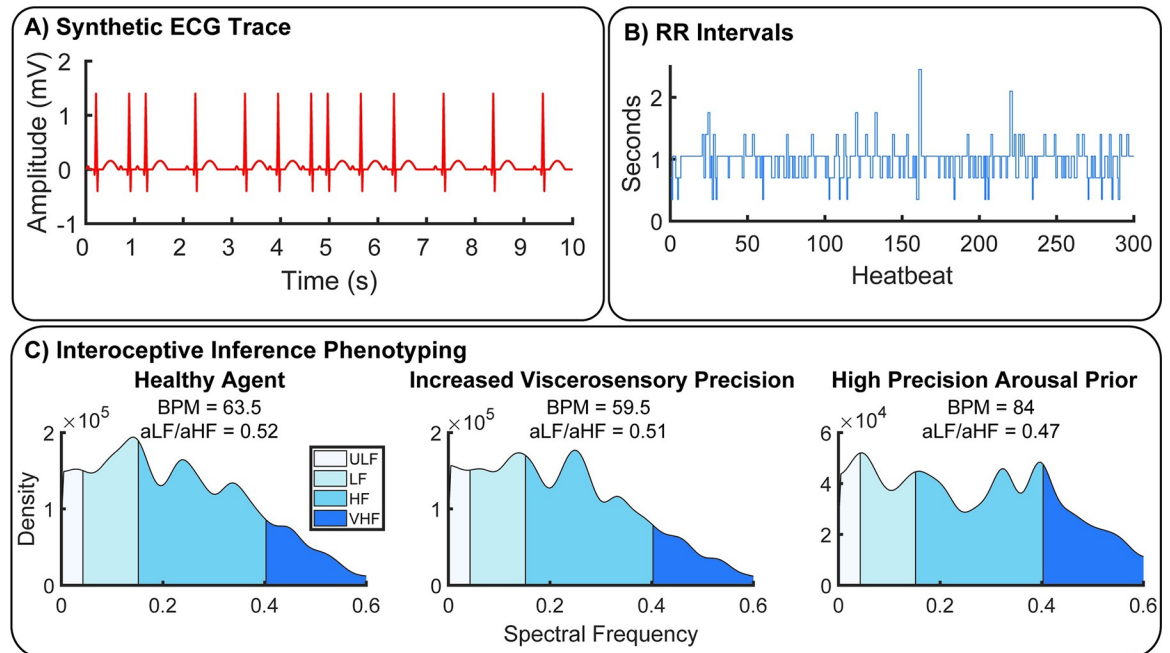


Fig 5. Synthetic Heart-Rate Variability (HRV) and Interoceptive Computational Phenotyping. To illustrate the potential of our approach as a generative model of physiological reactivity, we produced synthetic heartbeat traces and analyzed these with a standard time-frequency approach under various canonical parameter settings. **A)** Synthetic ECG traces produced by convolving a standard QRS-wave function with systole events generated by our model. **B)** These were then transformed into RR-intervals by assuming an 350ms sampling rate, **C)** Power spectra of RR-intervals were calculated using Welch's method and categorized as ultra-low (ULF), low (LF), high (HF), and very high frequency (VHF) bands for each simulated agent. Physiological responses were then summarized in terms of beats-per-minute (BPM) and sympathovagal balance (ratio of area under curve for each frequency band, aLF/aHF) [74]. To illustrate the potential of our approach for interoceptive computational phenotyping, we simulated three different agents—one with healthy interoceptive inference (bottom left), another with extremely precise visceral sensations (bottom middle), and another with overly precise priors for the aroused (sympathetic) policy (bottom right). These each produce unique interoceptive inference 'fingerprints'; i.e., the individual patterns of heart-rate variability produced by these parameter settings. In this example, extremely precise visceral sensations reduce heart-rate and shift overall peak frequency to the high-frequency domain, whereas overly precise arousal priors induce strong heart-rate acceleration coupled with attenuated ultra-low and ultra-fast oscillations. In the future, these idiosyncratic patterns could be used to identify maladaptive interoceptive inference from heart-rate data.

<https://doi.org/10.1371/journal.pcbi.1010490.g005>

hypothesis that fluctuations in both low-and high- frequency synthetic heart rate variability can be produced by altering the balance of interoceptive sensory precision versus the prior precision for the aroused sympathetic policy, and 2) to illustrate how generative modelling of interoceptive active inference can be used to phenotype maladaptive inference parameters from observed heart-rate data (i.e., interoceptive inference phenotyping). For this analysis, we simulated 1000 trials under three canonical parameter settings designed to resemble potential neuropsychiatric phenotypes of interest: healthy interoception ($\alpha = 0.8$, $\beta = 0.8$, prior probability of parasympathetic policy = 55%), extremely precise interoceptive sensation ($\alpha = 0.8$, $\beta = 1$, prior probability of parasympathetic policy = 65%), and hyper-precise arousal priors ($\alpha = 0.8$, $\beta = 0.8$, prior probability of sympathetic policy = 75%). This prior probability is specified in the E-vector, which provides a prior probability for the two alternative policies. The values used here were $MDP.E = [11/20, 9/20]$ for the baseline policy (slight parasympathetic preference), $MDP.E = [13/20, 7/20]$ for the parasympathetic policy and $MDP.E = [5/20, 15/20]$ for the sympathetic policy.

The resulting time-series of systole events from each agent were then convolved with a canonical QRS-wave response function and transformed into normalized beat-to-beat RR-

intervals. To normalize the (arbitrary) sampling rate of each time-series, we assigned a 350ms repetition time (TR) for each state of the MDP simulation, such that the healthy agent had a heart rate of approximately 60 BPM. The time intervals between successive synthetic R-peaks were then calculated. As the RR interval data is unevenly sampled, the time series was linearly interpolated. The power spectrum was then estimated using Welch's method. In line with conventional HRV analysis, the power spectra were then categorized into four frequency bands corresponding to ultra-low (0–0.04 Hz), low (0.04–0.15 Hz), high (0.15–0.4 Hz), and very-high (> 0.4 Hz) frequency categories. Finally, to summarize the physiological response of each agent, we calculated the beats per minute (BPM) and the ratio between low and high frequency components (LF/HF), i.e., sympathovagal gain or balance. Sympathovagal balance is thought to index the balance of sympathetic and vagal outflows and is frequently implicated in stress and other psychophysiological and clinical disorders [74,75] but see [76–78] for critique. Sympathovagal balance was calculated as the ratio of area under the curve (AUC) for low and high-frequency HRV; AUC^{LF} / AUC^{HF} .

Results

Simulated physiology and perceptual active inference

To explore the face validity of our model, we simulated the basic psychophysiological behavior of our active inference agent. This involved simulating a fixed series of stimuli (states) in which the heartbeat was forced to remain relaxed—and only non-arousing (flower) stimuli were presented. On the 15th trial, an unexpected spider stimulus was presented, and the simulation continued for a further 85 trials. Thus, by evaluating the evolution of the agent's synthetic interoceptive physiology and exteroceptive beliefs, before and after the quiescent baseline period, we hoped to reproduce and illuminate well-known psychophysiological phenomena such as the defensive startle reflex [79,80] in which the heart-rate accelerates greatly following an unexpected, threatening stimulus.

This analysis, illustrated in Fig 3, revealed several interesting aspects of interoceptive active inference. Over 60 simulations there was a clear and robust increase in heart-rate acceleration, following the presentation of the unexpected or novel threat stimulus. During subsequent experiences of its own heartbeat and arousing versus non-arousing stimuli, this response habituates, resulting in a gradual heart-rate deceleration from the evoked cardiac response. This robust modulation of heart-rate was accompanied by a jump from an expected probability of encountering a spider of about 25% to almost 65% following the spider presentation. This combined response of both the heartbeat and fear-expectations is further underscored by the curious oscillation of cardiac states and the expected probability of observing a spider; note the uptick in expectations of approximately 5% on each systole event (denoted by the pink dotted line on Fig 3, right panel). In sum, it is as if the agent “hallucinates” or misperceives a threatening affective outcome both when baseline arousal is high, and when the linkage between interoceptive states and outcomes is severed.

A simple explanation for this result is that, during presentation of a stream of flowers, we can confidently infer a safe external environment. This accounts for the relatively low probability of inferred spiders in the earlier part of the plot. However, during systole, attenuated integration of exteroceptive data leads to greater uncertainty. Going from a confident inference in the absence of a spider to a more uncertain inference necessarily increases the probability of a scary environment during this cardiac phase. This offers a simple perspective on previous experimental work suggesting that fear-stimuli are potentiated when presented in synchrony with the heart [5,32]; namely, that a mechanism underlying this effect can be found in the link between cardiac active inference and affective expectations. In short, under generative models of an embodied world—in which sensory sampling depends upon fast fluctuations

in bodily states—there is a necessary dependency of Bayesian belief updating (i.e., perceptual inference) across all modalities on interoception.

When comparing these effects in the healthy agent to our sample of ‘lesion patients’, a few sensible but counter-intuitive consequences ensue. In the physiological domain, when presented with the unexpected arousal stimulus, the lesioned agent shows a blunted cardiac acceleration response, which remains diminished throughout the simulated trials. This blunting effect is mirrored for negative affective expectations in the immediate post-stimulus (e.g., trials 15–20) period, further underlining the close link between visceral and exteroceptive inference in our agent. The reason for this blunting is found in the variable exteroceptive precision anticipated during different cardiac phases (see also Fig 4B). A visual impression—consistent with a spider—is highly informative during diastole but must be treated with suspicion during the sensory-attenuated systolic phase. This implies a blunting of belief-updating in response to a spider, when we are unsure of cardiac phase (compared to when we are confident of a diastolic phase). A further interesting result is found when examining the controlled baseline period (trials 0–15); baseline fear expectations in the interoceptive lesion group are actually slightly enhanced by about 5–10% posterior probability. This lends an interesting, embodied twist to the literature on ‘circular inference’, psychosis and hallucinations [81,82], suggesting that the disruption of interoceptive precision may be one mechanism underlying hallucinations, particularly those that are affective and/or somatic in nature.

Simulating the influence of sensory precision on belief uncertainty

We next performed a series of simulations to tease apart how interoceptive and exteroceptive precision (and their disruption) influence variables important for metacognition such as the uncertainty in our agent’s beliefs. To do so, we first measured the Shannon entropy for interoceptive and exteroceptive inferences (summed across both factors of posterior beliefs) under a full range of precision settings from 0.5–1. To highlight the oscillatory nature of cardiac effects, we then calculated the same entropy measure separately for each cardiac state (early diastole, late diastole, systole). Finally, we compared these ‘healthy’ simulations to extreme degradations in sensory precision (exteroceptive and interoceptive ‘lesions’), to better understand how disruptions of each modality are integrated in perceptual uncertainty, a key first-order input for metacognitive inference. These lesions can be read as a simulation of total cardiac deafferentation through vagotomy (as was sometimes performed for stomach ulcers) [83], or in the case of perturbed higher-order arousal priors, lesions of the insular cortex [84,85].

This analysis revealed first of all that, in our simplified model, uncertainty is largely influenced by the unimodal precision of each domain. For both exteroceptive and interoceptive inferences, the slope of the uncertainty gradient (Fig 4A and 4C) was predominantly characterized by degradations in the precision of the corresponding modality. However, this modularity is not complete; exteroceptive uncertainty is at its lowest when interoceptive and exteroceptive precision are maximal. Similarly, although interoceptive uncertainty is largely driven by interoceptive precision, small interactions with exteroceptive precision can be observed in the plotted uncertainty gradient. One interesting isomorphism, however, is that overall interoceptive uncertainty is less affected by exteroceptive precision. This is likely due to that fact that in our model, the cardiac cycle directly modulates exteroceptive precision, whereas exteroceptive states only indirectly modulate interoceptive responses, via policy selection.

This intricate relationship of the cardiac cycle and uncertainty is further teased apart in Fig 4B, which shows clearly that exteroceptive confidence oscillates with each phase of the heart-beat, being highest at diastole. This is an unsurprising feature of our model: on each diastole, exteroceptive sensory precision drops effectively to null. Interestingly however, average

exteroceptive uncertainty is modulated in a fairly linear fashion by visceral and exteroceptive lesions: average entropy is increased modestly by lesioning interoceptive precision and more robustly by exteroceptive lesions. Whereas interoceptive lesions caused the greatest increase in interoceptive entropy, exteroceptive lesions seem to exert a specific effect of unbinding entropy from the individual cardiac state, again mirroring the isomorphic representation of these states in uncertainty. This is a sensible finding, as the manipulation leads to relatively high uncertainty in the mapping between hidden states and outcomes during all cardiac phases, not just during the previously attenuated systolic phase. This sort of chronic hypo-arousal—as a consequence of a failure to contextually modulate precision—is not unlike that which may underwrite the negative symptoms of schizophrenia or depression [86–89].

Synthetic Heart-Rate Variability (HRV) and embodied computational phenotyping

In our final set of simulations, we illustrated how the interoceptive inference approach developed here offers a new means for analyzing and interpreting fluctuations in observed physiological data. Our goal here was to demonstrate the potential for generative modelling and ‘embodied computational phenotyping’; i.e., the identification of specific parameters of brain-body interaction underlying maladaptive interoceptive inference in psychiatric and other health-harming disorders; e.g., [90]. This approach could for example, be expanded to model either brain-body interaction during explicit interoceptive decisions [91,92], or to cases of passive brain-body interaction, i.e., during experiments presenting stimuli at specific phases of the heart, breath, or stomach (see, e.g., [93]).

To this end, we generated synthetic cardiac data by convolving our train of cardiac events with an ECG response waveform. Following standard methods, we then calculated the normalized beat-to-beat intervals and performed a time-frequency analysis of the resulting RR-interval data. By repeating this analysis for a ‘healthy’ agent under normative values, an agent with overly precise interoceptive signals (i.e., $\beta = 1$), and an agent with overly precise prior beliefs about its own arousal, we illustrate how individual HRV fingerprints are linked to unique patterns of interoceptive active inference.

This analysis showed that, despite the exceedingly simple (biomechanically speaking) conditions of our model, sensible and interesting patterns of heart-rate variability emerge for different combinations of interoceptive sensory and prior precision. Specifically, we found that whereas the healthy agent exhibited a relatively relaxed profile in terms of heart rate and sympathovagal balance (BPM = 63.5, peak frequency = 0.14 Hz)—predominated by low versus high frequency oscillations (aLF/aHF = 0.52)—an agent with extremely precise visceral sensations exhibited a mild downshift in heart-rate (BPM = 59.5) coupled with an overall increase in high-frequency oscillations (aLF/aHF = 0.51, peak frequency = 0.25). In contrast, the agent with overly precise arousal priors showed a strong bimodal modulation of both ultra-low and ultra-high frequencies HRV (peak frequencies = 0.04 Hz & 0.34 Hz, respectively), coupled with a strong increase in heart-rate (BPM = 84) and high versus low-frequency outflow (aLF/aHF = 0.47). These results speak to the unique role of different active inference parameters in producing highly idiosyncratic patterns of HRV. In the future, our model may be enhanced to subserve computational phenotyping of individual differences and/or patient subgroups categorized by the balance of visceral precision and arousal policy priors from raw HRV data alone.

Discussion

In the present work, we have introduced a formal model of interoceptive inference as applied to exteroceptive, affective, and metacognitive signals. Through a variety of simulations, we

demonstrated that this model can reproduce a variety of psychological and physiological phenomena, each of which speak to a unique domain of the burgeoning interoceptive inference literature [27,29,33,35], and the application of interoceptive inference to computational psychiatry [42,44,46,47]. This formulation of interoceptive inference reproduces some of the finer details of physiological responses to arousing stimuli that, crucially, are emergent properties under the simple assumption interoceptive and exteroceptive states are coupled to one another.

The form of the generative model and proposed neurobiological belief updating scheme used in this paper are generic: exactly the same scheme has been used to simulate a whole range of processes, from neuroeconomic games to scene construction and attentional neglect [49,94]. The key aspect of the generative model introduced here is that the quality (i.e., precision) of sensory information depends upon fluctuations in (inferred) autonomic states. This simple fact underwrites all of the phenomenology illustrated above; both in terms of simulated physiology and accompanying belief updates. The explicit inclusion of interoception into active inference licenses us to talk about ‘affect’ in the sense that affective inference is thought to emerge under models that generate multimodal predictions that encompass the interoceptive domain. Furthermore, casting everything as inference enables a metacognitive stance on belief updating, in the sense that one can quantify uncertainty invested in beliefs about states of the body, states of the world and, indeed, states of (autonomic) action.

In particular, we show that by simulating periodic attenuation of exteroceptive sensory inputs by the cardiac cycle, affective, fear-like expectations become intrinsically linked to afferent interoceptive signals through a startle reflex-like phenomenon. This linkage not only induces oscillatory synchrony between the heartbeat and exteroceptive behavior, but also propagates to metacognitive uncertainty (i.e., the entropy of posterior beliefs). This latter finding speaks to numerous reports of metacognitive bias (e.g., confidence-accuracy dissociation) by illustrating how the precision of interoceptive states can directly influence exteroceptive uncertainty [1,95,96], which is presumed to be integrated by higher-order inferences in many process models of metacognition [66–68,70]. By simulating synthetic heart-rate variability (HRV) responses, we further illustrated how idiosyncratic patterns of aberrant interoceptive precision-weighting can be recovered through generative modelling of physiological responses, opening the door to computational phenotyping of disordered brain-body interaction in the spirit of (Schwartenbeck P and K Friston 2016). Further, we found that simulated lesions of visceral information increased our synthetic agent’s anticipation of encountering a fear-stimulus, highlighting the close relationship between aberrant interoceptive inference and a kind of exteroceptive ‘hallucination’. In what follows, we outline some of what we view as the most promising future directions for this work, sketch a proposed neuroanatomy underlying our model, and point out a few limitations for consideration.

By focusing on the periodic nature of the cardiac cycle, and concomitant influences on exteroceptive perception, our goal was to provide a proof-of-principle, illustrating how visceral and exteroceptive signals could be combined under active inference. Our aim was not to suggest that our model provides a definitive view of interoceptive inference or metacognition; indeed, we view the present work as a starting point that can be taken forward in a variety of research directions, some of which are outline below. In particular, we note that our model focuses primarily on interactions of a sensory nature between interoceptive and exteroceptive channels. It demonstrates that both perceptual beliefs and associated uncertainties are coupled under this model in interesting ways, which would presumably influence higher order emotional (e.g., whether one experiences fear or not) or metacognitive (e.g., whether one is certain if a heartbeat has occurred or not) inferences in hierarchical models, with temporal depth [37,47,48,70,91]. However, we note that our model does not explicitly parameterize these deep

temporal processes that are likely to be a fruitful area for future formal and empirical development.

In this paper, we formalized the hypothesis that frequently reported effects of cardiac timing on perception could arise as a function of periodic sensory attenuation—but the reader should feel encouraged to test their own competing hypotheses within the openly available MDP framework. Our intention here was also not to prioritize cardiac-brain interaction over e.g., gastric or respiratory cycles, but instead to provide a toy example, to show how these systems may be subjected to formal analyses. This was motivated by the large predominance of research on cardiac-brain interaction (see e.g., [5] for review); however, we do anticipate that the periodic attenuation of sensory precision by visceral signals is likely to provide a general explanation of brain-body interaction in other viscerosensory and visceromotor domains.

Neurophysiologically, the principal means by which cardiac signals influence the central nervous system is through the afferent cardiac baroreceptors. These pressure-sensitive neurons, located primarily in the aorta and carotid artery, are triggered by the systolic pressure wave generated when the heart contracts. Far from being restricted to homeostatic function only, it was first reported (nearly a century ago) that afferent baroreceptor outputs induce a general inhibitory effect on cortical processing [3,12,97]. These findings were later extended by Lacey and Lacey [7] who proposed the “neurovisceral afferent integration hypothesis”, positing that cardiac acceleration and deceleration serve to respectively disengage or engage with an exteroceptive stimulus via cortical inhibition.

In parallel, the soviet psychologist Evgeny Sokolov proposed that novelty-evoked heart-rate deceleration was a core component of the ‘orienting reflex’ [80]. By reducing overall cardiac output, this reflex served to limit the contribution of cardiac signals to cortical noise boosting overall signal-to-noise ratio. Sokolov (1963) first described the orienting reflex as an ‘embodied’ mechanism for boosting to signal-to-noise and thus enhancing processing of the oddball stimulus. The reflex consists primarily of the rapid deployment of saccades to the oddball stimulus, freezing of the muscles of the head and neck so as to orient the visual organs towards the stimulus, and an immediate cardiac deceleration. In light of their presumed inhibitory influence (Lacy and Lacy, 1968), the cardiac deceleration was thought to primarily reduce cortical noise; when coupled with the other bodily components of the response it was thought that effective overall signal would be maximized. In contrast, Sokolov theorized that the defensive startle reflex, in which an extremely strong (e.g., the loud bang of a starting gun) or unexpectedly aversive (e.g., the sudden presentation of a spider) stimulus evokes cardiac acceleration—facilitated the disengagement of cortical processing, to initiate fight-or-flight responses. These theories in turn sparked a wave of empirical studies attempting to link cardio-acceleration and deceleration responses to increased or decreased exteroceptive sensitivity, which continues to this day [2,4,5,8–10,98–103].

While these findings highlight the intricate relationship between cardiac timing and exteroceptive psychophysics, so far, a consistent pattern of findings (e.g., sensory signal enhancement and/or inhibition) has failed to emerge (see Elliott, 1972 for one critique). A cursory review of this literature reveals evidence for both exteroceptive enhancement and suppression, depending upon the specific nature of the exteroceptive stimuli (i.e., whether they are inherently aversive, sociocultural, or neutral in nature), the context of the arousal (including specific stimulus and response timing), and other psychophysiological moderators; such as age, gender, and overall physical fitness. Accordingly, more recent proposals have focused on more modality-specific exteroceptive enhancement by cardiac signals. For example, that cardiac-exteroceptive effects specifically potentiate fear or threat signals [32] or the generation of a subjective first-person viewpoint [22].

We offer a unique synthesis of these views, expressed in terms of interoceptive inference. In our model, the cyclic influence of the heart on exteroception is exerted primarily through the attenuation of sensory precision on each systolic contraction, which in turn influences the selected arousal policy as determined by the agent's preferences. The coupling of sensory attenuation to the cardiac cycle endorses the notion that baroreceptors exert an inhibitory influence on the brain. Beyond this direct effect, our model can also be understood in light of the well-known relationship between intrinsic noise fluctuations in the brain and cardio-respiratory cycles [13,104]. Physiological oscillations exert non-neuronal influences on spontaneous brain activity via a variety of more or less direct causal influences; for example, at each heart beat visual input to the retina is briefly attenuated by a pulsatile blood inflow. Similarly, with each cardio-respiratory cycle, fluctuations in cerebral pulsatile motion and blood pressure induce neurons to spontaneously fire, shaping the 'infraslow' brain dynamics [13,105,106] that influence the overall global dynamics of neural excitability and connectivity [107,108]. Our suggestion is that, insofar as the brain must model its own dynamic noise trajectories as a function of active self-inference, non-neuronal sources of variability such as inscribed by visceral rhythms must be incorporated within the brain's generative model of its own percepts. Interoceptive fluctuations are thus an important influence over the precision of exteroceptive sensory channels, and interoception is itself the means by which the brain infers and controls its own pathway through these precision trajectories. The modelling introduced here can thus be expanded beyond the cardiac domain to the more general problem of modelling how spontaneous fluctuations in neurovisceral cycles (including heart-rate variability) influence information processing and behaviour.

What then, explains the lack of consistent results within the cardiac timing literature? In contrast to the binary on/off hypotheses proposed by Lacey or Sokolov, our simulations highlight the context-sensitive manner by which ascending visceral signals modulate the precision of both interoceptive and exteroceptive inferences. For example, our simulation of the startle response (illustrated in Fig 3) clearly indicates that the functional impact of cardio-ballistic responses is coupled to the agent's baseline prior expectations, as well as the overall precision of active inference and policy selection. In this sense, whether a specific cardiac response is likely to potentiate or inhibit a specific domain (e.g., fear) depends upon the specific weighting of arousal policy priors, the precision of incoming exteroceptive and interoceptive sensations, and the linkages thereof as determined by the task itself. In other words, the specific balance of prior beliefs and sensory information, in a given cognitive or affective domain, must be addressed before one can predict the exact directionality of an interoceptive effect on perception, or vice versa. Here, we modelled the generation of arousal as a function of hyper-parameters governing the preferred autonomic policy. In the future this can be unpacked further by examining the divergence between prior and posterior beliefs about these policies (e.g., through inferred epistemic value). Through Landauer's principle, this divergence may be equated with the associated metabolic costs of computation and the conceptual notion of interoceptive self-modelling [35,109,110].

It is also interesting to consider the relationship between this model and theories in psychology, such as Mowrer's 2-factor theory [111], which posits that affective learning such as threat conditioning depends upon an interaction between simple Pavlovian conditioning and the general law of effect. Active inference posits a similar kind of circular causality, so these two factors are mutually interlinked. That is to say, you infer that you are scared based on some set of physiological prediction errors and contextual priors (e.g., I see a bear, and my heart is racing, so I must be afraid), but also that fear is itself a hypothesis which engages physiological reflexes through active-inference (I am running away, therefore I should be aroused and scared). In this sense, the model presented here could be seen as applying the two-factor approach to a particular interaction between instrumental and Pavlovian forms of learning.

The computational neuroanatomy of interoceptive inference

Having addressed the construct validity of our model, we now speculate as to some likely neuronal substrates of the message passing implied by variational inference. Interoceptive inference can be broken down into four core functional domains: basic sensory-motor control, conscious interoceptive (perceptual) awareness, metacognitive monitoring, and hedonic (intrinsic) value. In our model, we focused primarily on the simplest possible implementation of interoceptive inference, corresponding to the sensory-motor domain (i.e., ascending and descending cardiac pathways) and their low-level interaction with exteroceptive inference, via neuromodulatory gain control. Future work will benefit from expanding upon our representation of uncertainty to include the computation of epistemic and/or intrinsic value as proxies for these higher-order interoceptive and metacognitive systems [112,113].

Accordingly, in our sketch of the putative neuroanatomy underlying cardiac active inference (Fig 6), we focus primarily on the neuronal substrates that inscribe low-level viscerosensory and visceromotor control, as well as some hierarchically superior regions related to emotional salience and interoceptive awareness. For simplicity, our model depicts only the minimal neuronal message passing scheme implied by our generative model; as such, we have omitted many of the intermediary relay nodes; e.g., in the thalamus and ventral visual stream. Afferent baroreceptor signals are transmitted along the ascending vagus to the rostrum of the nucleus tractus solitarius [115,116]. From here, ascending viscerosensory signals are projected via brainstem and midbrain nuclei to the thalamus, somatosensory cortex, and posterior insula [117,118]; ascending cardio-sensory outcomes are thus encoded in the NTS and then passed to the posterior insular cortex (PIC) as inferred interoceptive states. The PIC has a well-known role as primary viscerosensory cortex; electrical stimulation of this area elicits phantom visceral sensations (e.g., pain, heart-rate acceleration) [119,120] and bolus isoproterenol infusions increase the intensity of cardiorespiratory sensations and concomitant PIC activations [121,122]. In parallel, visual sensory outcomes are passed via the second cranial nerve to the superior colliculus, where they inform exteroceptive inference in the amygdala, which is well-situated to process salient emotional stimuli [123,124]. These interoceptive and exteroceptive expectations then converge in the anterior insular cortex (AIC), where they inform the selection of the appropriate autonomic policy. Finally, the selected policy is passed down the hierarchy via descending pathways (e.g., by von Economo neurons), to eventually engage the rostral nucleus ambiguus and descending vagus, decelerating the heart-rate when the relaxed policy is selected. Collectively, the scheme represents a multimodal reflex arc interlinking exteroceptive and interoceptive domains to specific patterns of cardio-ballistic responses.

Within this scheme, we suggest that the rostral anterior cingulate (ACC) controls the precision of ascending visceral outcomes and inferred interoceptive states via neuromodulatory gain control [125,126]. Further, interoceptive and exteroceptive state precisions here interact indirectly through global neuromodulatory influences, possibly through regulation of noradrenaline by the ACC (via descending influence on the locus coeruleus). Neurobiologically and functionally speaking, the AIC and ACC share similar profiles; both are densely populated with Von Economo neurons (VENs), which are well-suited for the long-range modulation of neural activity across the cortex [127], and also contain diverse populations of noradrenergic, dopaminergic, and opioidergic neurons. Both regions further share an integrative connectivity structure, with projections to both lower-level visceral-motor brainstem nuclei and higher-order regions implicated in decision-making, metacognition, and self-awareness, such as the ventromedial and dorsomedial prefrontal cortices [128–132]. However, the AIC is more densely interconnected with the PIC whereas the ACC is more closely related to uncertainty and decision-making. On this basis, we propose that whereas the AIC integrates the visceral

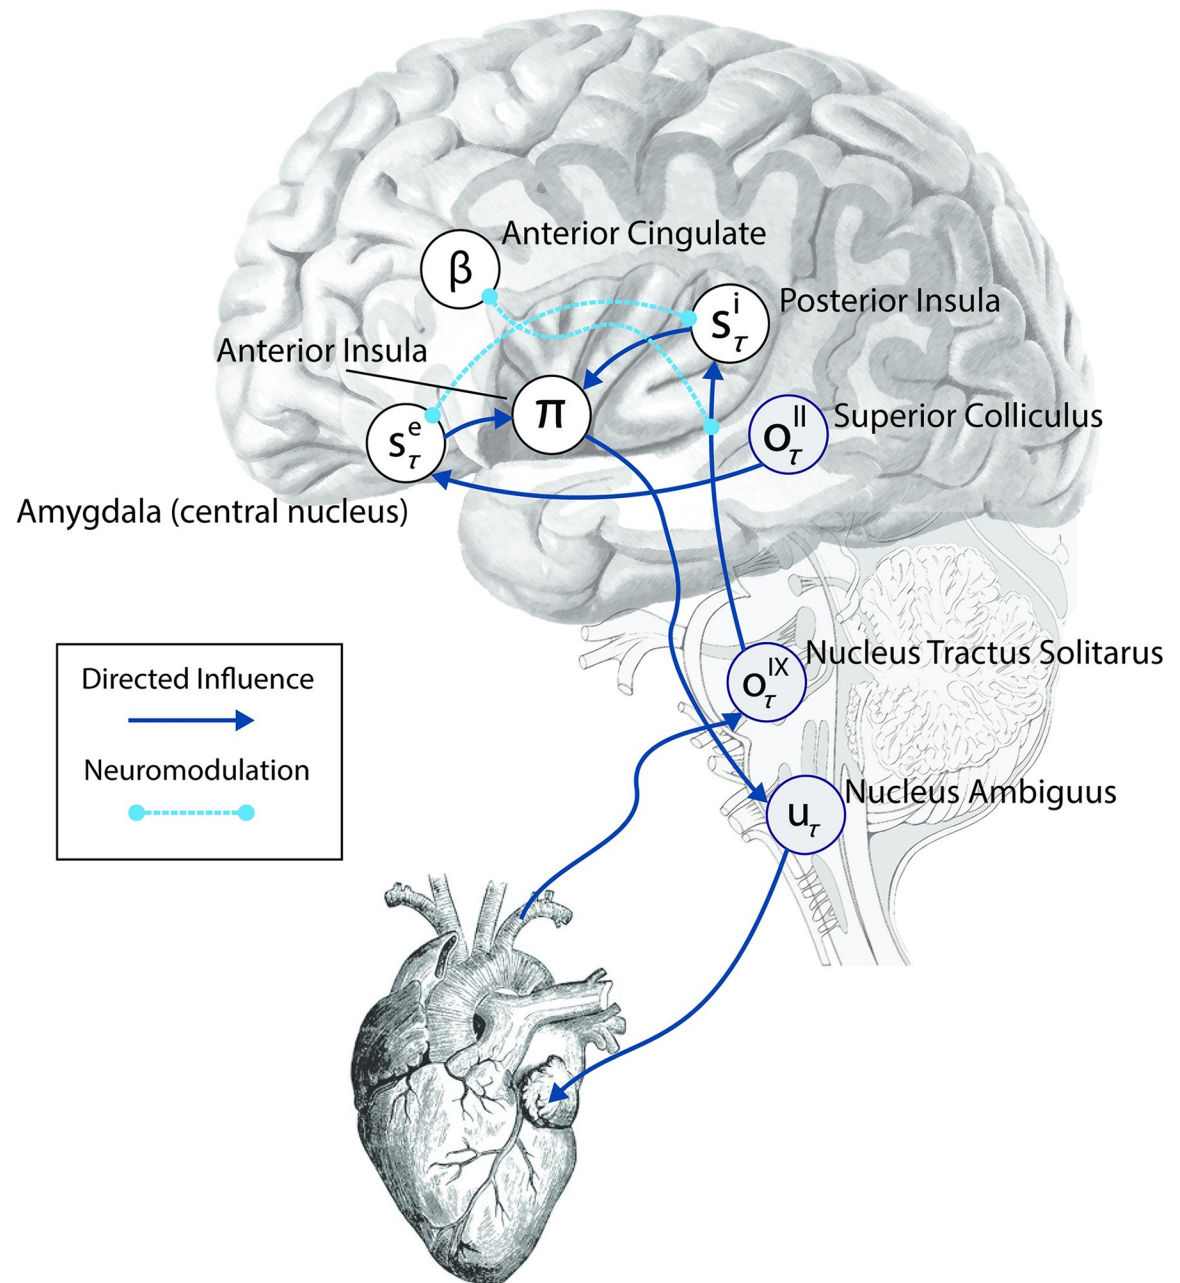


Fig 6. Computational neuroanatomy of interoception. The schematic above shows the form of the neuronal message passing implied by active inference for the generative model depicted in Fig 2. We have related this to the anatomical networks that could implement these inferences. The sensory observations in our simulations are visual and interoceptive (cardiac). These sensations are carried by cranial nerves II and IX respectively. Cranial nerve II targets the superior colliculus in the midbrain. This structure sends short latency visual data to the amygdala, which is well placed to make inferences about emotionally salient stimuli. The amygdala additionally receives downstream visual data from the ventral visual stream in the temporal lobe. Cranial nerve IX carries information from the carotid sinus baroreceptors to the nucleus tractus solitarus in the brainstem. This nucleus communicates with the posterior insula (via thalamic and PAG relays); the anterior cingulate monitors and controls the precision of this ascending visceral information via neuromodulation, possibly via feedback through noradrenergic pathways (not shown). The posterior insula and amygdala interact with one another but also project to the anterior insula. This targets the nucleus ambiguus (via brainstem relays such as the periaqueductal gray), which gives rise to the vagus (X) nerve. The vagus nerve targets neurons in the cardiac plexus that project to both the sinoatrial node and the atrioventricular node of the heart, slowing its rhythm. The nucleus tractus solitarus additionally participates in a reflex loop implicating the sympathetic control of the cardiac cycle, but this is omitted for simplicity. The functional anatomy suggested here implies the anterior insula might play a similar computational role in autonomic policy selection to the basal ganglia in selection of policies involving the skeletal muscles [114]. Note that inscribed directed influences (blue arrows), are not assumed to be monosynaptic—for simplicity, many intermediary relay nodes have been omitted.

<https://doi.org/10.1371/journal.pcbi.1010490.g006>

and exteroceptive states required for the regulation of arousal policies, the ACC is likely to regulate the gain or precision of these interactions. It is worth noting that this model may explain the widespread, seemingly unspecific activation profiles of these areas [133,134], as the generative model specified here suggests both form part of an integrated hierarchical circuit by which interoceptive and exteroceptive states interact: e.g., either through the regulation of arousal policies or through the modulation of ascending viscerosensory precision.

What about metacognitive or reward-related interoceptive processes? Although we do not model these higher-order functions, the model can be expanded to include the explicit representation of policy uncertainty and epistemic value as the mechanisms underlying metacognitive self-inference; i.e., the integrative self-model that combines exteroceptive and interoceptive predictions into a conscious schema [39,68]. In this case, we would expect that the VMPFC and DLPFC are likely to be engaged in inferences about variables (e.g., those derived from expected free energy such as epistemic and intrinsic value) that contextualize the inferences performed by the AIC and ACC over longer timescales [55,60].

Limitations and future directions

The model and simulations presented here represent a minimal proof-of-principle demonstrating how cyclic interactions of interoceptive and exteroceptive perception could arise directly from the principles of active inference. Here, our primary goal was to move the literature beyond purely conceptual analyses of 'interoceptive inference', to provide a formal model sub-serving direct hypothesis testing. As such, we focus primarily on reproducing commonly reported phenomena, rather than empirical cross-validation or biological plausibility. The model as presented is thus a demonstration of formal theory, and substantive further work will be needed to develop an observation model of empirical heart-brain interactions or associated physiological data: c.f., [46,135].

Along these lines, we imagine a variety of interesting further applications of the modelling developed herein. For example, the present MDP scheme could be expanded to include more biologically realistic cardiac parameters, or to include other visceral modalities such as gastric or respiratory fluctuations. Similarly, the exteroceptive states modelled here could be adapted to a variety of experimental tasks to capture embodied influences on, for example, active spatial navigation [136–138], active reward learning [139,140], interaction between the cardiac cycle and ballistic saccades [15,17,50], or metacognitive inferences [1,65,112]. These and other future directions will hopefully guide a newly embodied approach to computational psychiatry, enabling the detailed phenotyping of clinical populations in terms of aberrant interoceptive inference.

Data and code availability

The underlying MDP scheme here is available as part of the open-access distribution of SPM12. A demonstration of the scheme can be accessed by typing `>>DEM` into the Matlab command prompt and selecting the **Interoception** demo from the graphical user interface that appears. Further, all the code required to generate the simulations and figures herein can be found at the following github page: <https://github.com/embodied-computation-group/cardiac-active-inference>. Full installation instructions can be found in the following code wrapper, which reproduces all figures and simulations: https://github.com/embodied-computation-group/cardiac-active-inference/blob/master/paper_wrapper.m

Author Contributions

Conceptualization: Micah Allen, Thomas Parr, Karl J. Friston.

Formal analysis: Micah Allen, Andrew Levy, Thomas Parr, Karl J. Friston.

Methodology: Micah Allen, Andrew Levy, Karl J. Friston.

Software: Andrew Levy, Thomas Parr, Karl J. Friston.

Supervision: Thomas Parr, Karl J. Friston.

Visualization: Micah Allen, Andrew Levy, Thomas Parr.

Writing – original draft: Micah Allen, Andrew Levy, Thomas Parr, Karl J. Friston.

Writing – review & editing: Micah Allen, Andrew Levy, Thomas Parr, Karl J. Friston.

References

1. Allen M, Frank D, Schwarzkopf DS, Fardo F, Winston JS, Hauser TU, et al. Unexpected arousal modulates the influence of sensory noise on confidence. Shan H, editor. *eLife*. 2016 Oct 25; 5:e18103. <https://doi.org/10.7554/eLife.18103> PMID: 27776633
2. Azevedo RT, Garfinkel SN, Critchley HD, Tsakiris M. Cardiac afferent activity modulates the expression of racial stereotypes. *Nat Commun*. 2017 Jan 17; 8:13854. <https://doi.org/10.1038/ncomms13854> PMID: 28094772
3. Bonvallet M, Bloch V. Bulbar control of cortical arousal. *Science*. 1961; 133(3459):1133–4. <https://doi.org/10.1126/science.133.3459.1133> PMID: 17754758
4. Cohen R, Lieb H, Rist F. Loudness judgments, evoked potentials, and reaction time to acoustic stimuli early and late in the cardiac cycle in chronic schizophrenics. *Psychiatry Res*. 1980; 3(1):23–9. [https://doi.org/10.1016/0165-1781\(80\)90044-x](https://doi.org/10.1016/0165-1781(80)90044-x) PMID: 6934554
5. Garfinkel SN, Minati L, Gray MA, Seth AK, Dolan RJ, Critchley HD. Fear from the Heart: Sensitivity to Fear Stimuli Depends on Individual Heartbeats. *J Neurosci*. 2014 May 7; 34(19):6573–82. <https://doi.org/10.1523/JNEUROSCI.3507-13.2014> PMID: 24806682
6. Hauser TU, Allen M, Rees G, Dolan RJ. Metacognitive impairments extend perceptual decision making weaknesses in compulsivity. *Sci Rep*. 2017 Jul 26; 7(1):6614. <https://doi.org/10.1038/s41598-017-06116-z> PMID: 28747627
7. Lacey BC, Lacey JI. Two-way communication between the heart and the brain: Significance of time within the cardiac cycle. *Am Psychol*. 1978; 33(2):99.
8. Park HD, Correia S, Ducorps A, Tallon-Baudry C. Spontaneous fluctuations in neural responses to heartbeats predict visual detection. *Nat Neurosci*. 2014 Apr; 17(4):612–8. <https://doi.org/10.1038/nn.3671> PMID: 24609466
9. Salomon R, Ronchi R, Dónz J, Bello-Ruiz J, Herbelin B, Martet R, et al. The Insula Mediates Access to Awareness of Visual Stimuli Presented Synchronously to the Heartbeat. *J Neurosci*. 2016 May 4; 36(18):5115–27. <https://doi.org/10.1523/JNEUROSCI.4262-15.2016> PMID: 27147663
10. Velden M, Juris M. Perceptual Performance as a Function of Intra-Cycle Cardiac Activity. *Psychophysiology*. 1975; 12(6):685–92. <https://doi.org/10.1111/j.1469-8986.1975.tb00075.x> PMID: 1187975
11. Zelano C, Jiang H, Zhou G, Arora N, Schuele S, Rosenow J, et al. Nasal Respiration Entrain Human Limbic Oscillations and Modulates Cognitive Function. *J Neurosci*. 2016 Dec 7; 36(49):12448–67. <https://doi.org/10.1523/JNEUROSCI.2586-16.2016> PMID: 27927961
12. Bonvallet M, Dell P, Hiebel G. Tonus sympathique et activité électrique corticale. *Electroencephalogr Clin Neurophysiol*. 1954; 6:119–44. [https://doi.org/10.1016/0013-4694\(54\)90011-5](https://doi.org/10.1016/0013-4694(54)90011-5) PMID: 13141926
13. Karavaev AS, Kiselev AR, Runnova AE, Zhuravlev MO, Borovkova EI, Prokhorov MD, et al. Synchronization of infra-slow oscillations of brain potentials with respiration. *Chaos Interdiscip J Nonlinear Sci*. 2018 Aug 1; 28(8):081102. <https://doi.org/10.1063/1.5046758> PMID: 30180638
14. Varga S, Heck DH. Rhythms of the body, rhythms of the brain: Respiration, neural oscillations, and embodied cognition. *Conscious Cogn*. 2017 Nov 1; 56:77–90. <https://doi.org/10.1016/j.concog.2017.09.008> PMID: 29073509
15. Galvez-Pol A, McConnell R, Kilner J. Active sampling during visual search is modulated by the cardiac cycle. *bioRxiv*. 2018;405902.
16. Kunzendorf S, Klotzsche F, Akbal M, Villringer A, Ohl S, Gaebler M. Active information sampling varies across the cardiac cycle. *Psychophysiology*. 2019;e13322. <https://doi.org/10.1111/psyp.13322> PMID: 30620083

17. Ohi S, Wohltat C, Kliegl R, Pollatos O, Engbert R. Microsaccades are coupled to heartbeat. *J Neurosci*. 2016; 36(4):1237–41. <https://doi.org/10.1523/JNEUROSCI.2211-15.2016> PMID: 26818511
18. Allen M, Varga S, Heck DH. Respiratory rhythms of the predictive mind. *Psychol Rev*. 2022;No Pagination Specified-No Pagination Specified. <https://doi.org/10.1037/rev0000391> PMID: 35980689
19. Herrero JL, Khuvis S, Yeagle E, Cerf M, Mehta AD. Breathing above the brainstem: Volitional control and attentional modulation in humans. *J Neurophysiol*. 2017; <https://doi.org/10.1152/jn.00551.2017> PMID: 28954895
20. Tort ABL, Ponsel S, Jessberger J, Yanovsky Y, Brankačk J, Draguhn A. Parallel detection of theta and respiration-coupled oscillations throughout the mouse brain. *Sci Rep*. 2018; 8(1):6432–6432. <https://doi.org/10.1038/s41598-018-24629-z> PMID: 29691421
21. Tort ABL, Brankačk J, Draguhn A. Respiration-Entrained Brain Rhythms Are Global but Often Overlooked. *Trends Neurosci*. 2018 Apr 1; 41(4):186–97. <https://doi.org/10.1016/j.tins.2018.01.007> PMID: 29429805
22. Azzalini D, Rebollo I, Tallon-Baudry C. Visceral Signals Shape Brain Dynamics and Cognition. *Trends Cogn Sci [Internet]*. 2019 Apr 30 [cited 2019 May 12]; Available from: <http://www.sciencedirect.com/science/article/pii/S1364661319300890> <https://doi.org/10.1016/j.tics.2019.03.007> PMID: 31047813
23. Rebollo I, Devauchelle AD, Béranger B, Tallon-Baudry C. Stomach-brain synchrony reveals a novel, delayed-connectivity resting-state network in humans. Critchley H, editor. *eLife*. 2018 Mar 21; 7: e33321. <https://doi.org/10.7554/eLife.33321> PMID: 29561263
24. Richter CG, Babo-Rebello M, Schwartz D, Tallon-Baudry C. Phase-amplitude coupling at the organism level: The amplitude of spontaneous alpha rhythm fluctuations varies with the phase of the infra-slow gastric basal rhythm. *NeuroImage*. 2017 Feb 1; 146:951–8. <https://doi.org/10.1016/j.neuroimage.2016.08.043> PMID: 27557620
25. Allen M. Unravelling the Neurobiology of Interoceptive Inference. *Trends Cogn Sci*. 2020 Apr 1; 24(4):265–6. <https://doi.org/10.1016/j.tics.2020.02.002> PMID: 32160563
26. Allen M, Legrand N, Correa CMC, Fardo F. Thinking through prior bodies: autonomic uncertainty and interoceptive self-inference. *Behav Brain Sci [Internet]*. 2020 ed [cited 2020 Aug 28]; 43. Available from: <https://www.cambridge.org/core/journals/behavioral-and-brain-sciences/article/thinking-through-prior-bodies-autonomic-uncertainty-and-interoceptive-selfinference/131DD3A5A01DE226099AAE7A64302A7F> <https://doi.org/10.1017/S0140525X19002899> PMID: 32460937
27. Allen M, Friston KJ. From cognitivism to autopoiesis: towards a computational framework for the embodied mind Synthese. 2018 Jun 1; 195(6):2459–82. <https://doi.org/10.1007/s11229-016-1288-5> PMID: 29887647
28. Apps MAJ, Tsakiris M. The free-energy self: A predictive coding account of self-recognition. *Neurosci Biobehav Rev*. 2014 Apr 1; 41:85–97. <https://doi.org/10.1016/j.neubiorev.2013.01.029> PMID: 23416066
29. Barrett LF, Simmons WK. Interoceptive predictions in the brain. *Nat Rev Neurosci*. 2015 Jul; 16(7):419–29. <https://doi.org/10.1038/nrn3950> PMID: 26016744
30. Chanes L, Barrett LF. Redefining the Role of Limbic Areas in Cortical Processing. *Trends Cogn Sci*. 2016 Feb 1; 20(2):96–106. <https://doi.org/10.1016/j.tics.2015.11.005> PMID: 26704857
31. Gallagher S, Allen M. Active inference, enactivism and the hermeneutics of social cognition. *Synthese*. 2018 Jun 1; 195(6):2627–48. <https://doi.org/10.1007/s11229-016-1269-8> PMID: 29887648
32. Garfinkel SN, Critchley HD. Threat and the Body: How the Heart Supports Fear Processing. *Trends Cogn Sci*. 2016 Jan 1; 20(1):34–46. <https://doi.org/10.1016/j.tics.2015.10.005> PMID: 26628111
33. Seth AK. Interoceptive inference, emotion, and the embodied self. *Trends Cogn Sci*. 2013 Nov 1; 17(11):565–73. <https://doi.org/10.1016/j.tics.2013.09.007> PMID: 24126130
34. Seth AK, Friston KJ. Active interoceptive inference and the emotional brain. *Philos Trans R Soc B Biol Sci*. 2016 Nov 19; 371(1708):20160007. <https://doi.org/10.1098/rstb.2016.0007> PMID: 28080966
35. Seth AK, Tsakiris M. Being a Beast Machine: The Somatic Basis of Selfhood. *Trends Cogn Sci*. 2018 Nov 1; 22(11):969–81. <https://doi.org/10.1016/j.tics.2018.08.008> PMID: 30224233
36. Smith R, Parr T, Friston KJ. Simulating Emotions: An Active Inference Model of Emotional State Inference and Emotion Concept Learning. *Front Psychol*. 2019; 10:2844. <https://doi.org/10.3389/fpsyg.2019.02844> PMID: 31920873
37. Smith R, Thayer JF, Khalsa SS, Lane RD. The hierarchical basis of neurovisceral integration. *Neurosci Biobehav Rev*. 2017 Apr 1; 75:274–96. <https://doi.org/10.1016/j.neubiorev.2017.02.003> PMID: 28188890

38. Ainley V, Apps MAJ, Fotopoulou A, Tsakiris M. 'Bodily precision': a predictive coding account of individual differences in interoceptive accuracy. *Philos Trans R Soc B Biol Sci.* 2016 Nov 19; 371(1708):20160003. <https://doi.org/10.1098/rstb.2016.0003> PMID: 28080962
39. Allen M, Tsakiris M. The body as first prior: Interoceptive predictive processing and the primacy. *Interoceptive Mind Homeost Aware.* 2018; 27.
40. Corcoran AW, Pezzulo G, Hohwy J. From Allostatic Agents to Counterfactual Cognisers: Active Inference, Biological Regulation, and the Origins of Cognition. 2020 Feb 29 [cited 2021 Mar 26]; Available from: <https://www.preprints.org/manuscript/201911.0083/v2>
41. Khalsa SS, Lapidus RC. Can Interoception Improve the Pragmatic Search for Biomarkers in Psychiatry? *Front Psychiatry* [Internet]. 2016 Jul 25 [cited 2020 Feb 18]; 7. Available from: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4958623/> <https://doi.org/10.3389/fpsy.2016.00121> PMID: 27504098
42. Owens AP, Allen M, Ondobaka S, Friston KJ. Interoceptive inference: From computational neuroscience to clinic. *Neurosci Biobehav Rev.* 2018 Jul 1; 90:174–83. <https://doi.org/10.1016/j.neubiorev.2018.04.017> PMID: 29694845
43. Petzschner FH, Garfinkel SN, Paulus MP, Koch C, Khalsa SS. Computational Models of Interoception and Body Regulation. *Trends Neurosci.* 2021 Jan 1; 44(1):63–76. <https://doi.org/10.1016/j.tins.2020.09.012> PMID: 33378658
44. Petzschner FH, Weber LAE, Gard T, Stephan KE. Computational Psychosomatics and Computational Psychiatry: Toward a Joint Framework for Differential Diagnosis. *Biol Psychiatry.* 2017 Sep 15; 82(6):421–30. <https://doi.org/10.1016/j.biopsych.2017.05.012> PMID: 28619481
45. Stephan KE, Manjaly ZM, Mathys CD, Weber LAE, Paliwal S, Gard T, et al. Allostatic Self-efficacy: A Metacognitive Theory of Dyshomeostasis-Induced Fatigue and Depression. *Front Hum Neurosci* [Internet]. 2016 [cited 2019 Oct 14];10. Available from: <https://www.frontiersin.org/articles/10.3389/fnhum.2016.00550/full> PMID: 27895566
46. Smith R, Mayeli A, Taylor S, Al Zoubi O, Naegel J, Khalsa SS. Gut inference: A computational modeling approach. *Biol Psychol.* 2021 Sep 1; 164:108152. <https://doi.org/10.1016/j.biopsycho.2021.108152> PMID: 34311031
47. Smith R, Kuplicki R, Feinstein J, Forthman KL, Stewart JL, Paulus MP, et al. A Bayesian computational model reveals a failure to adapt interoceptive precision estimates across depression, anxiety, eating, and substance use disorders. *PLOS Comput Biol.* 2020 Dec 14; 16(12):e1008484. <https://doi.org/10.1371/journal.pcbi.1008484> PMID: 33315893
48. Hesp C, Smith R, Parr T, Allen M, Friston KJ, Ramstead MJD. Deeply Felt Affect: The Emergence of Valence in Deep Active Inference. *Neural Comput.* 2021 Feb 1; 33(2):398–446. https://doi.org/10.1162/neco_a_01341 PMID: 33253028
49. Friston KJ, FitzGerald T, Rigoli F, Schwartenbeck P, Pezzulo G. Active inference: a process theory. *Neural Comput.* 2017; 29(1):1–49. https://doi.org/10.1162/NECO_a_00912 PMID: 27870614
50. Mirza MB, Adams RA, Mathys CD, Friston KJ. Scene Construction, Visual Foraging, and Active Inference. *Front Comput Neurosci.* 2016; 10:56. <https://doi.org/10.3389/fncom.2016.00056> PMID: 27378899
51. Da Costa L, Parr T, Sajid N, Veselic S, Neacsu V, Friston K. Active inference on discrete state-spaces: A synthesis. *J Math Psychol.* 2020 Dec 1; 99:102447. <https://doi.org/10.1016/j.jmp.2020.102447> PMID: 33343039
52. Attias H. Planning by probabilistic inference. In: *Proc of the 9th Int Workshop on Artificial Intelligence and Statistics.* 2003.
53. Botvinick M, Toussaint M. Planning as inference. *Trends Cogn Sci.* 2012 Oct 1; 16(10):485–8. <https://doi.org/10.1016/j.tics.2012.08.006> PMID: 22940577
54. Parr T, Markovic D, Kiebel SJ, Friston KJ. Neuronal message passing using Mean-field, Bethe, and Marginal approximations. *Sci Rep.* 2019 Feb 13; 9(1):1889. <https://doi.org/10.1038/s41598-018-38246-3> PMID: 30760782
55. Friston KJ, Parr T, de Vries B. The graphical brain: Belief propagation and active inference. *Netw Neurosci.* 2017 Jun 7; 1(4):381–414. https://doi.org/10.1162/NETN_a_00018 PMID: 29417960
56. Barto A, Mirolli M, Baldassarre G. Novelty or surprise? *Front Psychol.* 2013; 4:907. <https://doi.org/10.3389/fpsyg.2013.00907> PMID: 24376428
57. Itti L, Baldi P. Bayesian surprise attracts human attention. *Vision Res.* 2009; 49(10):1295–306. <https://doi.org/10.1016/j.visres.2008.09.007> PMID: 18834898
58. Oudeyer PY, Kaplan F. What is intrinsic motivation? A typology of computational approaches. *Front Neurobotics.* 2009; 1:6.
59. Schmidhuber J. Formal theory of creativity, fun, and intrinsic motivation. *IEEE Trans Auton Ment Dev.* 2010; 2(3):230–47.

60. Friston K, Rigoli F, Ognibene D, Mathys C, Fitzgerald T, Pezzulo G. Active inference and epistemic value. *Cogn Neurosci*. 2015; 6(4):187–214. <https://doi.org/10.1080/17588928.2015.1020053> PMID: 25689102
61. Al E, Iliopoulos F, Forschack N, Nierhaus T, Grund M, Motyka P, et al. Heart–brain interactions shape somatosensory perception and evoked potentials. *Proc Natl Acad Sci*. 2020 May 12; 117(19):10575–84. <https://doi.org/10.1073/pnas.1915629117> PMID: 32341167
62. Barnett WH, Baekey DM, Paton JFR, Dick TE, Wehrwein EA, Molkov YI. Heartbeats entrain breathing via baroreceptor-mediated modulation of expiratory activity. *Exp Physiol*. 2021; 106(5):1181–95. <https://doi.org/10.1113/EP089365> PMID: 33749038
63. Leganes-Fonteneau M, Buckman JF, Suzuki K, Pawlak A, Bates ME. More than meets the heart: systolic amplification of different emotional faces is task dependent. *Cogn Emot*. 2021 Feb 17; 35(2):400–8. <https://doi.org/10.1080/02699931.2020.1832050> PMID: 33073669
64. Conant RC, Ashby W. Every good regulator of a system must be a model of that system. *Int J Syst Sci*. 1970 Oct 1; 1(2):89–97.
65. Hauser TU, Allen M, Purg N, Moutoussis M, Rees G, Dolan RJ. Noradrenaline blockade specifically enhances metacognitive performance. *eLife*. 2017 10;6. <https://doi.org/10.7554/eLife.24901> PMID: 28489001
66. Fleming SM. Awareness as inference in a higher-order state space. *Neurosci Conscious* [Internet]. 2020 Jan 1 [cited 2021 Mar 23]; 2020(niz020). Available from: <https://doi.org/10.1093/nc/niz020>
67. Fleming SM, Daw ND. Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychol Rev*. 2016; 124(1):91.
68. Nikolova N, Waade PT, Friston KJ, Allen M. What Might Interoceptive Inference Reveal about Consciousness? *Rev Philos Psychol* [Internet]. 2021 Aug 14 [cited 2021 Aug 16]; Available from: <https://doi.org/10.1007/s13164-021-00580-3>
69. Whyte CJ. Integrating the global neuronal workspace into the framework of predictive processing: Towards a working hypothesis. *Conscious Cogn*. 2019 Aug 1; 73:102763. <https://doi.org/10.1016/j.concog.2019.102763> PMID: 31228697
70. Whyte CJ, Smith R. The predictive global neuronal workspace: A formal active inference model of visual consciousness. *Prog Neurobiol*. 2021 Apr 1; 199:101918. <https://doi.org/10.1016/j.pneurobio.2020.101918> PMID: 33039416
71. Friston KJ, Wiese W, Hobson JA. Sentience and the Origins of Consciousness: From Cartesian Duality to Markovian Monism. *Entropy*. 2020 May; 22(5):516. <https://doi.org/10.3390/e22050516> PMID: 33286288
72. Limanowski J. Enacting Proprioceptive Predictions in the Rubber Hand Illusion. *Front Hum Neurosci*. 2022 Feb 16; 16:839890. <https://doi.org/10.3389/fnhum.2022.839890> PMID: 35250522
73. Limanowski J, Friston K. Attenuating oneself: An active inference perspective on “selfless” experiences. *Philos Mind Sci*. 2020 Mar 24; 1(1):1–16.
74. Malliani A, Pagani M, Lombardi F, Cerutti S. Cardiovascular neural regulation explored in the frequency domain. *Circulation*. 1991 Aug; 84(2):482–92. <https://doi.org/10.1161/01.cir.84.2.482> PMID: 1860193
75. Strigo A, Craig AD. Interoception, homeostatic emotions and sympathovagal balance. *Philos Trans R Soc B Biol Sci*. 2016 Nov 19; 371(1708):20160010. <https://doi.org/10.1098/rstb.2016.0010> PMID: 28080968
76. Eckberg Dwain L. Sympathovagal Balance. *Circulation*. 1997 Nov 4; 96(9):3224–32.
77. Heathers JAJ. Sympathovagal balance from heart rate variability: an obituary. *Exp Physiol*. 2012; 97(4):556–556. <https://doi.org/10.1113/expphysiol.2011.063867> PMID: 22525665
78. Quintana DS, Heathers JAJ. Considerations in the assessment of heart rate variability in biobehavioral research. *Front Psychol* [Internet]. 2014 [cited 2022 Aug 3];5. Available from: <https://www.frontiersin.org/articles/10.3389/fpsyg.2014.00805> PMID: 25101047
79. Graham FK, Clifton RK. Heart-rate change as a component of the orienting response. *Psychol Bull*. 1966; 65(5):305–20. <https://doi.org/10.1037/h0023258> PMID: 5325894
80. Sokolov EN. Higher nervous functions: The orienting reflex. *Annu Rev Physiol*. 1963; 25(1):545–80. <https://doi.org/10.1146/annurev.ph.25.030163.002553> PMID: 13977960
81. Denève S, Jardri R. Circular inference: mistaken belief, misplaced trust. *Curr Opin Behav Sci*. 2016 Oct 1; 11:40–8.
82. Powers AR, Mathys C, Corlett PR. Pavlovian conditioning–induced hallucinations result from overweighting of perceptual priors. *Science*. 2017 Aug 11; 357(6351):596–600. <https://doi.org/10.1126/science.aan3458> PMID: 28798131

83. Tzeng YC, Larsen PD, Galletly DC. Mechanism of cardioventilatory coupling: insights from cardiac pacing, vagotomy, and sinoaortic denervation in the anesthetized rat. *Am J Physiol-Heart Circ Physiol*. 2007 Apr 1; 292(4):H1967–77. <https://doi.org/10.1152/ajpheart.01049.2006> PMID: 17172271
84. Khalsa SS, Rudrauf D, Feinstein JS, Tranel D. The pathways of interoceptive awareness. *Nat Neurosci*. 2009 Dec; 12(12):1494–6. <https://doi.org/10.1038/nn.2411> PMID: 19881506
85. Wang X, Wu Q, Egan L, Gu X, Liu P, Gu H, et al. Anterior insular cortex plays a critical role in interoceptive attention. Stephan KE, Frank MJ, Stephan KE, Faull OK, Allen M, editors. *eLife*. 2019 Apr 15; 8:e42265. <https://doi.org/10.7554/eLife.42265> PMID: 30985277
86. Barrett LF, Quigley KS, Hamilton P. An active inference theory of allostasis and interoception in depression. *Philos Trans R Soc B Biol Sci*. 2016 Nov 19; 371(1708):20160011. <https://doi.org/10.1098/rstb.2016.0011> PMID: 28080969
87. Jeganathan J, Breakspear M. An active inference perspective on the negative symptoms of schizophrenia. *Lancet Psychiatry*. 2021 Aug 1; 8(8):732–8. [https://doi.org/10.1016/S2215-0366\(20\)30527-7](https://doi.org/10.1016/S2215-0366(20)30527-7) PMID: 33865502
88. Moratti S, Rubio G, Campo P, Keil A, Ortiz T. Hypofunction of Right Temporoparietal Cortex During Emotional Arousal in Depression. *Arch Gen Psychiatry*. 2008 May 1; 65(5):532–41. <https://doi.org/10.1001/archpsyc.65.5.532> PMID: 18458205
89. Paulus MP, Stein MB. Interoception in anxiety and depression. *Brain Struct Funct*. 2010 Jun 1; 214(5):451–63. <https://doi.org/10.1007/s00429-010-0258-9> PMID: 20490545
90. Peters A, McEwen BS, Friston K. Uncertainty and stress: Why it causes diseases and how it is mastered by the brain. *Prog Neurobiol*. 2017 Sep 1; 156:164–88. <https://doi.org/10.1016/j.pneurobio.2017.05.004> PMID: 28576664
91. Legrand N, Nikolova N, Correa C, Brændholt M, Stuckert A, Kildahl N, et al. The heart rate discrimination task: a psychophysical method to estimate the accuracy and precision of interoceptive beliefs. *Biol Psychol*. 2021 Dec 11; 108239. <https://doi.org/10.1016/j.biopsycho.2021.108239> PMID: 34902450
92. Nikolova N, Harrison O, Toohey S, Brændholt M, Legrand N, Correa C, et al. The respiratory resistance sensitivity task: An automated method for quantifying respiratory interoception and metacognition. *Biol Psychol*. 2022 Apr 1; 170:108325. <https://doi.org/10.1016/j.biopsycho.2022.108325> PMID: 35358604
93. Wolpert N, Tallon-Baudry C. Coupling between the phase of a neural oscillation or bodily rhythm with behavior: Evaluation of different statistical procedures. *NeuroImage*. 2021 Aug 1; 236:118050. <https://doi.org/10.1016/j.neuroimage.2021.118050> PMID: 33848619
94. Parr T, Friston KJ. The Computational Anatomy of Visual Neglect. *Cereb Cortex*. 2018 Feb 1; 28(2):777–90. <https://doi.org/10.1093/cercor/bhx316> PMID: 29190328
95. Boldt A, de Gardelle V, Yeung N. The impact of evidence reliability on sensitivity and bias in decision confidence. *J Exp Psychol Hum Percept Perform*. 2017; 43(8):1520–31. <https://doi.org/10.1037/xhp0000404> PMID: 28383959
96. Spence ML, Dux PE, Arnold DH. Computations underlying confidence in visual perception. *J Exp Psychol Hum Percept Perform*. 2016; 42(5):671–82. <https://doi.org/10.1037/xhp0000179> PMID: 26594876
97. Koch E. Die irradiation der pressoreceptorischen kreislaufreflexe. *J Mol Med*. 1932; 11(6):225–7.
98. Delfini LF, Campos JJ. Signal Detection and the “Cardiac Arousal Cycle.” *Psychophysiology*. 1972; 9(5):484–91.
99. Edwards L, Ring C, McIntyre D, Winer JB, Martin U. Sensory detection thresholds are modulated across the cardiac cycle: Evidence that cutaneous sensibility is greatest for systolic stimulation. *Psychophysiology*. 2009; 46(2):252–6. <https://doi.org/10.1111/j.1469-8986.2008.00769.x> PMID: 19207192
100. Elliott R. The significance of heart rate for behavior: A critique of Lacey’s hypothesis. 1972;
101. Ghione S. Hypertension-associated hypalgesia: Evidence in experimental animals and humans, pathophysiological mechanisms, and potential clinical consequences. *Hypertension*. 1996; 28(3):494–504.
102. Sandman CA, McCanne TR, Kaiser DN, Diamond B. Heart rate and cardiac phase influences on visual perception. *J Comp Physiol Psychol*. 1977; 91(1):189. <https://doi.org/10.1037/h0077302> PMID: 838914
103. Saxon SA. Detection of near threshold signals during four phases of cardiac cycle. *Ala J Med Sci*. 1970; 7(4):427. PMID: 5532273
104. Birn RM. The role of physiological noise in resting-state functional connectivity. *Neuroimage*. 2012; 62(2):864–70. <https://doi.org/10.1016/j.neuroimage.2012.01.016> PMID: 22245341

105. Golanov EV, Yamamoto S, Reis DJ. Spontaneous waves of cerebral blood flow associated with a pattern of electrocortical activity. *Am J Physiol-Regul Integr Comp Physiol*. 1994 Jan 1; 266(1):R204–14. <https://doi.org/10.1152/ajpregu.1994.266.1.R204> PMID: 8304543
106. Zanatta P, Toffolo GM, Sartori E, Bet A, Baldanzi F, Agarwal N, et al. The human brain pacemaker: Synchronized infra-slow neurovascular coupling in patients undergoing non-pulsatile cardiopulmonary bypass. *NeuroImage*. 2013 May 15; 72:10–9. <https://doi.org/10.1016/j.neuroimage.2013.01.033> PMID: 23357071
107. Fox MD, Snyder AZ, Vincent JL, Raichle ME. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*. 2007; 56(1):171–84. <https://doi.org/10.1016/j.neuron.2007.08.023> PMID: 17920023
108. Fox MD, Raichle ME. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat Rev Neurosci*. 2007; 8(9):700. <https://doi.org/10.1038/nrn2201> PMID: 17704812
109. Kiverstein J. Free Energy and the Self: An Ecological–Enactive Interpretation. *Topoi* [Internet]. 2018 Apr 12 [cited 2019 Mar 26]; Available from: <https://doi.org/10.1007/s11245-018-9561-5>
110. Limanowski J, Blankenburg F. Minimal self-models and the free energy principle. *Front Hum Neurosci* [Internet]. 2013 [cited 2021 Mar 26]; 7. Available from: https://www.frontiersin.org/articles/10.3389/fnhum.2013.00547/full?utm_source=newsletter&utm_medium=email&utm_campaign=Neurology-w40-2013 PMID: 24062658
111. Mowrer OH. Two-factor learning theory: summary and comment. *Psychol Rev*. 1951; 58(5):350–4. <https://doi.org/10.1037/h0058956> PMID: 14883248
112. Friston K, Lin M, Frith CD, Pezzulo G, Hobson JA, Ondobaka S. Active Inference, Curiosity and Insight. *Neural Comput*. 2017 Aug 4; 29(10):2633–83. https://doi.org/10.1162/neco_a_00999 PMID: 28777724
113. Parr T, Friston KJ. Uncertainty, epistemics and active inference. *J R Soc Interface*. 2017 Nov 1; 14(136):20170376. <https://doi.org/10.1098/rsif.2017.0376> PMID: 29167370
114. Friston KJ, Rosch R, Parr T, Price C, Bowman H. Deep temporal models and active inference. *Neurosci Biobehav Rev*. 2018 Jul 1; 90:486–501. <https://doi.org/10.1016/j.neubiorev.2018.04.004> PMID: 29747865
115. Mifflin SW, Felder RB. Synaptic mechanisms regulating cardiovascular afferent inputs to solitary tract nucleus. *Am J Physiol-Heart Circ Physiol*. 1990 Sep 1; 259(3):H653–61. <https://doi.org/10.1152/ajpheart.1990.259.3.H653> PMID: 2204275
116. Miura M, Reis DJ. The role of the solitary and paramedian reticular nuclei in mediating cardiovascular reflex responses from carotid baro- and chemoreceptors. *J Physiol*. 1972; 223(2):525–48. <https://doi.org/10.1113/jphysiol.1972.sp009861> PMID: 5039285
117. Cechetto DF, Saper CB. Evidence for a viscerotopic sensory representation in the cortex and thalamus in the rat. *J Comp Neurol*. 1987; 262(1):27–45. <https://doi.org/10.1002/cne.902620104> PMID: 2442207
118. Craig AD. How do you feel? Interoception: the sense of the physiological condition of the body. *Nat Rev Neurosci*. 2002 Aug; 3(8):655–66. <https://doi.org/10.1038/nrn894> PMID: 12154366
119. Chouchou F, Mauguière F, Vallayer O, Catenoix H, Isnard J, Montavont A, et al. How the insula speaks to the heart: Cardiac responses to insular stimulation in humans. *Hum Brain Mapp*. 2019 Feb 28;
120. Oppenheimer SM, Gelb A, Girvin JP, Hachinski VC. Cardiovascular effects of human insular cortex stimulation. *Neurology*. 1992 Sep 1; 42(9):1727. <https://doi.org/10.1212/wnl.42.9.1727> PMID: 1513461
121. Hassanpour MS, Yan L, Wang DJJ, Lapidus RC, Arevian AC, Simmons W. Kyle, et al. How the heart speaks to the brain: neural activity during cardiorespiratory interoceptive stimulation. *Philos Trans R Soc B Biol Sci*. 2016 Nov 19; 371(1708):20160017.
122. Khalsa SS, Rudrauf D, Sandesara C, Olshansky B, Tranel D. Bolus isoproterenol infusions provide a reliable method for assessing interoceptive awareness. *Int J Psychophysiol*. 2009 Apr 1; 72(1):34–45. <https://doi.org/10.1016/j.ijpsycho.2008.08.010> PMID: 18854201
123. Anderson AK, Phelps EA. Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature*. 2001 May; 411(6835):305–9. <https://doi.org/10.1038/35077083> PMID: 11357132
124. Liddell BJ, Brown KJ, Kemp AH, Barton MJ, Das P, Peduto A, et al. A direct brainstem–amygdala–cortical ‘alarm’ system for subliminal signals of fear. *NeuroImage*. 2005 Jan 1; 24(1):235–43. <https://doi.org/10.1016/j.neuroimage.2004.08.016> PMID: 15588615

125. Fardo F, Auksztulewicz R, Allen M, Dietz MJ, Roepstorff A, Friston KJ. Expectation violation and attention to pain jointly modulate neural gain in somatosensory cortex. *NeuroImage*. 2017 Jun 1; 153:109–21. <https://doi.org/10.1016/j.neuroimage.2017.03.041> PMID: 28341164
126. Feldman H, Friston K. Attention, Uncertainty, and Free-Energy. *Front Hum Neurosci* [Internet]. 2010 [cited 2019 Mar 16]; 4. Available from: <https://www.frontiersin.org/articles/10.3389/fnhum.2010.00215/full>
127. Allman JM, Tetreault NA, Hakeem AY, Manaye KF, Semendeferi K, Erwin JM, et al. The von Economo neurons in the frontoinsula and anterior cingulate cortex. *Ann N Y Acad Sci*. 2011; 1225(1):59–71. <https://doi.org/10.1111/j.1749-6632.2011.06011.x> PMID: 21534993
128. Allen M, Glen JC, Müllensiefen D, Schwarzkopf DS, Fardo F, Frank D, et al. Metacognitive ability correlates with hippocampal and prefrontal microstructure. *NeuroImage*. 2017 Apr 1; 149:415–23. <https://doi.org/10.1016/j.neuroimage.2017.02.008> PMID: 28179164
129. Allen M, Fardo F, Dietz MJ, Hillebrandt H, Friston KJ, Rees G, et al. Anterior insula coordinates hierarchical processing of tactile mismatch responses. *NeuroImage*. 2016 Feb 15; 127:34–43. <https://doi.org/10.1016/j.neuroimage.2015.11.030> PMID: 26584870
130. Fleming SM, Dolan RJ. The neural basis of metacognitive ability. *Philos Trans R Soc B Biol Sci*. 2012 May 19; 367(1594):1338–49. <https://doi.org/10.1098/rstb.2011.0417> PMID: 22492751
131. Menon V, Uddin LQ. Saliency, switching, attention and control: a network model of insula function. *Brain Struct Funct*. 2010 Jun 1; 214(5):655–67. <https://doi.org/10.1007/s00429-010-0262-0> PMID: 20512370
132. Ullsperger M, Harsay HA, Wessel JR, Ridderinkhof KR. Conscious perception of errors and its relation to the anterior insula. *Brain Struct Funct*. 2010 Jun 1; 214(5):629–43.
133. Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD. Large-scale automated synthesis of human functional neuroimaging data. *Nat Methods*. 2011 Aug; 8(8):665–70. <https://doi.org/10.1038/nmeth.1635> PMID: 21706013
134. Chang LJ, Yarkoni T, Khaw MW, Sanfey AG. Decoding the Role of the Insula in Human Cognition: Functional Parcellation and Large-Scale Reverse Inference. *Cereb Cortex*. 2013 Mar 1; 23(3):739–49. <https://doi.org/10.1093/cercor/bhs065> PMID: 22437053
135. Schwartenbeck P, Friston K. Computational Phenotyping in Psychiatry: A Worked Example. *eNeuro*. 2016 Jul 1; 3(4):ENEURO.0049-16.2016. <https://doi.org/10.1523/ENEURO.0049-16.2016> PMID: 27517087
136. Kaplan R, Friston KJ. Planning and navigation as active inference. *Biol Cybern*. 2018 Aug; 112(4):323–43. <https://doi.org/10.1007/s00422-018-0753-2> PMID: 29572721
137. Lockmann ALV, Tort ABL. Nasal respiration entrains delta-frequency oscillations in the prefrontal cortex and hippocampus of rodents. *Brain Struct Funct*. 2018 Jan; 223(1):1–3. <https://doi.org/10.1007/s00429-017-1573-1> PMID: 29222724
138. Lockmann ALV, Laplagne DA, Tort ABL. Olfactory bulb drives respiration-coupled beta oscillations in the rat hippocampus. *Eur J Neurosci*. 2018; 48(8):2663–73. <https://doi.org/10.1111/ejn.13665> PMID: 28833629
139. FitzGerald THB, Dolan RJ, Friston K. Dopamine, reward learning, and active inference. *Front Comput Neurosci*. 2015; 9:136. <https://doi.org/10.3389/fncom.2015.00136> PMID: 26581305
140. Marshall AC, Gentsch A, Blum AL, Broering C, Schütz-Bosbach S. I feel what I do: Relating interoceptive processes and reward-related behavior. *NeuroImage*. 2019 May 1; 191:315–24. <https://doi.org/10.1016/j.neuroimage.2019.02.032> PMID: 30776528