

## Inference on homeostatic belief precision

Ozan Unal<sup>a,c,\*</sup>, Orhun Caner Eren<sup>a</sup>, Göktuğ Alkan<sup>a</sup>, Frederike Hermi Petzschner<sup>a</sup>, Yu Yao<sup>a</sup>,  
Klaas Enno Stephan<sup>a,b</sup>

<sup>a</sup> Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich & ETH Zurich, 8032 Zurich, Switzerland

<sup>b</sup> Max Planck Institute for Metabolism Research, 50931 Cologne, Germany

<sup>c</sup> Computer Vision Lab (CVL), ETH Zurich, 8092 Zürich, Switzerland

### ARTICLE INFO

#### Keywords:

Interoception  
Homeostasis  
Allostasis  
Allostatic self-efficacy  
Translational neuro-modeling  
Computational psychosomatics

### ABSTRACT

Interoception and homeostatic/allostatic control are intertwined branches of closed-loop brain-body interactions (BBI). Given their importance in mental and psychosomatic disorders, establishing computational assays of BBI represents a clinically important but methodologically challenging endeavor. This technical note presents a novel approach, derived from a generic computational model of homeostatic/allostatic control that underpins (meta) cognitive theories of affective and psychosomatic disorders. This model views homeostatic setpoints as probability distributions (“homeostatic beliefs”) whose parameters determine regulatory efforts and change dynamically under allostatic predictions. In particular, changes in homeostatic belief precision, triggered by anticipated threats to homeostasis, are thought to alter cerebral regulation of bodily states. Here, we present statistical procedures for inferring homeostatic belief precision from measured bodily states and/or regulatory (action) signals. We analyze the inference problem, derive two alternative estimators of homeostatic belief precision, and apply our method to simulated data. Our proposed approach may prove useful for assessing BBI in individual subjects.

### 1. Introduction

Biological organisms face the fundamental challenge of ensuring homeostasis: their physical states have to be kept within certain ranges in order to avoid violating boundary conditions of life. In a dynamically changing world, this requires adaptive behavior, driven by computations which take into account current and future states of both the body and the external environment. To achieve this, brain-body interactions (BBI) are essential. BBI consist of complex mechanisms for conveying information about numerous dimensions of bodily state (e.g., temperature, osmolality and oxygenation of blood, levels of metabolically and immunologically important molecules; for reviews, see (Critchley & Harrison, 2013; Khalsa et al., 2018)) from the body to the brain, and using this information to deploy regulatory mechanisms in a reactive manner (homeostatic control) or in an anticipatory fashion (allostatic control; Sterling, 2012).

It has been suggested that deciphering the mechanisms of BBI is not only important for understanding the physiological aspects of homeostatic/allostatic regulation, but that human cognition is grounded and constrained by the fundamental necessity of ensuring bodily

homeostasis (Stephan et al., 2016; Turner, 2019; Ziemke, 2016). For this reason, BBI play a key role in theories of mental and psychosomatic disorders (e.g. Badoud & Tsakiris, 2017; Barca & Pezzulo, 2020; Bonaz et al., 2021; Garfinkel et al., 2016; Henningsen et al., 2018; Khalsa et al., 2018; Paulus, Feinstein, & Khalsa, 2019; Petzschner, Weber, Gard, & Stephan, 2017; Quadt, Critchley, & Garfinkel, 2018; Schulz et al., 2020; Smith, Badcock, & Friston, 2021; Stephan et al., 2016). (Here, we refer to “psychosomatic disorders” as including both the perception of aversive bodily states in the absence of objectifiable pathology and the occurrence of somatic disorders due to unsuitable descending regulatory signals from the brain.) From a clinical perspective, it would therefore be highly desirable to have methods that can assess characteristics of BBI in human patients, in a non-invasive and quantitative fashion, as a basis for differentiating between alternative disease mechanisms (differential diagnosis). Additionally, it is possible that the same interoceptive dysfunction may be present in multiple clinical conditions (Smith et al., 2020).

A major challenge is that most (albeit not all) instantiations of BBI are organized as closed loops where homeostatic/allostatic actions change the state of bodily variables whose changes triggered regulation

\* Corresponding author at: Computer Vision Lab (CVL), ETH Zurich, 8092 Zürich, Switzerland.

E-mail address: [ozan.unal@vision.ee.ethz.ch](mailto:ozan.unal@vision.ee.ethz.ch) (O. Unal).

<https://doi.org/10.1016/j.biopsycho.2021.108190>

Received 7 February 2021; Received in revised form 7 September 2021; Accepted 10 September 2021

Available online 20 September 2021

0301-0511/© 2021 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

in the first place. This means that, in disease, pathological changes can propagate within the closed loops between body and brain, and it is difficult to distinguish between primary and secondary pathological changes. One potential approach to address this problem are “computational assays”: these are generative models of measured (physiological and/or cognitive) data which can be inverted to infer on likely mechanisms (Stephan & Mathys, 2014).

Computational assays of BBI do not exist yet. They differ from their counterparts in cognition or neurophysiology (Browning et al., 2020; Stephan, Iglesias, Heinze, & Diaconescu, 2015) in that they need to consider the closed-loop problem and account for both physiological (bodily) and neuronal/cognitive (brain) data. The development of suitably structured computational assays together with controlled perturbations of bodily states has been proposed as a major goal for computational psychosomatics (Petzschner et al., 2017). In this paper, we present an initial methodological step in this direction.

BBI comprise multiple processes (Petzschner et al., 2017; Petzschner, Garfinkel, Paulus, Koch, & Khalsa, 2021), including afferent signaling from the body (interoception), inference on the underlying bodily states (interoception), predicting their future evolution (forecasting), regulatory processes (homeostatic and allostatic control), and monitoring the performance of BBI loops (metacognition). In this paper, we do not consider all components of BBI but focus on the question how one might estimate a single, but fundamental, property of regulatory processes: the sensitivity of autonomic reflexes to (viscero)sensory inputs from the body. This sensitivity is conceptualized via the precision of probabilistic representations (homeostatic beliefs) at the lowest level of hierarchically arranged systems for allostatic control, as understood from a Bayesian (active inference) perspective.

In order to estimate homeostatic belief precision, we present a novel approach that is based on a generic framework of BBI and homeostatic/allostatic control, the “Bayesian Allostatis” (BA) model by Stephan et al. (2016). Clearly, this is not the only framework of BBI and homeostatic/allostatic control that could be used; in particular, other schemes based on active inference offer sophisticated proposals for studying BBI (Allen, Levy, Parr, & Friston, 2019; Gu, Hof, Friston, & Fan, 2013; Paulus et al., 2019; Pezzulo, Rigoli, & Friston, 2015; Seth, 2013; Seth & Friston, 2016; Smith et al., 2020). For our purposes, the advantage of the BA framework is that it provides a mathematically explicit proposal for the lowest level of a hierarchical control system, i.e. a reflex arc, while not requiring a specification of the processes that govern inference and forecasting processes at higher levels.

In brief, the BA model by Stephan et al. (2016) assumes that the brain represents homeostatic setpoints in a probabilistic fashion - i.e. as probability distributions or “homeostatic beliefs” about desired bodily states - and “defends” these beliefs against deviations (induced by perturbations that induce dyshomeostasis) by eliciting regulatory actions. Here, action selection follows the same principle as in active inference (Friston, Daunizeau, Kilner, & Kiebel, 2010): actions are chosen to fulfill beliefs (here, homeostatic beliefs about bodily states), thus minimizing the brain’s surprise (in an information-theoretic sense) about sensory inputs. In the context of bodily regulation, this principle minimises the variability of bodily states and keeps them close to homeostatic setpoints; in other words, the entropy of bodily states is minimized (for a more detailed discussion, see Stephan et al., 2016). For bodily regulation, actions could comprise, for example, descending neural signals to the autonomic nervous system or the emission of hormones via hypothalamus and the pituitary gland.

Unlike active inference, the BA model does not specify the algorithmic nature of allostatic control but describes how the outcomes of (unspecified) allostatic computations lead to actions: via a change of the parameters (sufficient statistics) of homeostatic beliefs represented in autonomic effector regions that govern reflex arcs for bodily regulation. For example, given a forecasting mechanism that identifies future perturbations of homeostasis and threats of bodily integrity, changing the expectation (mean) of homeostatic beliefs triggers anticipatory

regulatory actions automatically. Furthermore, and of particular importance for this paper, changing the precision of homeostatic beliefs (i.e., the inverse variance of the desired range of bodily states) determines how rapidly cerebral regulatory actions will kick in when sensory inputs indicate a deviation of actual from expected bodily states (prediction error). (For a visualization of these mechanisms, see Fig. 2 below.).

These concepts are clinically relevant and have informed theories of specific clinical symptoms (e.g. fatigue; Manjaly et al., 2019; Stephan et al., 2016), disorders (e.g. depression: “allostatic self-efficacy”; Stephan et al., 2016) as well as theories of psychotherapeutic interventions (e.g. mindfulness-based cognitive therapy, MBCT; Manjaly & Iglesias, 2020). More generally, it has been argued that methods for characterizing homeostatic beliefs in individual patients could be of major utility for differential diagnosis (Petzschner et al., 2017; Stephan et al., 2016).

Here, we focus on the specific question how one might estimate the precision of homeostatic beliefs - which, in the BA model, determine the sensitivity of reflex arcs and thus the deployment of regulatory actions - from measurements of bodily states and/or regulatory signals. Changes of the precision of homeostatic beliefs are particularly important in situations where a perturbation or threat of bodily integrity can be predicted but not its exact nature or direction. In this situation, a targeted anticipatory action is precluded: in the absence of predictions about what exactly will happen, any anticipatory action that shifts setpoints could make matters worse. However, a remaining allostatic control option is to change the precision of homeostatic beliefs, which alters the reactivity of reflex arcs and thus the emission of regulatory responses in response to sensory inputs from the body. Notably, in this paper, we are not concerned with the characterization of anticipatory allostatic processes per se. Instead, we propose ways of estimating changes of homeostatic belief precision that transform such anticipatory processes into actions, by altering the sensitivity of autonomic reflexes to interoceptive signals.

The paper is structured as follows. We first present a brief mathematical summary of the BA model by Stephan et al. (2016) which represents the backbone of our approach. We then analyze the problem of inference in the context of this model. Based on this analysis, we introduce two methods for estimating homeostatic belief precision from measurements of both bodily states and regulatory signals. (In the [Supplementary Material](#), we present full derivations and variations of our approach that only require information about bodily states or regulatory signals.) Finally, we illustrate the performance of our methods using simulated data.

## 2. Methods

### 2.1. The Bayesian Allostatis (BA) model

The approach presented in this paper is based on a probabilistic model of BBI and allostasis that was previously introduced in the context of the “allostatic self-efficacy” theory of fatigue and depression (Stephan et al., 2016). Here, we refer to this model as the “Bayesian Allostatis” (BA) model. The BA model is based on the fundamental principle of homeostasis: living systems interact with their environment in order to keep their physical states within certain ranges that do not violate biophysical boundary conditions of life (e.g. conditions that would prevent certain chemical reactions). In other words, living systems attempt to minimize the entropy of their internal states, at the expense of increased entropy outside the system (Schrödinger, 1943).

The key difference of the BA model to classical textbook models of homeostatic control is that setpoints are neither assumed to be fixed entities nor that they are deterministic. Instead, in accordance with the general notion of the “Bayesian brain” (Friston, 2005; Knill & Pouget, 2004), the BA model assumes that setpoints (1) are represented by the brain in a probabilistic fashion and (2) that setpoints can change dynamically. In other words, the BA model views setpoints as potentially

time-varying probability distributions over bodily states. From the cognitive neuroscience perspective, these distributions can be thought of as beliefs, held by the brain, about the states which the body should inhabit or, equivalently, as hypotheses about the states the body should be in. In this setting, minimizing the entropy of bodily states (by choosing appropriate regulatory actions) corresponds to minimizing interoceptive surprise (i.e. information-theoretic (Shannon) surprise about sensory inputs that the brain receives about bodily states) or, equivalently, maximizing the evidence for the hypothesis of being in homeostasis. Generally, the BA model is conceptually related to (and partially inspired by) the framework of active inference (Adams, Shipp, & Friston, 2013) but is (deliberately) less specific and differs in a number of aspects. For example, it does not specify how exactly perceptual inference and allostatic predictions take place, and it does not make any assumptions about specific approximations to surprise (such as the free energy approximation in active inference; Friston, 2009). Instead, it focuses on a generic description of a modular cognitive architecture that describes how homeostatic reflex arcs are controlled by allostatic predictions, as explained in the following.

The basic architecture of the BA model consists of 2 intertwined loops, a lower and a higher level (see Fig. 1). (Strictly speaking, the model includes a third, metacognitive, level which, however, is not of relevance for this paper; see Stephan et al. (2016) for details.) The lower level consists of a closed-loop reflex arc between a given bodily state  $x$  and a brain structure that (1) receives sensory inputs ( $y$ ) about this bodily state, (2) compares it to a probabilistic setpoint (homeostatic belief), resulting in a precision-weighted prediction error (pwPE) that determine the emission of (3) descending regulatory signals, with the goal of minimizing interoceptive surprise or, put differently, that future sensory inputs become more similar to what is expected under the homeostatic belief. The higher level, which receives the same sensory inputs, is thought to perform inference (about current bodily states) and forecasting (about future bodily states) as a basis for anticipatory control signals that change the homeostatic belief at the lower level and thus alters the reactivity of the reflex arc. Notably, the existing formulation of the BA model only provides a mathematical formulation of the lower level and leaves the exact implementation of the higher level open. This modular structure is useful for the purpose of this paper - to develop estimators of the precision of homeostatic beliefs—since it allows us to focus on the simpler lower level (which can be modeled with relatively few assumptions) without having to make any assumptions about the operations and mathematical formulations of the more complex higher level.

The lower level of the BA model consists of three components: a likelihood function that describes how bodily states lead to noisy sensory inputs received by the brain, a prior distribution over bodily states (homeostatic beliefs), and an “action function” describing how regulatory actions change bodily states. (For simplicity, in the following, we will consider the case of a single bodily state.) Importantly, the form of these functions in the BA is fairly general, and relatively few assumptions are necessary. Specifically, the assumptions made are that (1) homeostatic setpoints are represented probabilistically in the brain as probability distributions, (2) these probability distributions as well as

sensory noise have a Gaussian form, and (3) the action function takes the form of a gradient descent. We emphasize that the statistical approach presented below does not require any additional assumptions and holds regardless of the choice of implementation details, such as the choice of observation functions, time constants, etc. In the following, we briefly summarize the key components of the BA model.

In the BA model, the homeostatic belief (prior distribution over bodily state) is assumed to be represented by a Gaussian distribution with mean  $\mu_{prior}$  and variance (inverse precision)  $\pi_{prior}^{-1}$ :

$$p(x) = \mathcal{N}(x; \mu_{prior}, \pi_{prior}^{-1}) \quad (1)$$

Through afferent channels, the brain receives a noisy measurement of the bodily states  $x$ , i.e. the viscerosensory input  $y$ :

$$y_t = g(x_t) + \eta_{d,t} \quad \text{with} \quad \eta_{d,t} \sim \mathcal{N}(0, \pi_d^{-1}) \quad (2)$$

Here,  $g$  denotes the mapping from bodily states to sensory inputs which the brain receives (i.e. a signal-generating or forward function, part of the likelihood in the brain’s generative model; compare Fig. 1) and  $\eta_{d,t}$  representing sensory channel noise (white Gaussian noise with variance  $\pi_d^{-1}$ ). To maintain homeostasis, the brain deploys regulatory processes in a reactive manner in forms of action sequences  $a$ , which aim to reduce any discrepancy between the actual sensory input and the input that is expected under the homeostatic belief. The action corresponds to a gradient descent on interoceptive (Shannon) surprise and depends on the precision-weighted prediction error, i.e.  $\pi_t(y_t - g(\mu_{prior}))$ :

$$a_t = -\pi_t(y_t - g(\mu_{prior})) \partial_x g \quad (3)$$

leading to a change of bodily state  $x$ :


$$\partial_t x = \lambda^{-1} f(a(t)) \quad (4)$$

In Eq. (4), for generality,  $\lambda$  represents the time constant of the specific regulatory action considered and  $f$  denotes a mapping from regulatory action to changes in  $x$ ; please see Stephan et al. (2016) for details. Importantly, for the method of estimating homeostatic belief precision we propose below, these quantities do not need to be known as long as measurements of the actions are available.

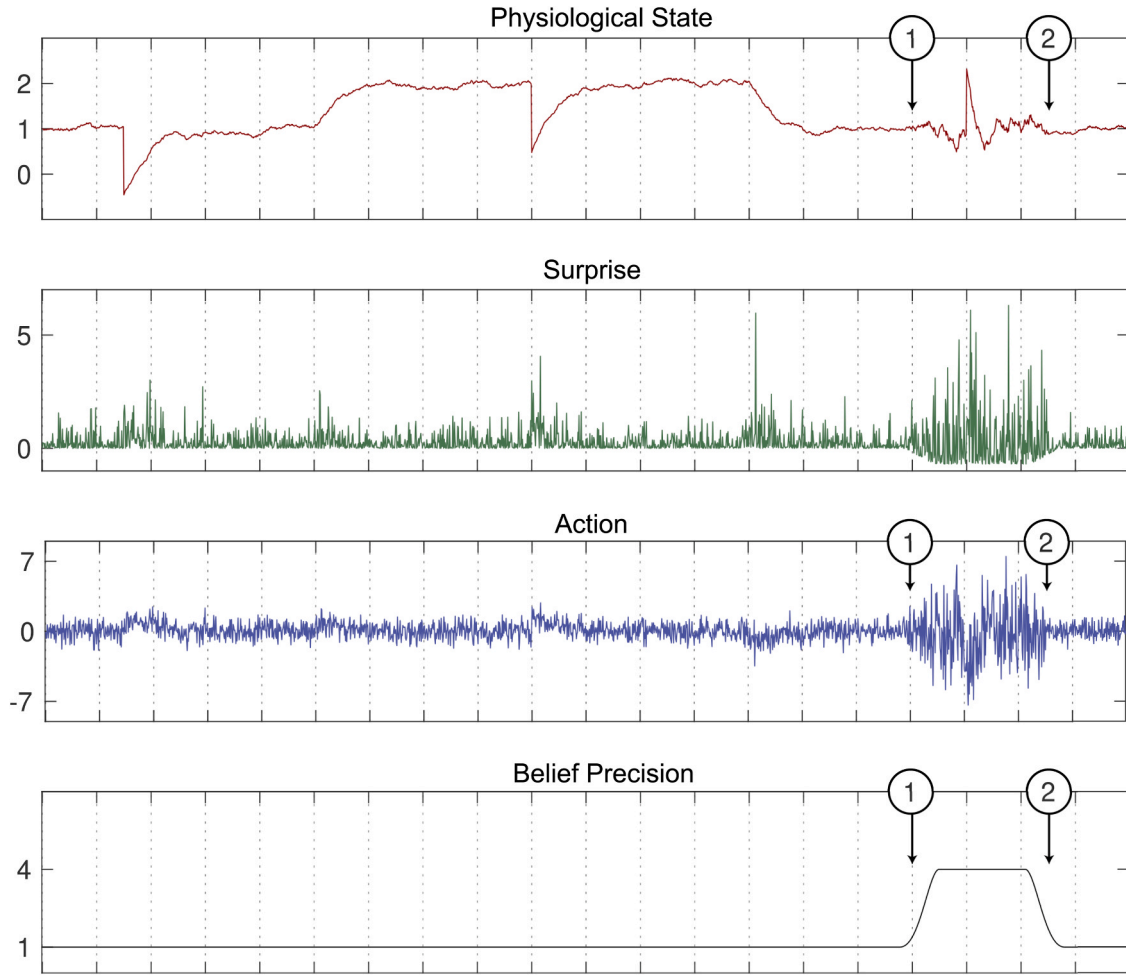
In this work, we consider the prior belief precision  $\pi_{prior}$  in Eq. (1) to be a dynamic quantity (that changes in time according to allostatic predictions) and denote it as  $\pi_t$  referencing the discrete time point  $t$ . A visualization of the model’s behavior can be seen in Fig. 2. The goal of this paper is to establish statistical methods for inferring the time-varying belief precision based on measurements of actions  $a$  and/or states  $x$ .

## 2.2. Analyzing the problem of inferring belief precision in the BA model

Writing the BA model as a Bayesian network results in the structure displayed by Fig. 3. The belief precision  $\pi_t$ , action  $a_t$  and viscerosensory input  $y_t$  form a  $v$ -structure which is encoded as  $\pi_t \perp\!\!\!\perp y_t | a_t$ . In other words, the belief precision  $\pi_t$  and the sensory input  $y_t$  are not



**Fig. 1.** The general structure of the Bayesian Allostater (BA) model which provides the foundation for the approach described in this paper. For explanations, please see main text. The figure is adapted from Fig. 5 in Stephan et al. (2016) under the Creative Commons Attribution License (CC BY).



**Fig. 2.** An extension to the simulated example of allostatic regulation of homeostatic control (with time-varying homeostatic belief precision) from Stephan et al. (2016). The x-axis denotes time, the y-axis arbitrary units. The upper panel shows an arbitrary physiological state  $x$ . The second and third panels show an approximation to interoceptive surprise (in terms of squared precision-weighted prediction error, *pwPE2*) and the associated action signal  $a$ , respectively. The bottom panel displays the time-varying belief precision. The initial two-thirds of these panels show how external perturbations and their anticipation lead to both reactive and anticipatory actions, accompanied by changes in interoceptive surprise (see Stephan et al. (2016) for details). For this paper, the period between ① and ② is important: here, homeostatic belief precision increases in response to an anticipated perturbation whose nature (direction) is unknown. It can be seen that once homeostatic belief precision increases, both interoceptive surprise and action signals increase drastically, reflecting the heightened reactivity of the reflex arc and its attempt to control very small fluctuations in sensory input. Prior to a real perturbation, this amounts to an attempt of controlling sensory noise. On the other hand, once the perturbation occurs (spike between ① and ② in upper panel), control is exerted much more swiftly compared to previous perturbations during the simulation. Technical details: In the bottom panel, between time points ① and ②, belief precision rises and falls according to a biased second order damped harmonic oscillator given by  $h(z) = \frac{d^2z}{dt^2} + 2\gamma \frac{dz}{dt} + \omega^2 z + b$ , with  $b$  denoting a bias term that centers the function around a given time point. The harmonic function is modified such that  $h(z) = 0$  for all terms before and after the minima and maxima to preserve the continuity of its derivative. For the simulations  $\pi_d = 2$ ,  $\pi_a = 100$ ,  $\lambda = 50$ ,  $f = g = 1$  were chosen.

independent if the action sequence  $a_t$  is observed. Thus, observing the action states  $a_{t'}$  and the state variable  $x_{t'} \forall t' \in [1, \dots, t]$ , opens up the structure which enables the inference of the belief precision  $\pi_t$  from the action  $a_{t-1}$  via  $(a_{t-1} \rightarrow x_t \rightarrow y_t \rightarrow a_t \leftarrow \pi_t)$ .

We start by computing the statistics of the measured noisy action  $a_t^m$  (denoted by the superscript  $m$ ) as follows:

$$\mathbb{E}[a_t^m | a_{1:t-1}^m; \pi_t] = -\pi_t (g(x_t) - g(\mu_{\text{prior},t})) \partial_x g \quad (5)$$

$$\text{Var}[a_t^m | a_{1:t-1}^m; \pi_t] = \pi_t^2 \phi^{-1} + \pi_a^{-1} \quad (6)$$

For a detailed derivation, please see Appendix A.1.

Solving Eq. (5) for the precision  $\pi$  results in an ill-posed problem statement as  $\mathbb{E}[a_t^m | a_{1:t-1}^m; \pi_t] / ((g(x_t) - g(\mu_{\text{prior},t})) \partial_x g) \rightarrow \infty$ , meaning the belief precision is not recoverable from the expectation of the action sequence under the condition that  $\mathbb{E}[x] \rightarrow \mu_{\text{prior}}$ .


On the other hand, the variance of the action is determined by the

scaled precision of the sensory data noise  $\phi$  and the belief precision squared. To infer the belief precision  $\pi_t$ , we propose a *Standard Deviation based Precision Estimation* (StdPE) method that utilizes Eq. (6) and is described in the next section.

### 2.3. Standard deviation based precision estimation (StdPE)

Assuming that (1) the regulatory actions are sampled from a differentiable continuous function with its derivative being bound by some constant and that (2) multiple measurements of actions can be obtained to increase the measurement precision  $\pi_a \rightarrow \infty$ , the belief precision  $\pi_t$  can be estimated up to a scale factor  $\sqrt{\phi}$  (see Appendix A.2 for mathematical details). Notably, this only requires measurements of actions  $a$  - without any knowledge or assumptions about other variables like the sensory data precision  $\pi_d$ , the time constant  $\lambda$ , the mapping  $g$  from bodily states to sensory inputs, or the function  $f$  determining the influence of actions





**Fig. 3.** The Bayesian Allostat (BA) model described in Stephan et al. (2016) displayed as a Bayesian network.  $x$ : bodily state;  $y$ : sensory input received by the brain;  $a$ : regulatory action;  $\pi$ : homeostatic belief precision;  $\lambda$ : time constant of action;  $t$ : time index.

on bodily states:

$$\pi_t \approx \sqrt{\phi} \cdot \sqrt{\text{Var}[a_{t-N+1:t}^m; \pi_t]} \quad (7)$$

Here,  $N$  is the size of the causal window, i.e. a window chosen such that it only contains the  $N$  data points from before the time point  $t$ . For the full derivation, please see Appendix A.2.

In essence, the precision is inferred by sliding a window over the data and computing the standard deviation of the windowed region. The standard deviation of the windowed action can be taken as an estimate for the scaled point-wise standard deviation of the last action point. It should be noted that the causal window we chose allows for the method to be performed in real time with the belief precision being inferred as new data arrives. When the method is applied for offline inference from previously acquired data, the window can be centered around the time frame  $t$  to reduce latency through sudden rises.

Although perhaps obvious, it is worth mentioning that the reliance of StdPE on variance in the data means that the approach would fail in the limit of  $\phi \rightarrow \infty$  (see Appendix A.1). However, this is not a practical concern because it would require that either the sensory signal received by the brain ( $\pi_d$ ) is completely free of noise or that the mapping  $g$  (from bodily states to sensory inputs) becomes extremely steep. Neither of these conditions can be plausibly expected to be fulfilled.

Eq. (7) provides an estimate of belief precision up to the scaling factor  $\sqrt{\phi}$ . While, under certain conditions, this is already useful for practical application (see Discussions for details), it is also possible, in principle, to obtain an exact estimate of belief precision. In the Supplementary Material we propose a method to estimate  $\phi$  from simultaneous measurements of  $x$  and  $y$  which enables exact estimates of the belief precision  $\pi_t$ .

#### 2.4. Log-evidence based precision estimation (LEPE)

A second strategy to infer the belief precision is to identify that value of belief precision which best explains the measured action sequence, in terms of maximizing its (log) evidence. In this section, we summarize the Log-Evidence Based Precision Estimation (LEPE) approach; the full derivation can be found in the Supplementary Material.

Assuming that the belief precision is constant within a given time window, the log-evidence can be computed from the joint probability distribution of all actions using the chain-rule. Thus, the belief precision  $\pi_t$  is inferred as the argument that maximizes the log-evidence. In other words, the estimated belief precision is taken from the most likely *a posteriori* distribution of the last action sequence.

It is worth emphasizing that while StdPE can be used to determine belief precision up to a scale factor without having to make any assumptions, for LEPE, one must know or assume sensory input precision

$\pi_d$  as well as the mapping  $g$  from bodily states  $x$  to the sensory inputs  $y$  which the brain receives (please see Eq. (2) and the term  $\phi$  in the derivation of LEPE in the Supplementary Material).

### 3. Results

In this section we provide simulations that demonstrate the two inference methods introduced in the Methods section, along with further investigations of the chosen parameters for the model and inference method. For simulations, we follow Stephan et al. (2016) and chose the state-input mapping ( $g$ ) and action-state mapping ( $f$ ) as identity functions, the time constant of the regulatory action  $\lambda = 50$ , the precision of the action measurement  $\pi_a = 100$  and the precision of the sensory channel  $\pi_d = 2$ . The window size for each method is chosen as  $N = 15$  time steps unless stated otherwise for a simulation with 2000 time steps. All simulations have been implemented in MATLAB R2019a (MathWorks Inc., Natwick, MA, USA). The code is publicly available at github.com/ouenal/inference-on-homeostatic-belief-precision.

#### 3.1. Results for StdPE

Fig. 4 shows the results of the StdPE method applied to the simulated case presented in Section 2.1.

In Fig. 4.a, the belief precision is estimated up to a scale factor with an unknown scale factor  $\phi$ . For an easier visual comparison, both the ground truth and the estimation are normalized onto the  $[0, 1]$  interval based on the means of their top-10 minimum and maximum values. As shown in this plot, StdPE can successfully infer the belief precision up to a scale factor.

The exact estimation of the belief precision  $\pi_t$  is shown in Fig. 4.b where we set the scale factor  $\phi$  to its simulated value. The results demonstrate that, in this specific simulation scenario, StdPE recovers the time-varying homeostatic belief precision accurately.

#### 3.2. Results for LEPE

Fig. 4 also shows the estimated belief precision for the simulated case using the LEPE method. As LEPE requires that the scale factor  $\phi$  is set to some value, in Fig. 4.a, we use an arbitrary guess ( $\phi = 1$ , %50 of its true value), and normalize the resulting belief precision estimation as in Section 3.1. In Fig. 4.b. we show the results where the scaled precision of the sensory data noise  $\phi$  is set to its true value ( $\phi = 2$ ). It can be seen that, in the chosen simulation scenario, LEPE infers the belief precision up to a scale factor when  $\phi$  is unknown (but a value is guessed), and to its exact value when  $\phi$  is known.

#### 3.3. A more challenging simulation

In this section we provide the graphical results for a more challenging simulation with more frequent precision changes while retaining the mean state prior  $\mu_{prior}$  curve and external perturbations to the system.

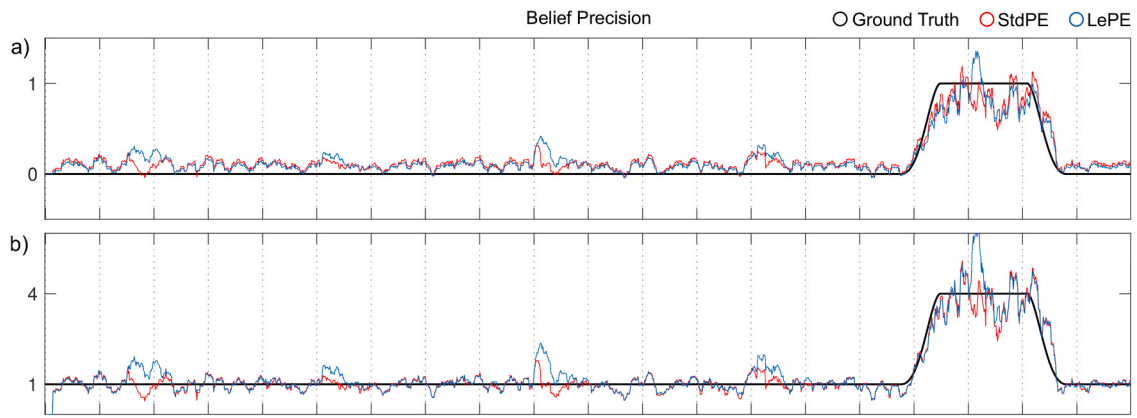
The simulated model is shown in Fig. 5 with the results for an unknown scale factor  $\phi$  in Fig. 6.a and a known  $\phi$  in Fig. 6.b.

#### 3.4. Quantitative comparison of StdPE and LEPE

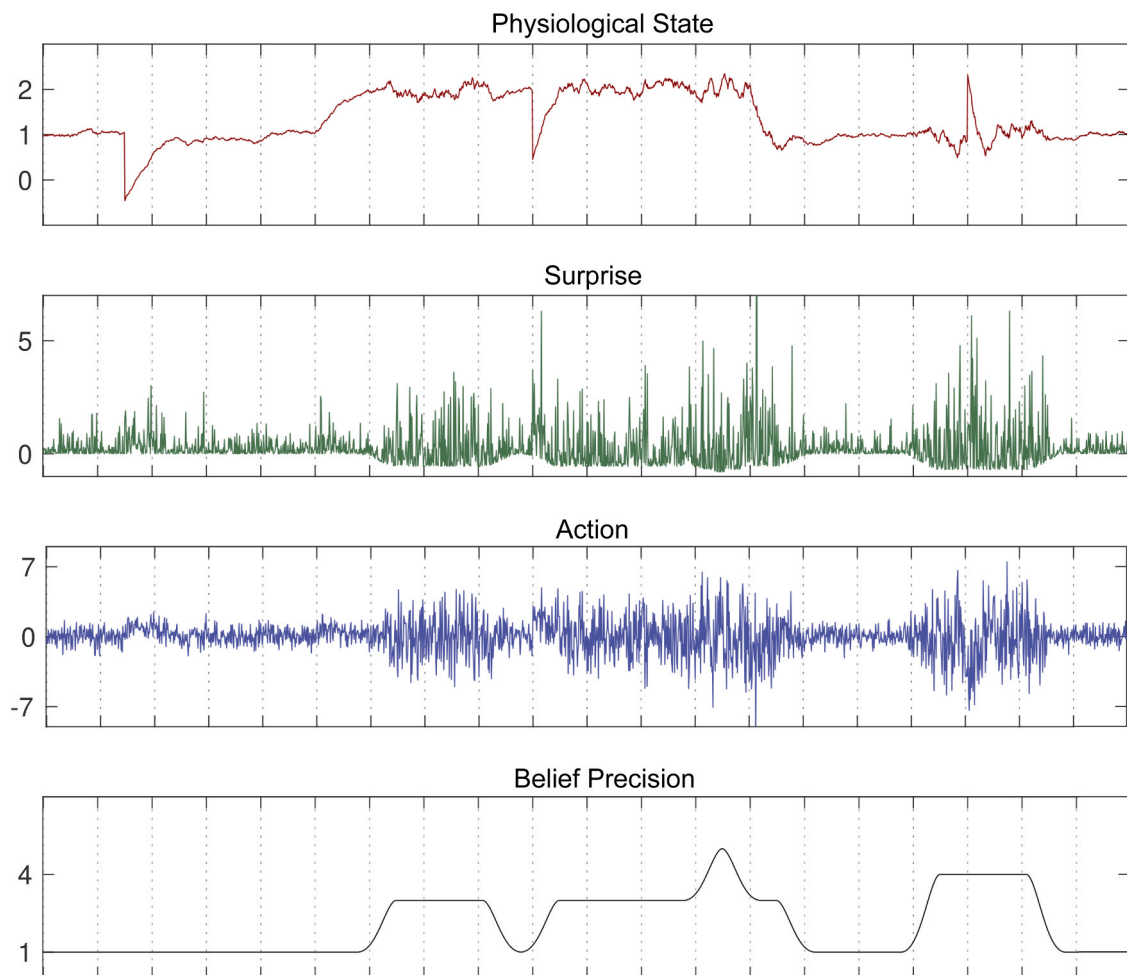
To quantitatively assess the quality of estimation, we compute the Mean Absolute Percentage Error (MAPE) of the estimated belief precision against its ground truth. The MAPE is given by:

$$\text{MAPE} = \frac{1}{T} \sum_{t=1:T} \left| \frac{\pi_t - \pi_t^{gt}}{\pi_t^{gt}} \right| \quad (8)$$

with the estimated belief precision  $\pi_t$ , ground truth belief precision  $\pi_t^{gt}$  and total time steps  $T$ . MAPE states the percentile error of each



**Fig. 4.** Belief precision inference results for StdPE and LEPE. Shown are the ground truth belief precision in black and the estimated belief precision in red for StdPE, and in blue for LEPE. Best viewed in color. (a) Scaled belief precision estimation. While StdPE can be used directly to estimate the belief precision up to a scale factor, LEPE requires a value to be guessed for  $\phi$ . In the given example this value is set to  $\phi = 1$  (simulation with  $\phi = \sqrt{2}$ ). For visual purposes the ground truth and resulting estimations are normalized onto the  $[0,1]$  interval based on the average of their maximum and minimum 10 values. (b) Exact belief precision estimation using StdPE and LEPE with  $\phi = \sqrt{2}$ .

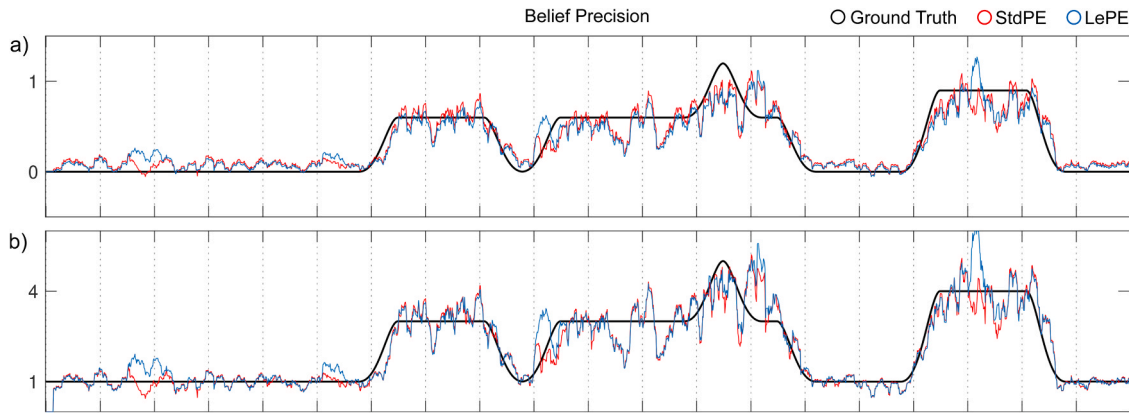


**Fig. 5.** A second simulated example of allostatic regulation of homeostatic control with more frequent belief precision changes than Fig. 2. For the simulations  $\pi_d = 2$ ,  $\pi_a = 100$ ,  $\lambda = 50$ ,  $f = g = 1$  were chosen.

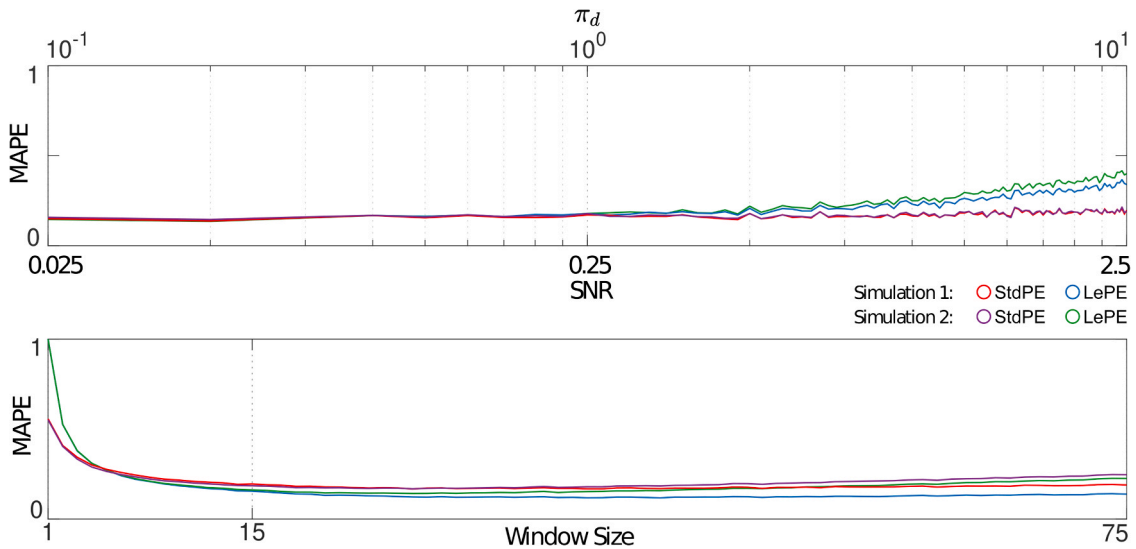
estimate compared to its ground truth, averaged across the entire sequence.

Fig. 7 (top panel) shows the MAPE of two simulations (Simulation 1: Section 2.1 and Simulation 2: Section 3.3) for both proposed methods with varying Signal-to-Noise-Ratios ( $SNR = \text{Var}[x]/\text{Var}[y] = \pi_d \text{Var}[x]$ ).

As seen, StdPE performs well for a wide range of SNR values, with  $MAPE < 18\%$  across all simulations. Similarly LEPE performs well for low SNR values of  $< 0.25$ , however, as SNR increases, so does the estimation error. As an inverse relation between estimation error and SNR is highly counterintuitive, it is important to clarify why this is observed.



**Fig. 6.** Belief precision inference results for StdPE and LEPE on Simulation 2 (Fig. 4). Shown are the ground truth belief precision in black and the estimated belief precision in red for StdPE, and in blue for LEPE. Best viewed in color. (a) Scaled belief precision estimation. While StdPE can be used directly to estimate the belief precision up to a scale factor, LEPE requires a value to be guessed for  $\phi$ . In the given example this value is set to  $\phi = 1$  (simulation with  $\phi = \sqrt{2}$ ). For visual purposes the ground truth and resulting estimations are normalized onto the [0,1] interval based on the average of their maximum and minimum 10 values. (b) Exact belief precision estimation using StdPE and LEPE with  $\phi = \sqrt{2}$ .



**Fig. 7.** Quantitative analysis of the inference methods based on the Mean Average Percentile Error (MAPE). Simulation 1 results (Fig. 2) for StdPE are shown in red and for LEPE are shown in blue. Simulation 2 results (Fig. 4) for StdPE are shown in purple and for LEPE in green. Top panel: MAPE visualized across varying Signal-to-Noise-Ratio (SNR) changes. SNR is computed as  $\text{SNR} = \pi_d \text{Var}[x]$ . Bottom panel: MAPE visualized across varying window sizes to show the influence of the window size. It can be seen that the window size  $N$  must be sufficiently large enough that the variance can be accurately estimated.

For the derivation of LEPE, it is assumed that through numerous measurements of the action  $a$ , the measurement noise is reduced sufficiently that it can be ignored. In the simulations, this is represented as  $\pi_a = 100$ . The simplification to a quadratic equation assumes that  $\phi\pi_a^{-1} \rightarrow 0$  and  $\phi^2\pi_a^{-1} \rightarrow 0$  hold. The second assumption, however, gets less appropriate if the sensory channel precision  $\pi_d$  increases to very high values, as the scale factor is given by  $\phi = \pi_d(\partial_x g)^2$ . In this case, the second assumption will, at some point, fail which then renders LEPE inaccurate.

We further show an additional investigation for the chosen window size  $N$  in Fig. 7 (bottom panel). As expected, a small window size fails to correctly identify the variance due to a lack of available data points. This is reflected by the initial rapid decrease in the MAPE as the window size increases. In contrast to this observation, a window size that is too large has the drawback of not being able to capture incoming precision changes. As the belief precision changes, the measured variance is affected by the previous measurements which can cause a latency in the estimation if the window size approaches the time scale of signal changes, or an underestimation of the belief precision if the signal length

is smaller than the window size.

#### 4. Discussion

In this technical note, we have introduced two statistical techniques for estimating the precision of homeostatic beliefs in the context of the Bayesian Allostat (BA) model by Stephan et al. (2016). We have illustrated the application of these methods to simulated data where ground truth (i.e. the actual temporal trajectory of belief precision) was known and examined their behavior across different levels of SNR. In the following, we discuss limitations and strengths of the current approach, how these could be applied to experimental measurements, and how it might be improved in the future.

First, we wish to emphasize that the present paper should be seen as an early - and still rather incomplete - step towards computational assays of BBI. This is true both conceptually - because we cover only a single aspect of BBI, i.e. homeostatic belief precision - and because our approach has only been tested on synthetic data so far, evaluating whether a known ground truth (the temporal trajectory of homeostatic

belief precision) can be recovered from simulated data. We would have liked to test our approach on empirical data but, unfortunately, do not yet have a suitable dataset that allows for a meaningful practical demonstration. This is because, as we explain below in more detail, applying our approach in practice usually requires measurements of regulatory (action) signals at a minimum. These measurements are not easy to obtain with standard methods (see below), should be acquired with a sufficient sampling rate (relative to the time constant of the reflex arc) and of sufficient duration to result in a larger number of data points. Furthermore, these data should be acquired under an experimental manipulation that elicits allostatic (i.e. anticipatory) control in response to an upcoming bodily perturbation of unknown direction whose predictability has a known start point and a known duration. The third condition is necessary to ensure one has well-defined predictions against which the method can be tested, i.e. a clearly defined temporal interval in which changes in homeostatic belief precision are expected. As explained in more detail below, while it is not impossible to apply the proposed approach when only measurements of bodily states are available, this would require additional experimental measurements in order to determine  $\lambda$  (the time constant of regulatory action) and  $f$  (the mapping from action to changes in bodily states).

The BA model and our statistical estimation approaches that rest on it are fairly general, and relatively few assumptions are necessary. However, one fundamental assumption of our model must be kept in mind: the notion that homeostatic setpoints are represented probabilistically in the brain. To our knowledge, this general assumption - which derives from the general "Bayesian brain" notion that the brain constructs representations of the world which reflect uncertainty (Knill & Pouget, 2004) - has not yet been examined in empirical (i.e. neurophysiological) studies. If this assumption turns out to be wrong and homeostatic setpoints are found to be represented by the brain as quantities without uncertainty, our model would lose its conceptual foundation.

Beyond the limitations discussed above, the methods we propose also have notable strengths: they are mathematically simple, extremely fast, and do not require complex optimization algorithms. Furthermore, the generic BA model they refer to is not limited to specific applications or particular process models of allostatic control, but applies to all types of BBI that represent a closed loop between sensory inputs and regulatory action. This generality may also create opportunities for testing predictions from specific process models such as active inference. As explained above, the BA model suggests two ways how allostatic predictions can be transformed into actions that regulate bodily states: (i) shifting the expectation (mean) of homeostatic beliefs, or (ii) changing their precision. The former case elicits anticipatory action while the latter changes the sensitivity of reflex arcs to sensory input from the body (i.e. the case considered by this paper). Changes in precision can also be detected by models based on active inference; for an example, see Smith et al. (2020). Predictions from a specific active inference model could thus potentially be cross-validated by the approach presented in this paper, for example, by comparing ratios in estimated precision before and after experimental interventions designed to elicit changes in the precision (not expectations) of homeostatic beliefs.

In practice, given the robust behavior of StdPE in our simulations across a wide range of SNR and since it requires less information, we recommend primarily using StdPE, instead of LEPE. In the following, we would like to describe how StdPE could be applied in practice, depending on which measurements are available. Generally, three different use cases can be distinguished:

(i) Measurements of regulatory (action) signals only are available:

Having measurements of action signals is sufficient to compute homeostatic belief precision up to a scaling factor  $\sqrt{\phi}$  (see Eq. (7)). This scaling factor depends on quantities that can be assumed to be stable within an individual, at least under the same measurement conditions (i.e. the precision of sensory inputs from

the body,  $\pi_d$ , and the mathematical form of the mapping from bodily states to sensory input,  $g(x)$ ; see Eq. (A.6) in the Appendix). Therefore, estimating homeostatic belief precision based on measurements of action signals alone could be useful to investigate the impact of therapeutic interventions (e.g. psychotherapy) within the same individual. By contrast, since  $\phi$  could differ across subjects, estimates of homeostatic belief precision obtained from regulatory signals alone should not be compared across individuals.

(ii) Measurements of bodily states only are available:

Computing homeostatic belief precision using the StdPE approach rests on information about action signals (see Eq. (7)). If only measurements of bodily state are available, one could, in principle, compute the action signal according to Eq. (16) in the Supplementary Material. However, this would require knowledge or assumptions about two unknown variables: (i)  $\lambda$ , the time constant of regulatory action, and (ii)  $f$ , the mapping from action to changes in bodily states. Both quantities would need to be known or estimated, for example, based on separate experimental measurements of regulatory signals and the ensuing changes in bodily states.

(iii) Measurements of both regulatory (action) signals and bodily states are available:

In this case, homeostatic belief precision can be computed fully (not only proportionally, as in the first case). Specifically, Section 2 (Eqs. (11–15)) in the Supplementary Material describes how the scaling factor  $\phi$  could be estimated that is required by Eq. (7) in the main text (Please note that, as shown in the Supplementary Material, estimating  $\phi$  also requires measurements of the sensory inputs  $y$ ).

Following this summary of different use cases, we now turn to the question what kind of bodily states and regulatory signals could be measured with existing and emerging techniques.

There are several bodily states that can be measured easily and repeatedly in time, in a manner that is non-invasive and safe, and whose homeostatic regulation is reasonably well understood. For example, this includes a number of cardiac and vascular variables that can be measured continuously, such as blood pressure, heart rate, or heart rate variability (HRV) (for a review of HRV, see (Acharya, Joseph, Kannathal, Lim, & Suri, 2006)). Similarly, a wide range of respiratory variables can be assessed, including frequency of breathing, tidal volume, inspiratory pressure, or partial pressure of oxygen or carbon dioxide. Electrodermal activity is another potentially interesting measure that can be obtained continuously and with relatively small technical effort (for review, see (Posada-Quintero & Chon, 2020)). Finally, time-varying concentrations of hormones from somatic glands, such as catecholamines or cortisol produced by the suprarenal gland, can be measured in discrete time steps on the order of minutes, using continuous blood sampling techniques; for examples, see (Bhake et al., 2019; Ward et al., 1983).

Measurements that reflect descending regulatory processes are technically considerably more challenging, although both neuronal and endocrine candidate signals exist. On the endocrine side, one could consider repeated plasma level measurements of pituitary hormones which determine activity of bodily glands, such as ACTH (e.g. (Sharma, Aoun, Wigham, Weist, & Veldhuis, 2014)). With regard to neural signals, it would be necessary to obtain recordings from structures that are purely efferent in function and, ideally, specifically affect a bodily state of interest. In principle, it is possible to directly record from sympathetic nerves in humans. This approach, called microneurography, involves the insertion of microelectrodes into a postganglionic sympathetic nerve bundle and can differentiate between sympathetic activity related to regulation of blood pressure versus thermoregulation (for review, see (Hart et al., 2017)). However, this approach is invasive and requires highly trained specialists. A less invasive approach would entail imaging techniques to record activity from organ-specific autonomic ganglia or



pre-ganglionic nuclei in the spinal cord. To our knowledge, non-invasive imaging of autonomic ganglia in humans has not been achieved yet. However, the new technique of optoacoustic imaging (e.g. (Gottschalk et al., 2019)) has interesting potential in this regard and, following appropriate validation, might be able to record activity in accessible autonomic ganglia. Imaging autonomic nuclei in the spinal cord should, in principle, be feasible using high-resolution spinal fMRI (for review, see (Tinnermann, Büchel, & Cohen-Adad, 2020)). However, we are not aware that such studies have already been conducted, and it is not clear which degree of organ-specificity can be achieved. Finally, it would be possible to use high-resolution fMRI to record activity from brainstem nuclei known to send regulatory efferents to the autonomic nervous system (e.g. periaqueductal gray, dorsal nucleus of vagus nerve, nucleus ambiguus). Here, the challenge is that, from the perspective of the BA model, visceromotor nuclei would be expected to contain both neurons representing setpoints (homeostatic beliefs) and neurons computing precision-weighted prediction errors (that drive regulatory efferent signals). This assumption is not experimentally verified, to our knowledge, but if it were correct, it would mean that fMRI signals from visceromotor nuclei could not be used in a “raw” fashion but would first need to be partitioned (e.g. using computational models) into signal components that are specific for homeostatic setpoints and for prediction errors, respectively.

Once suitable experimental measures of bodily states and regulatory signals are available, the approach presented in this paper might have useful clinical applications. Before discussing these, it is worth remembering that, according to the BA model, the precision of homeostatic beliefs determines how quickly and vigorously the brain will respond to bodily inputs that signal dyshomeostasis. If this precision is low, regulatory action is sluggish and there is a risk that severe perturbations of bodily states may not be corrected sufficiently quickly. On the other hand, a high homeostatic belief precision enables the brain to respond very rapidly; simply speaking, this is because the range of “permissible” (expected) bodily states becomes very narrow and even small changes in bodily states lead to reflex-like actions that “defend” the homeostatic belief. This, however, comes at the cost that regulation now becomes sensitive to small changes in sensory inputs from the body, and sensory noise alone may be sufficient to drive the system into a state where it continuously emits regulatory actions in the absence of any meaningful perturbations. This scenario is visible in Fig. 2 where high-amplitude action signals, driven by interoceptive surprise, begin to emerge as soon as homeostatic belief precision is tightened in anticipation of an upcoming perturbation (see the red “spike” in the top panel between

points ① and ②), leading to a period with massive amounts of regulatory action that is simply driven by sensory noise. This constellation is reminiscent of the symptomatology in affective and psychosomatic disorders, e.g. anxiety or stress disorders, where beliefs about upcoming threats may trigger descending neural and endocrine influences that lead to prolonged activation of the sympathetic nervous system, with possible adverse consequences for cardiovascular, immunological and metabolic systems (compare the concept of “allostatic load” (McEwen, 1998)) Peters, McEwen, and Friston (2017).

With these considerations in mind, the BA model by Stephan et al. (2016) highlights that homeostatic beliefs represent prime targets for psychotherapeutic or psychoeducational interventions. For example, mindfulness-based techniques (compare (Manjaly & Iglesias, 2020)) or hypnotherapeutic interventions with a bodily focus may prove useful to alter homeostatic beliefs and thus change the reactivity of brain-based reflex arcs for homeostatic control. In this context, subject-specific estimates of homeostatic belief precision and its change to controlled perturbations of bodily states might usefully guide clinical decisions. For example, accurate estimates of homeostatic belief precision could be useful for differential diagnosis and choosing individualized treatments, provided reference distributions for a population are known. Another option is to use estimates of homeostatic belief precision for evaluating the success of psychotherapeutic treatments. For the latter case, an estimate of homeostatic belief precision could even be useful when it cannot be computed exactly but only up to a scaling factor (compare Eq. (7) and subsequent discussion). This is because the estimate only depends on quantities (i.e. precision of sensory channels and the mapping from bodily states to sensory inputs; see Eq. (6) and the definition of  $\phi$  directly after Eq. (A.6)) that can be expected to be stable properties of an individual, at least as long measurements are taken for the same bodily state and experimental setting.

The above considerations highlight the clinical potential of accurate estimates of BBI. In this paper, we have presented some early and small steps towards these goals but have only been able to evaluate the proposed methods on synthetic data. How well the proposed methods work in practice remains to be tested empirically once suitable datasets, which fulfill the criteria described above, become available.

## Acknowledgements

This work was supported by the René and Susanne Braginsky Foundation (KES) and the University of Zurich (KES).

## Appendix A. Derivations of equations

### Appendix A.1. Statistics of Action $a$ : Derivation of Eqs. (5–6)

Inserting Eq. (2) into 3 yields the following expression for the action sequence:

$$a_t = -\pi_t(g(x_t) + \eta_{d,t} - g(\mu_{prior}))\partial_x g \quad (A.1)$$

The act of measuring the action  $a_t$  and bodily state  $x_t$  introduces two additional noise terms that are included in the measured action  $a_t^m$  and measured state  $x_t^m$ , respectively.

$$a_t^m = a_t + \eta_{a,t} \quad \text{with} \quad \eta_{a,t} \sim \mathcal{N}(0, \pi_a^{-1}) \quad (A.2)$$

$$x_t^m = x_t + \eta_{x,t} \quad \text{with} \quad \eta_{x,t} \sim \mathcal{N}(0, \pi_x^{-1}) \quad (A.3)$$

which results in the following formulation:

$$a_t^m = -\pi_t(g(x_t) + \eta_{d,t} - g(\mu_{prior}))\partial_x g + \eta_{a,t} \quad (A.4)$$

Using the fact that the sum of Gaussian random variables also follows a Gaussian distribution, the measured regulatory action  $a_t^m$  can be fully described by its expectation and variance. Given that the action states  $a_{t'}^m$  and the state variable  $x_{t'}^m \quad \forall t' \in [1, \dots, t-1]$  are observed, the statistics of the measured action  $a_t^m$  are computed as:

$$\begin{aligned}
\mathbb{E}[a_t^m | a_{1:t-1}^m; \pi_t] &= \mathbb{E}[-\pi_t(g(x_t) + \eta_{d,t} - g(\mu_{prior,t}))\partial_x g + \eta_{a,t}] \\
&= -\pi_t(\mathbb{E}[g(x_t)] - \mathbb{E}[g(\mu_{prior,t})])\partial_x g \\
&= -\pi_t(g(x_t) - g(\mu_{prior,t}))\partial_x g \\
&= -\pi_t\theta_t
\end{aligned} \tag{A.5}$$

with  $\theta_t = (g(x_t) - g(\mu_{prior,t}))\partial_x g$  denoting the prediction error (i.e. the difference between the actual sensory input and the sensory input expected under the prior mean), scaled by the derivative of the mapping from bodily states to sensory inputs,  $\partial_x g$ .

Furthermore,

$$\begin{aligned}
\text{Var}[a_t^m | a_{1:t-1}^m; \pi_t] &= \text{Var}[-\pi_t(g(x_t) + \eta_{d,t} - g(\mu_{prior,t}))\partial_x g + \eta_{a,t}] \\
&= \pi_t^2 (\text{Var}[\eta_{d,t}](\partial_x g)^2 + \text{Var}[\eta_{a,t}]) \\
&= \pi_t^2 \phi^{-1} + \pi_a^{-1}
\end{aligned} \tag{A.6}$$

with  $\phi^{-1} = \pi_d^{-1}(\partial_x g)^2$  denoting the transformation of the sensory input precision through the second derivative of  $g$ , the mapping from bodily states to sensory inputs (see [Supplementary Material](#) on how this term can be estimated).

## Appendix A.2. StdPE: Derivation of Eq. (7)

$$\begin{aligned}
\pi_t &= \sqrt{\phi(\text{Var}[a_t^m | a_{1:t-1}^m; \pi_t] - \pi_a^{-1})} \\
\Rightarrow \pi_t &\approx \sqrt{\phi \text{Var}[a_{t-N+1:t}^m; \pi_t] - \phi \pi_a^{-1}}
\end{aligned} \tag{A.7}$$

Multiple measurements of the action sequence  $a_t$  can be used to increase the measurement precision  $\pi_a$ , as the sum of white Gaussian noise terms will converge to their mean  $\mu_a = 0$ . When a large number of measurements are available, assuming that the precision of the measurement noise  $\pi_a \rightarrow \infty$ , [Eq. \(A.7\)](#) simplifies to:

$$\pi_t \approx \sqrt{\phi} \cdot \sqrt{\text{Var}[a_{t-N+1:t}^m; \pi_t]} \tag{A.8}$$

with  $\sqrt{\phi}$  denoting the scaling factor. Thus the precision at time point  $t$  can be estimated as the square root of the variance of the windowed action sequence up to a scaling factor. Details on how this scaling factor can be determined can be found in the [Supplementary Material](#).

## Appendix B. Supporting information

Supplementary data associated with this article can be found in the online version at [doi:10.1016/j.biopsycho.2021.108190](https://doi.org/10.1016/j.biopsycho.2021.108190).

## References

- Acharya, U. R., Joseph, K. P., Kannathal, N., Lim, C. M., & Suri, J. S. (2006). Heart rate variability: A review. *Medical and Biological Engineering and Computing*, 44(12), 1031–1051. <https://doi.org/10.1007/s11517-006-0119-0>
- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, 218(3), 611–643. <https://doi.org/10.1007/s00429-012-0475-5>
- Allen, M., Levy, A., Parr, T., & Friston, K. J. (2019). In the body's eye: The computational anatomy of interoceptive inference. *BioRxiv*, Article 603928. <https://doi.org/10.1101/603928>
- Badoud, D., & Tsakiris, M. (2017). From the body's viscera to the body's image: Is there a link between interoception and body image concerns? *Neuroscience & Biobehavioral Reviews*, 77, 237–246. <https://doi.org/10.1016/j.neubiorev.2017.03.017>
- Barca, L., & Pezzulo, G. (2020). Keep your interoceptive streams under control: An active inference perspective on anorexia nervosa. *Cognitive, Affective, & Behavioral Neuroscience*, 20, 427–440. <https://doi.org/10.3758/s13415-020-00777-6>
- Bhake, R., Kluckner, V., Stassen, H., Russell, G., Leendertz, J., Stevens, K., ... Lightman, S. (2019). Continuous free cortisol profiles—circadian rhythms in healthy men. *The Journal of Clinical Endocrinology & Metabolism*, 104(12), 5935–5947. <https://doi.org/10.1210/je.2019-00449>
- Bonaz, B., Lane, R. D., Oshinsky, M. L., Kenny, P. J., Sinha, R., Mayer, E. A., & Critchley, H. D. (2021). Diseases, disorders, and comorbidities of interoception. *Trends in Neurosciences*, 44(1), 39–51. <https://doi.org/10.1016/j.tins.2020.09.009>
- Browning, M., Carter, C. S., Chatham, C., Den Ouden, H., Gillan, C. M., Baker, J. T., Chekroud, A. M., Cools, R., Dayan, P., Gold, J., Goldstein, R. Z., Hartley, C. A., Kepecs, A., Lawson, R. P., Mourao-Miranda, J., Phillips, M. L., Pizzagalli, D. A., Powers, A., Rindskopf, D., Roiser, J. P., Schmack, K., Schiller, D., Sebold, M., Stephan, K. E., Frank, M. J., Huys, Q., & Paulus, M. (2020). Realizing the clinical potential of computational psychiatry: Report from the banbury center meeting, february 2019. *Biological Psychiatry*, 88(2), e5–e10. <https://doi.org/10.1016/j.biopsych.2019.12.026>
- Critchley, H. D., & Harrison, N. A. (2013). Visceral influences on brain and behavior. *Neuron*, 77(4), 624–638. <https://doi.org/10.1016/j.neuron.2013.02.008>
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456), 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Friston, K. (2009). The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, 13(7), 293–301. <https://doi.org/10.1016/j.tics.2009.04.005>
- Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: A free-energy formulation. *Biological Cybernetics*, 102(3), 227–260. <https://doi.org/10.1007/s00422-010-0364-z>
- Garfinkel, S. N., Tiley, C., O'Keefe, S., Harrison, N. A., Seth, A. K., & Critchley, H. D. (2016). Discrepancies between dimensions of interoception in autism: Implications for emotion and anxiety. *Biological Psychology*, 114, 117–126. <https://doi.org/10.1016/j.biopsycho.2015.12.003>
- Gottschalk, S., Degtyaruk, O., Mc Larney, B., Rebling, J., Hutter, M. A., Deán-Ben, X. L., ... Razansky, D. (2019). Rapid volumetric optoacoustic imaging of neural dynamics across the mouse brain. *Nature Biomedical Engineering*, 3(5), 392–401. <https://doi.org/10.1038/s41551-019-0372-9>
- Gu, X., Hof, P. R., Friston, K. J., & Fan, J. (2013). Anterior insular cortex and emotional awareness. *Journal of Comparative Neurology*, 521(15), 3371–3388. <https://doi.org/10.1002/cne.23368>
- Hart, E. C., Head, G. A., Carter, J. R., Wallin, B. G., May, C. N., Hamza, S. M., ... Osborn, J. W. (2017). Recording sympathetic nerve activity in conscious humans and other mammals: Guidelines and the road to standardization. *American Journal of Physiology. Heart and Circulatory Physiology*, 312, 1031. <https://doi.org/10.1152/ajpheart.00703.2016>
- Henningsen, P., Gündel, H., Kop, W. J., Löwe, B., Martin, A., Rief, W., Rosmalen, J., Schröder, A., van der Feltz-Cornelis, C., Van den Bergh, O., & EURONET-SOMA, G. (2018). Persistent physical symptoms as perceptual dysregulation: A neuropsychobehavioral model and its clinical implications. *Psychosomatic Medicine*, 80(5), 422–431. <https://doi.org/10.1097/PSY.0000000000000588>
- Khalsa, S. S., Adolphs, R., Cameron, O. G., Critchley, H. D., Davenport, P. W., Feinstein, J. S., Feusner, J. D., Garfinkel, S. N., Lane, R. D., Mehling, W. E., Mehling, W. E., Meuret, A. E., Nemeroff, C. B., Oppenheimer, S., Petzschner, F. H., Pollatos, O., Rhudy, J. L., Schramm, L. P., Simmons, W. K., Stein, M. B., Stephan, K. E., Van den Bergh, O., Van Diest, I., von Leupoldt, A., Paulus, M. P., & Interoception Summit, p (2018). Interoception and mental health: A roadmap.

- Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 3(6), 501–513. <https://doi.org/10.1016/j.bpsc.2017.12.004>
- Knill, D. C., & Pouget, A. (2004). The bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12), 712–719. <https://doi.org/10.1016/j.tins.2004.10.007>
- Manjaly, Z.-M., Harrison, N. A., Critchley, H. D., Do, C. T., Stefanics, G., Wenderoth, N., ... Stephan, K. E. (2019). Pathophysiological and cognitive mechanisms of fatigue in multiple sclerosis. *Journal of Neurology, Neurosurgery & Psychiatry*, 90(6), 642–651. <https://doi.org/10.1136/jnnp-2018-320050>
- Manjaly, Z.-M., & Iglesias, S. (2020). A computational theory of mindfulness based cognitive therapy from the “Bayesian Brain” perspective. *Frontiers in Psychiatry*, 11, 404. <https://doi.org/10.3389/fpsy.2020.00404>
- McEwen, B. S. (1998). Stress, adaptation, and disease: Allostasis and allostatic load. *Annals of the New York Academy of Sciences*, 840(1), 33–44. <https://doi.org/10.1111/j.1749-6632.1998.tb09546.x>
- Paulus, M. P., Feinstein, J. S., & Khalsa, S. S. (2019). An active inference approach to interoceptive psychopathology. *Annual Review of Clinical Psychology*, 15, 97–122. <https://doi.org/10.1146/annurev-clinpsy-050718-095617>
- Peters, A., McEwen, B. S., & Friston, K. (2017). Uncertainty and stress: Why it causes diseases and how it is mastered by the brain. *Progress in Neurobiology*, 156, 164–188. <https://doi.org/10.1016/j.pneurobio.2017.05.004>
- Petzschner, F. H., Garfinkel, S. N., Paulus, M. P., Koch, C., & Khalsa, S. S. (2021). Computational models of interoception and body regulation. *Trends in Neurosciences*, 44(1), 63–76. <https://doi.org/10.1016/j.tins.2020.09.012>
- Petzschner, F. H., Weber, L. A., Gard, T., & Stephan, K. E. (2017). Computational psychosomatics and computational psychiatry: Toward a joint framework for differential diagnosis. *Biological Psychiatry*, 82(6), 421–430. <https://doi.org/10.1016/j.biopsych.2017.05.012>
- Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, 134, 17–35. <https://doi.org/10.1146/annurev-clinpsy-050718-095617>
- Posada-Quintero, H. F., & Chon, K. H. (2020). Innovations in electrodermal activity data collection and signal processing: A systematic review. *Sensors*, 20(2), 479. <https://doi.org/10.3390/s20020479>
- Quadt, L., Critchley, H. D., & Garfinkel, S. N. (2018). The neurobiology of interoception in health and disease. *Annals of the New York Academy of Sciences*, 1428(1), 112–128. <https://doi.org/10.1111/nyas.13915>
- Schrödinger, E. (1943). What is life? Lectures delivered under the auspices of the Dublin Institute for Advanced Studies.
- Schulz, A., Rost, S., Flasiński, T., Dierolf, A. M., Lutz, A. P., Münch, E. E., Mertens, V.-C., Witthöft, M., & Vögele, C. (2020). Distinctive body perception mechanisms in high versus low symptom reporters: A neurophysiological model for medically-unexplained symptoms. *Journal of Psychosomatic Research*, 137, Article 110223. <https://doi.org/10.1016/j.jpsychores.2020.110223>
- Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17(11), 565–573. <https://doi.org/10.1016/j.tics.2013.09.007>
- Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 371(1708), Article 20160007. <https://doi.org/10.1098/rstb.2016.0007>
- Sharma, A. N., Aoun, P., Wigham, J. R., Weist, S. M., & Veldhuis, J. D. (2014). Estradiol, but not testosterone, heightens cortisol-mediated negative feedback on pulsatile ACTH secretion and ACTH approximate entropy in unstressed older men and women. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 306(9), R627–R635. <https://doi.org/10.1152/ajpregu.00551.2013>
- Smith, R., Badcock, P., & Friston, K. J. (2021). Recent advances in the application of predictive coding and active inference models within clinical neuroscience. *Psychiatry and Clinical Neurosciences*, 75(1), 3–13. <https://doi.org/10.1111/pcn.13138>
- Smith, R., Kuplicki, R., Feinstein, J., Forthman, K. L., Stewart, J. L., Paulus, M. P., ... Khalsa, S. S. (2020). A bayesian computational model reveals a failure to adapt interoceptive precision estimates across depression, anxiety, eating, and substance use disorders. *PLoS Computational Biology*, 16(12), Article 1008484. <https://doi.org/10.1371/journal.pcbi.1008484>
- Stephan, K. E., Iglesias, S., Heinze, J., & Diaconescu, A. O. (2015). Translational perspectives for computational neuroimaging. *Neuron*, 87(4), 716–732. <https://doi.org/10.1016/j.neuron.2015.07.008>
- Stephan, K. E., Manjaly, Z. M., Mathys, C. D., Weber, L. A., Paliwal, S., Gard, T., Tittgemeyer, M., Fleming, S. M., Haker, H., Seth, A. K., Seth, A. K., & Petzschner, F. H. (2016). Allostatic self-efficacy: A metacognitive theory of dyshomeostasis-induced fatigue and depression. *Frontiers in Human Neuroscience*, 10, 550. <https://doi.org/10.3389/fnhum.2016.00550>
- Stephan, K. E., & Mathys, C. (2014). Computational approaches to psychiatry. *Current Opinion in Neurobiology*, 25, 85–92. <https://doi.org/10.1016/j.conb.2013.12.007>
- Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior*, 106(1), 5–15. <https://doi.org/10.1016/j.physbeh.2011.06.004>
- Tinnermann, A., Büchel, C., & Cohen-Adad, J. (2021). Cortico-spinal imaging to study pain. *NeuroImage*, 224, Article 117439. <https://doi.org/10.1016/j.neuroimage.2020.117439>
- Turner, J. S. (2019). Homeostasis as a fundamental principle for a coherent theory of brains. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 374(1774), Article 20180373. <https://doi.org/10.1098/rstb.2018.0373>
- Ward, M. M., Mefford, I. N., Parker, S. D., Chesney, M. A., Taylor, B. C., Keegan, D. L., & Barchas, J. D. (1983). Epinephrine and norepinephrine responses in continuously collected human plasma to a series of stressors. *Psychosomatic Medicine*, 45(6), 471–486. <https://doi.org/10.1097/00006842-198312000-00002>
- Ziemke, T. (2016). The body of knowledge: On the role of the living body in grounding embodied cognition. *Biosystems*, 148, 4–11. <https://doi.org/10.1016/j.biosystems.2016.08.005>