



Uncertainty in perception and the Hierarchical Gaussian Filter

Christoph D. Mathys^{1,2,3,4,*}, Ekaterina I. Lomakina^{3,4,5}, Jean Daunizeau⁶, Sandra Iglesias^{3,4}, Kay H. Brodersen^{3,4}, Karl J. Friston¹ and Klaas E. Stephan^{1,3,4}

¹ Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, London, UK

² Max Planck UCL Centre for Computational Psychiatry and Ageing Research, London, UK

³ Translational Neuromodeling Unit, Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Zurich, Switzerland

⁴ Laboratory for Social and Neural Systems Research (SNS Lab), Department of Economics, University of Zurich, Zurich, Switzerland

⁵ Department of Computer Science, ETH Zurich, Zurich, Switzerland

⁶ Institut du Cerveau et de la Moelle Épinrière, Hôpital Pitié Salpêtrière, Paris, France

Edited by:

Hauke R. Heekeren, Freie Universität Berlin, Germany

Reviewed by:

Dirk Ostwald, Freie Universität Berlin, Germany

Mateus Joffily, Centre National de la Recherche Scientifique, France

*Correspondence:

Christoph D. Mathys, Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK
e-mail: c.mathys@ucl.ac.uk

In its full sense, perception rests on an agent's model of how its sensory input comes about and the inferences it draws based on this model. These inferences are necessarily uncertain. Here, we illustrate how the Hierarchical Gaussian Filter (HGF) offers a principled and generic way to deal with the several forms that uncertainty in perception takes. The HGF is a recent derivation of one-step update equations from Bayesian principles that rests on a hierarchical generative model of the environment and its (in)stability. It is computationally highly efficient, allows for online estimates of hidden states, and has found numerous applications to experimental data from human subjects. In this paper, we generalize previous descriptions of the HGF and its account of perceptual uncertainty. First, we explicitly formulate the extension of the HGF's hierarchy to any number of levels; second, we discuss how various forms of uncertainty are accommodated by the minimization of variational free energy as encoded in the update equations; third, we combine the HGF with decision models and demonstrate the inversion of this combination; finally, we report a simulation study that compared four optimization methods for inverting the HGF/decision model combination at different noise levels. These four methods (Nelder–Mead simplex algorithm, Gaussian process-based global optimization, variational Bayes and Markov chain Monte Carlo sampling) all performed well even under considerable noise, with variational Bayes offering the best combination of efficiency and informativeness of inference. Our results demonstrate that the HGF provides a principled, flexible, and efficient—but at the same time intuitive—framework for the resolution of perceptual uncertainty in behaving agents.

Keywords: uncertainty, volatility, Bayesian inference, hierarchical modeling, filtering, free energy, learning, decision-making

INTRODUCTION

Perception has long been proposed to take place in the context of prediction (Helmholtz, 1860). This entails that agents have a model of the environment which generates their sensory input. Probability theory formally prescribes how agents should learn about their environment from sensory information, given a model. This rests on sequential updating of beliefs according to Bayes' theorem, where beliefs represent inferences about hidden states of the environment in the form of posterior probability distributions. It is this process that we refer to as perception. Beliefs about hidden states are inherently uncertain. This uncertainty has two sources. First, even when states are constant, the amount of sensory information will in general be too little to infer them exactly. This has been referred to as *informational uncertainty* or *estimation uncertainty* (Payzan-LeNestour and Bossaerts, 2011). The second source of uncertainty

is the possibility that states change with time, i.e., *environmental uncertainty*.

Various models have suggested how an agent may deal with an environment fraught with both kinds of uncertainty (e.g., Yu and Dayan, 2003, 2005; Nassar et al., 2010; Payzan-LeNestour and Bossaerts, 2011; Wilson et al., 2013). Here, we discuss an alternative approach that derives closed form update equations for the hidden states, and crucially, for the uncertainty about them, by variational inversion of a generic hierarchical generative model that reflects the time-varying structure of the environment in its higher levels. This derivation has the advantage that the resulting updates optimize a clearly defined objective function, namely variational free energy. Since this quantity is an approximation to surprise (i.e., to the negative log-probability of sensory input), the updates are optimal in the sense that they minimize surprise, given an agent's individual model. Furthermore,

the updates explicitly reflect informational and environmental uncertainty.

Our approach makes use of a framework which assumes that agents have an internal generative model of their sensory input. This model is generative in the sense that it describes how sensory inputs are generated by the external world. It does this by assigning a probability (the likelihood) to each sensory input given states (which vary with time) and parameters (which are constant in time) and by completing this with a prior probability distribution for states and parameters. While the purpose of the model is to predict input emanating from the external world, it is internal in the sense that it reflects the agent's *beliefs* about how sensory inputs are generated by the external world.

While Bayesian belief updating is optimal from the perspective of probability theory, it requires computing complicated integrals which are not tractable analytically and difficult to evaluate in real time. Although some attempts to design a Bayesian model of how biological agents learn in a changing environment were remarkably successful in explaining empirical behavior (Behrens et al., 2007, 2008), they were restricted by the computational burden imposed by these models and the assumption that the learning process was identical across subjects. Recently, however, theoretical advances have enabled computationally efficient approximations to exact Bayesian inference during learning (e.g., Friston, 2009; Daunizeau et al., 2010a,b) and have furnished a basis for biologically plausible mechanisms that might underlie belief updating in the brain. These approaches rest on variational Bayesian techniques which optimize a free-energy bound on the surprise about sensory inputs, given a model of the environment, and represent a special case of the general “Bayesian brain” hypothesis (Dayan et al., 1995; Knill and Pouget, 2004; Körding and Wolpert, 2006; Friston, 2009; Doya et al., 2011). This hypothesis has been highly influential in recent years, shaping concepts of brain function and inspiring the design of many specific computational models (see Friston and Dolan, 2010, for review). However, for practical applications to empirical data, a general purpose modeling framework has been lacking that would allow for straightforward “off the shelf” implementations of models explaining trial-wise empirical data (e.g., behavioral responses, eye movements, evoked response amplitude in EEG etc.) from the Bayesian brain perspective. This is in contrast to reinforcement learning (RL) models which, due to their simplicity and computational efficiency, have found widespread application in experimental neuroscience, for example, in the analysis of functional magnetic resonance imaging (fMRI) and behavioral data (for reviews, see Daw and Doya, 2006; O’Doherty et al., 2007).

To fill this gap and provide a generic, robust and flexible framework for analysis of trial-wise data from the Bayesian brain perspective, we recently introduced the Hierarchical Gaussian Filter (HGF), a hierarchical Bayesian model Mathys et al. (2011) in which states evolve as coupled Gaussian random walks, such that each state determines the step size of the evolution of the next lower state (for examples of applications, cf. Iglesias et al., 2013; Joffily and Coricelli, 2013; Vossel et al., 2013). Based on a mean field approximation to the full Bayesian solution, we derived analytic update equations whose form resembles RL updates, with dynamic learning rates and precision-weighted prediction errors.

These highly efficient update equations made our approach well suited for filtering purposes, i.e., predicting the value of (and, crucially, the uncertainty about) a hidden and moving quantity based on all information acquired up to a certain point. Our original formulation (Mathys et al., 2011) only contained three levels; here, we extend the HGF explicitly to any number of levels and show that the update equations maintain the same form across all levels because they are derived on the basis of the same coupling. Furthermore, the derivation of the variational energies involved in the inversion is given in much more detail than in Mathys et al. (2011). It is important to note that “perceptual uncertainty” has a broader meaning here than in Mathys et al. (2011), where it was used more narrowly for that part of the informational uncertainty that relates to sensory input.

Furthermore, in this paper we describe how the HGF is applied in the “observing the observer” framework developed in (Daunizeau et al., 2010a,b). This framework is based on a clear separation of two model components: First, the agent’s perception of (inference about) its environment, i.e., the posterior estimates provided by the agent’s model of how its sensory input is generated. Second, the agent’s observed actions (i.e., decisions or responses) which are (probabilistic) consequences of the agent’s beliefs about its environment. We call the first model *perceptual*, while the second is the *decision* or *response* model.

The “observing the observer” framework is meta-Bayesian in that it enables Bayesian inference (by an observer or experimenter) on Bayesian inference (by a subject). It requires four elements: (1) a generative model of sensory inputs (i.e., a perceptual model), (2) a computationally efficient and robust method for model inversion, (3) a loss function for actions depending on the inferred state, and (4) a decision model. A specific suggestion for the first two elements is contained in Mathys et al. (2011). In what follows, we extend and generalize this description, discussing specifically the nature of the coupling between levels, choice of coordinates at higher levels, and how to deal with sensory inputs that arrive at irregular time intervals.

In the following section of this paper, we set out our theoretical framework. We first define the HGF model formally. We then proceed with its variational inversion, which gives us closed one-step update equations. Next, we show how the HGF can serve as a perceptual model for any decision model that provides a mapping from the HGF’s representations of the environment to the probability of an observed decision. We then derive an objective function whose optimization leads to maximum-a-posterior (MAP) estimates for the parameters of the HGF and the decision model, followed by a short discussion of how the choice of decision model affects which HGF parameters can be estimated.

In the next section, we turn to examples and simulations. We first deal with categorical outcomes and sensory uncertainty. To complement this, we introduce a decision model for binary choices and use it to give an example of model inversion and comparison based on two different but closely related decision models. We do this by estimating model parameters from empirical data in a single subject, juxtaposing the two different response models (which do and do not take into account the uncertainty of beliefs) and the ensuing differences in inferred state trajectories. Next, we introduce a decision model for a one-armed bandit. As

in the preceding examples of decision models, we base the decision rule on the agent's expected loss under a given loss function. In the last part of this section, we report the results of a simulation study which demonstrates the feasibility of accurate model inversion, i.e., inferring known HGF parameter values from observed decisions. This test of model inversion was performed under different levels of noise and using four different optimization methods (Markov chain Monte Carlo, the Nelder–Mead simplex algorithm, Gaussian process-based global optimization, and variational Bayes).

Together, the theoretical derivations and simulation results provided in this paper generalize the framework of the HGF and demonstrate its utility for estimating individual approximations to Bayes-optimality from observed decision-making under uncertainty.

THEORETICAL FRAMEWORK

THE HIERARCHICAL GAUSSIAN FILTER (HGF)

The goal of the model introduced in Mathys et al. (2011) is simple and general: to describe how an agent learns about a continuous uncertain quantity (i.e., random variable) x that moves. One generic way of describing this motion is a Gaussian random walk:

$$x^{(k)} \sim \mathcal{N}(x^{(k-1)}, \vartheta), k = 1, 2, \dots \quad (1)$$

where k is a time index, and $x^{(k-1)}$ and ϑ are the mean and variance (not standard deviation) of a Gaussian distribution, respectively. In this formulation, the volatility in x is governed by the positive constant ϑ (in this paper, we define volatility as the variance of a time series per unit of time); however, there is in principle no reason to assume that volatility is constant. To allow for changes in volatility, we replace ϑ by a positive function f of a second random variable, x_2 , while x becomes x_1 :

$$x_1^{(k)} \sim \mathcal{N}(x_1^{(k-1)}, f(x_2)) \quad (2)$$

We may now further assume that x_2 performs a Gaussian random walk of its own, with a constant variance ϑ , so that the model for x_2 is the same as the one for x in Equation 1. This adding of levels of Gaussian random walks coupled by their variances can now continue up to any number n of levels in the hierarchy, as illustrated in **Figure 1**. At each level i , the coupling to the next highest level $i + 1$ is given by a positive function $f_i(x_{i+1})$ which represents the variance or step size of the random walk:

$$x_i^{(k)} \sim \mathcal{N}(x_i^{(k-1)}, f_i(x_{i+1})), i = 1, \dots, n - 1. \quad (3)$$

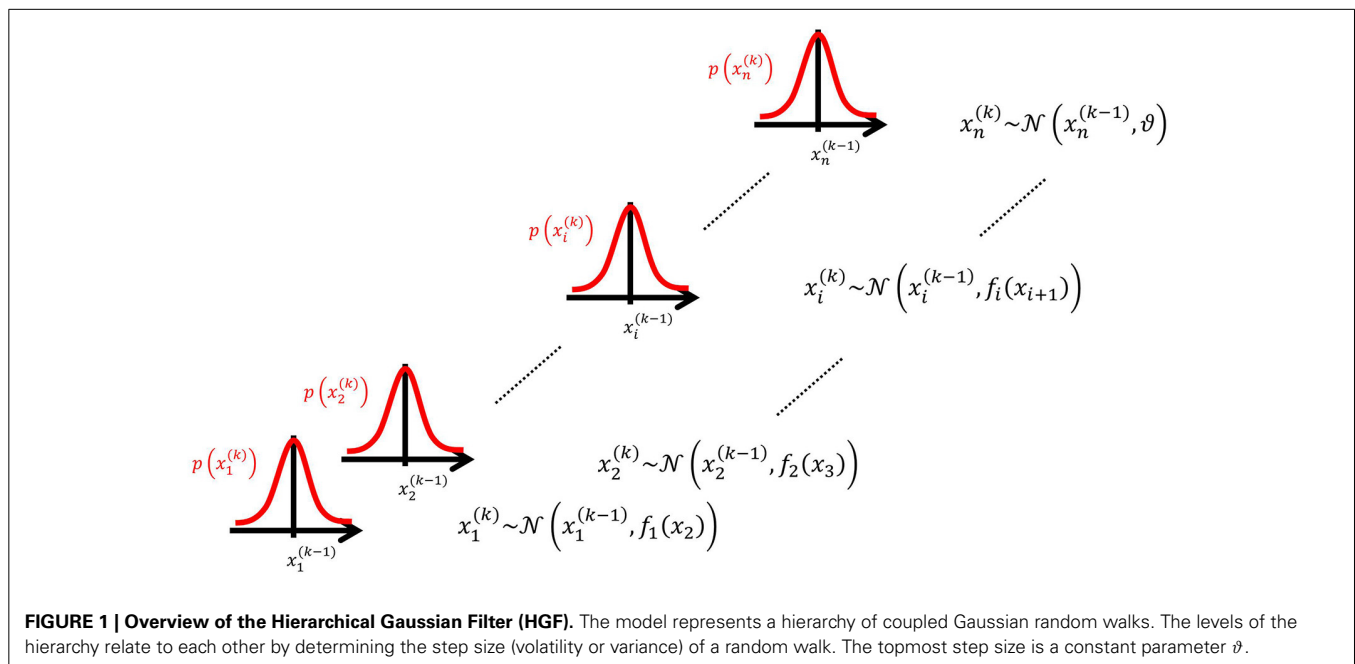
At the top level, instead of f_n , we have ϑ :

$$x_n^{(k)} \sim \mathcal{N}(x_n^{(k-1)}, \vartheta), \vartheta > 0. \quad (4)$$

To complete our model, we still need to define the f_i in Equation 3. A flexible and straightforward approach to this is to allow any positive analytic f_i , but to expand it in powers to first order to give it a simple functional form. However, since f_i has to be positive, we cannot approximate it by expanding it directly. Instead, dropping indices for clarity, we expand its logarithm (for details, see Appendix A), which motivates our definition of coupling between levels:

$$f_i(x_{i+1}) \stackrel{\text{def}}{=} \exp(\kappa_i x_{i+1} + \omega_i) \quad (5)$$

As we will see below, this form of coupling has the additional advantage of enabling the derivation of simple one-step update equations under a mean-field approximation.



By further assuming that inputs (observations) $u^{(k)}$ are generated by means of a Gaussian emission distribution of the form.

$$u^{(k)} \sim \mathcal{N}(x_1^{(k)}, \hat{\pi}_u^{-1}), \quad (6)$$

where $\hat{\pi}_u$ denotes the precision of the emission distribution, the model defined by Equations 3 and 5 can be used for online prediction of $x_1^{(k)}$, i.e., *filtering*. Considering that the model consists of a hierarchy of Gaussian random walks, this motivates why we refer to it as the *HGF*.

To illustrate this, we might imagine a time series of financial data. $u^{(k)}$ could be the observed returns of a particular security. $x_1^{(k)}$ then is the underlying quantity (the true return after observation noise $\hat{\pi}_u^{-1}$ has been filtered out). ω_i is the tonic (i.e., time-invariant) part of the (log-)volatility of x_1 , while κx_2 is the phasic (i.e., time-varying) part. Accounting for the scaling by κ , x_2 is now the scaled phasic log-volatility of x_1 . This pattern then repeats until the top of the hierarchy is reached. One of the advantages of the HGF is that volatility is captured hierarchically: not only returns are volatile, but also their volatility, and the volatility of the volatility, etc.

APPROXIMATE INVERSION FOR SENSORY INPUT AT IRREGULAR INTERVALS

Sensing takes place at the bottom of the hierarchy: u is the agent's sensory input. To allow for input that comes at irregular intervals, we can multiply the variance of the random walks at all levels by the time $t^{(k)}$ that elapses between the arrival of inputs $u^{(k-1)}$ and $u^{(k)}$:

$$x_i^{(k)} \sim \mathcal{N}(x_i^{(k-1)}, t^{(k)} f_i(x_{i+1})), i = 1, \dots, n-1. \quad (7)$$

This proportionality of the variance to time reflects the fact that the mean squared distance of a quantity performing a Gaussian random walk from its origin is proportional to time (cf. the connection between Gaussian random walks, Brownian motion, and the heat equation; Evans, 2010). For inputs at regular intervals, we may simply set $t^{(k)} = 1$ for all k , effectively removing $t^{(k)}$ from the model.

We can now derive update equations using the variational inversion method introduced in Mathys et al. (2011). This approximate inversion assumes Gaussian posteriors at all levels with means μ_i and precisions (inverse variances) π_i :

$$x_i^{(k)} | u^{(1)}, \dots, u^{(k)}, \chi \sim \mathcal{N}(\mu_i^{(k)}, (\pi_i^{(k)})^{-1}), \quad (8)$$

where $\chi \stackrel{\text{def}}{=} \{\kappa_1, \omega_1, \dots, \kappa_{n-1}, \omega_{n-1}, \vartheta\}$. This is an approximation because the true posterior distribution $p(x_i^{(k)} | u^{(1)}, \dots, u^{(k)}, \chi)$ will deviate somewhat from a Gaussian shape. A discussion of the variational nature and the implications of this approximation can be found in Mathys et al. (2011). The sufficient statistics μ_i and π_i are the quantities that are updated after each new input u according to the following equations (cf. the

Discussion, where we give a natural interpretation of them in terms of learning rate and prediction error):

$$\mu_i^{(k)} = \hat{\mu}_i^{(k)} + \frac{1}{2} \kappa_{i-1} v_{i-1}^{(k)} \frac{\hat{\pi}_{i-1}^{(k)}}{\pi_i^{(k)}} \delta_{i-1}^{(k)} \quad (9)$$

$$\pi_i^{(k)} = \hat{\pi}_i^{(k)} + \frac{1}{2} \left(\kappa_{i-1} v_{i-1}^{(k)} \hat{\pi}_{i-1}^{(k)} \right)^2 \left(1 + \left(1 - \frac{1}{v_{i-1}^{(k)} \pi_{i-1}^{(k-1)}} \right) \delta_{i-1}^{(k)} \right) \quad (10)$$

with

$$v_i^{(k)} \stackrel{\text{def}}{=} \begin{cases} t^{(k)} \exp(\kappa_i \mu_{i+1}^{(k-1)} + \omega_i), & i = 1, \dots, n-1 \\ t^{(k)} \vartheta, & i = n \end{cases} \quad (11)$$

$$\hat{\mu}_i^{(k)} \stackrel{\text{def}}{=} \mu_i^{(k-1)} \quad (12)$$

$$\hat{\pi}_i^{(k)} \stackrel{\text{def}}{=} \frac{1}{\sigma_i^{(k-1)} + v_i^{(k)}} \quad (13)$$

$$\delta_i^{(k)} \stackrel{\text{def}}{=} \frac{\sigma_i^{(k)} + (\mu_i^{(k)} - \hat{\mu}_i^{(k)})^2}{\sigma_i^{(k-1)} + v_i^{(k)}} - 1 \quad (14)$$

Variances (i.e., inverse precisions) are denoted by $\sigma_i^{(k)} = 1/\pi_i^{(k)}$. Note that irregular intervals between inputs are fully accounted for by the factor $t^{(k)}$ in Equation 11. While the updates of Equations 9 and 10 apply to all but the first level, they are different at the first level:

$$\mu_1^{(k)} = \hat{\mu}_1^{(k)} + \frac{\hat{\pi}_u}{\pi_1^{(k)}} \delta_u^{(k)} \quad (15)$$

$$\pi_1^{(k)} = \hat{\pi}_1^{(k)} + \hat{\pi}_u \quad (16)$$

with

$$\delta_u^{(k)} \stackrel{\text{def}}{=} u^{(k)} - \hat{\mu}_1^{(k)} \quad (17)$$

where $\hat{\mu}_1^{(k)}$ and $\hat{\pi}_1^{(k)}$ are defined by Equations 12 and 13. The different form of the updates at the first level arises because, at the first level, the direction of inference is from u to x_1 , which appears in the *mean* of the Gaussian in Equation 6, while at all higher levels, the direction of inference is from x_i to x_{i+1} , which appears in the *variance* of the Gaussian in Equation 3. This results in the updates being driven by different kinds of prediction errors: value prediction errors (VAPes) at the first level, volatility (i.e., variance) prediction errors (VOPEs) at all higher levels. We elaborate this distinction in the Discussion below.

The details of this inversion are given in Appendix B. The notation chosen here emphasizes the role of precisions in the updating process more than that used in (Mathys et al., 2011); they are, however, equivalent.

MAXIMUM-A-POSTERIORI (MAP) PARAMETER ESTIMATION

Given the above update equations, initial representations $\lambda^{(0)} \stackrel{\text{def}}{=} \{\mu_1^{(0)}, \pi_1^{(0)}, \dots, \mu_n^{(0)}, \pi_n^{(0)}\}$ (i.e., the means μ_i and precisions π_i of the states x_i at time $k = 0$), and priors on the perceptual parameters χ , we could invert the model (i.e., estimate the values of χ that lead to least aggregate surprise about sensory inputs) from sensory inputs alone; this would provide us with state trajectories and parameters which represent an ideal Bayesian agent, where “ideal” means experiencing the least surprise about sensory inputs. However, our goal is usually different; it is to estimate subject-specific parameters (which encode the individual’s approximation to Bayes-optimality) from his/her observed behavior, as formalized in the “observing the observer” framework (Daunizeau et al., 2010a,b). To achieve this goal, we will now bring the HGF into this framework. This requires the introduction of a response model which links the agent’s current estimates $\lambda^{(k)} \stackrel{\text{def}}{=} \{\mu_1^{(k)}, \sigma_1^{(k)}, \dots, \mu_n^{(k)}, \sigma_n^{(k)}\}$ of states to expressed decisions $y^{(k)}$ and which also contains subject-specific parameters $\xi \stackrel{\text{def}}{=} \{\xi_1, \xi_2, \dots\}$. For example, a useful response model (e.g., Iglesias et al., 2013), which we also use in our simulation study below, is the unit-square sigmoid, which maps the predictive probability $m^{(k)}$ that the next outcome will be 1 onto the probabilities $p(y^{(k)} = 1)$ and $p(y^{(k)} = 0)$ that the agent will choose response 1 or 0, respectively:

$$p(y | m, \xi) = \left(\frac{m^\xi}{m^\xi + (1 - m)^\xi} \right)^y \cdot \left(\frac{(1 - m)^\xi}{m^\xi + (1 - m)^\xi} \right)^{1-y}, \quad (18)$$

where, for clarity, we have omitted time indices on y and m . For this to serve as a response model for the inversion of the HGF, the predictive probability $m = m(\lambda)$ has to be a function of the quantities λ the HGF keeps track of. This model is explained and discussed in detail below. **Figure 4** is a graphical representation of it. For our present purposes, the only important point is that it contains a parameter ξ that captures how deterministically y is associated with m . The higher ξ , the more likely the agent is to choose the option that is more likely according to its current belief. Since m is a deterministic function of λ , we can write $p(y | m, \xi) = p(y | \lambda, \xi)$.

In general, the joint distribution for observations (i.e., decisions) and parameters of an HGF-based decision model takes the form

$$p(y, \chi, \lambda^{(0)}, \xi | u) = p(\chi, \lambda^{(0)}, \xi) \prod_{k=1}^K p(y^{(k)} | \lambda^{(k)}(\chi, \lambda^{(0)}, u), \xi) \quad (19)$$

where $u \stackrel{\text{def}}{=} \{u^{(1)}, \dots, u^{(K)}\}$ and $y \stackrel{\text{def}}{=} \{y^{(1)}, \dots, y^{(K)}\}$ are the inputs and responses at time points $k = 1$ to $k = K$, respectively, and $\lambda^{(k)}(\chi, \lambda^{(0)}, u)$ are the sufficient statistics $\lambda^{(k)} \stackrel{\text{def}}{=} \{\mu_1^{(k)}, \sigma_1^{(k)}, \dots, \mu_n^{(k)}, \sigma_n^{(k)}\}$ of the hidden states of the HGF at time k . The inputs u are given because the agent and its observer both know them and the agent uses them to invert its HGF model, resulting in the trajectories $\lambda^{(k)}$. This makes it clear that the decision model is not a generative model of sensory input, while the HGF is. Strictly speaking, the decision model does not use the perceptual model directly. Instead, the decision model uses the perceptual model indirectly via its inversion, where input is given. It is also noteworthy that the sets of inputs u and observations y are finite here ($k = 1, \dots, K$), while the HGF is open-ended (cf. $k = 1, 2, \dots$ in Equation 1).

The goal now is to find an expression for the maximum-a-posteriori (MAP) estimate for the parameters $\xi \stackrel{\text{def}}{=} \{\chi, \lambda^{(0)}, \xi\}$. The MAP estimate ξ^* of ξ is defined as

$$\xi^* \stackrel{\text{def}}{=} \arg \max_{\xi} p(\xi | y, u), \quad (20)$$

We unpack this in Appendix E to make it tractable:

$$\xi^* = \arg \max_{\xi} \left(\sum_k \ln p(y^{(k)} | \lambda^{(k)}(\chi, \lambda^{(0)}, u), \xi) + \ln p(\xi) \right) \quad (21)$$

The objective function $Z(\xi)$ that needs to be maximized is therefore the log-joint probability density of the parameters ξ and responses y given inputs u :

$$Z(\xi | u, y) \stackrel{\text{def}}{=} \ln p(\xi, y | u) = \sum_k \ln p(y^{(k)} | \lambda^{(k)}(\chi, \lambda^{(0)}, u), \xi) + \ln p(\xi) \quad (22)$$

While the response model gives $p(y^{(k)} | \lambda^{(k)}, \xi)$, the perceptual model (i.e., the HGF) provides the representations $\lambda^{(k)}(\chi, \lambda^{(0)}, u)$ by way of its update equations. The last missing part in Equation 22 is the prior distribution $p(\xi)$. This will be discussed below. There are many alternative optimization procedures to implement the maximization in Equation 21. We have compared four in the simulations discussed below.

Finally, one important point in relation to model inversion is model identifiability which we discuss in detail in Appendix F. In brief, when the posterior mean of the state at level i (i.e., μ_i) is included in the response model and thus affects measured behavior, all quantities at that level $\mu_i^{(0)}$, κ_i , and ω_i can be estimated. If μ_i is not included in the response model, it is advisable to fix two of the three parameters $\mu_i^{(0)}$, κ_i , and ω_i , reflecting a particular choice of origin and scale on x_i . This avoids an overparameterized model.

PRIORS AND TRANSFORMED PARAMETER SPACES

A crucial part of Bayesian inference is the specification of a prior distribution, in our case $p(\xi)$. There is no principled reason why the priors on the different elements of ξ should not be independent; therefore, we may assume the following factorization:

$$p(\xi) = p(\vartheta) \prod_i p(\kappa_i) p(\omega_i) p(\mu_i^{(0)}) p(\sigma_i^{(0)}) \prod_j p(\zeta_j). \quad (23)$$

$p(\zeta_j)$ depends on the response model chosen and will have to be dealt with on a case-by-case basis (see below), but the remaining marginal priors are generic and will be discussed in what follows.

The most straightforward case are the priors on ω_i . Since ω_i can take values on the whole real line, it can be estimated in its native space with a (possibly wide) Gaussian prior:

$$\omega_i \sim \mathcal{N}(\mu_{\omega_i}, \sigma_{\omega_i}). \quad (24)$$

The same applies to $\mu_i^{(0)}$:

$$\mu_i^{(0)} \sim \mathcal{N}(\mu_{\mu_i^{(0)}}, \sigma_{\mu_i^{(0)}}). \quad (25)$$

$\sigma_i^{(0)}$ has a natural lower bound at zero since it is a variance. We can avoid non-positive values by estimating $\sigma_i^{(0)}$ in log-space. That is, we use a log-Gaussian prior:

$$\ln \sigma_i^{(0)} \sim \mathcal{N}(\mu_{\ln \sigma_i^{(0)}}, \sigma_{\ln \sigma_i^{(0)}}). \quad (26)$$

Just like $\sigma_i^{(0)}$, ϑ is a variance and has a lower bound at zero. In addition to the lower bound, it is desirable to have an upper bound on ϑ because, for a ϑ too large, the assumptions underlying the derivation of the update equations of the HGF no longer hold. Specifically, for large values of ϑ it is possible to get updates that push the precision π_n at the top level below zero, indicating that the agent knows “less than nothing” about x_n . In less extreme cases, a large ϑ may allow μ_n to jump to very high levels, giving rise to improbable inference trajectories (cf. Equations 9 and 11). This is due to a violation of the assumption that the variational energies $I(x_i)$ are nearly quadratic (see Mathys et al., 2011, for details).

To avoid such violations, it is sensible to place an upper bound on ϑ in addition to the lower bound at zero. This can be achieved by estimating ϑ in “logit-space,” a logistic sigmoid transformation of native space with a variable upper bound $a > 0$:

$$\begin{aligned} \text{logit}_a(x) &\stackrel{\text{def}}{=} \ln \left(\frac{x}{a-x} \right); \\ \Rightarrow x &= \frac{a}{1 + \exp(-\text{logit}_a(x))} \end{aligned} \quad (27)$$

In that space, the prior on ϑ can then be taken as

$$\text{logit}_{a_\vartheta} \vartheta \sim \mathcal{N}(\mu_{\text{logit}_{a_\vartheta} \vartheta}, \sigma_{\text{logit}_{a_\vartheta} \vartheta}). \quad (28)$$

While κ_i can in principle take any real value, flipping the sign of κ_i is equivalent to flipping that of x_{i+1} (cf. Equation 5). It is therefore useful to adopt the convention that all $\kappa_i > 0$. This leads to

the intuitive relation that a greater x_{i+1} means a greater variability in x_i ; in other words, this makes f_i in Equation 3 a monotonically increasing function. A second useful constraint on κ_i is that it is bounded above, for the same reason as ϑ . Consequently, we evaluate κ_i in logit-space with the following priors:

$$\text{logit}_{a_{\kappa_i}} \kappa_i \sim \mathcal{N}(\mu_{\text{logit}_{a_{\kappa_i}} \kappa_i}, \sigma_{\text{logit}_{a_{\kappa_i}} \kappa_i}). \quad (29)$$

The exact specification of the above priors can vary, depending on the experimental context and instructions given to the subject (e.g., whether or not to expect a volatile environment). In most cases, a choice of κ_i and ϑ with upper bounds at or below 2 will be sensible. In cases where there is little movement in x_n (the topmost x), a choice of ϑ closer to 0 than 1 will be appropriate. Notably, given that choosing a different prior amounts to having a different model, the choice between alternative priors can be evaluated using model comparison (cf. Stephan et al., 2009).

EXAMPLES AND SIMULATIONS

CATEGORICAL OUTCOMES AND SENSORY UNCERTAINTY

In the formulation above, x_1 performs a Gaussian random walk on a continuous scale. However, states of an agent’s environment that generate sensory input are often categorical, in the simplest case binary (e.g., present/absent). This fact can be accommodated in the perceptual model above by making the base level, x_1 , binary (we omit time indices k unless they are needed to avoid confusion):

$$x_1 \in \{0, 1\}. \quad (30)$$

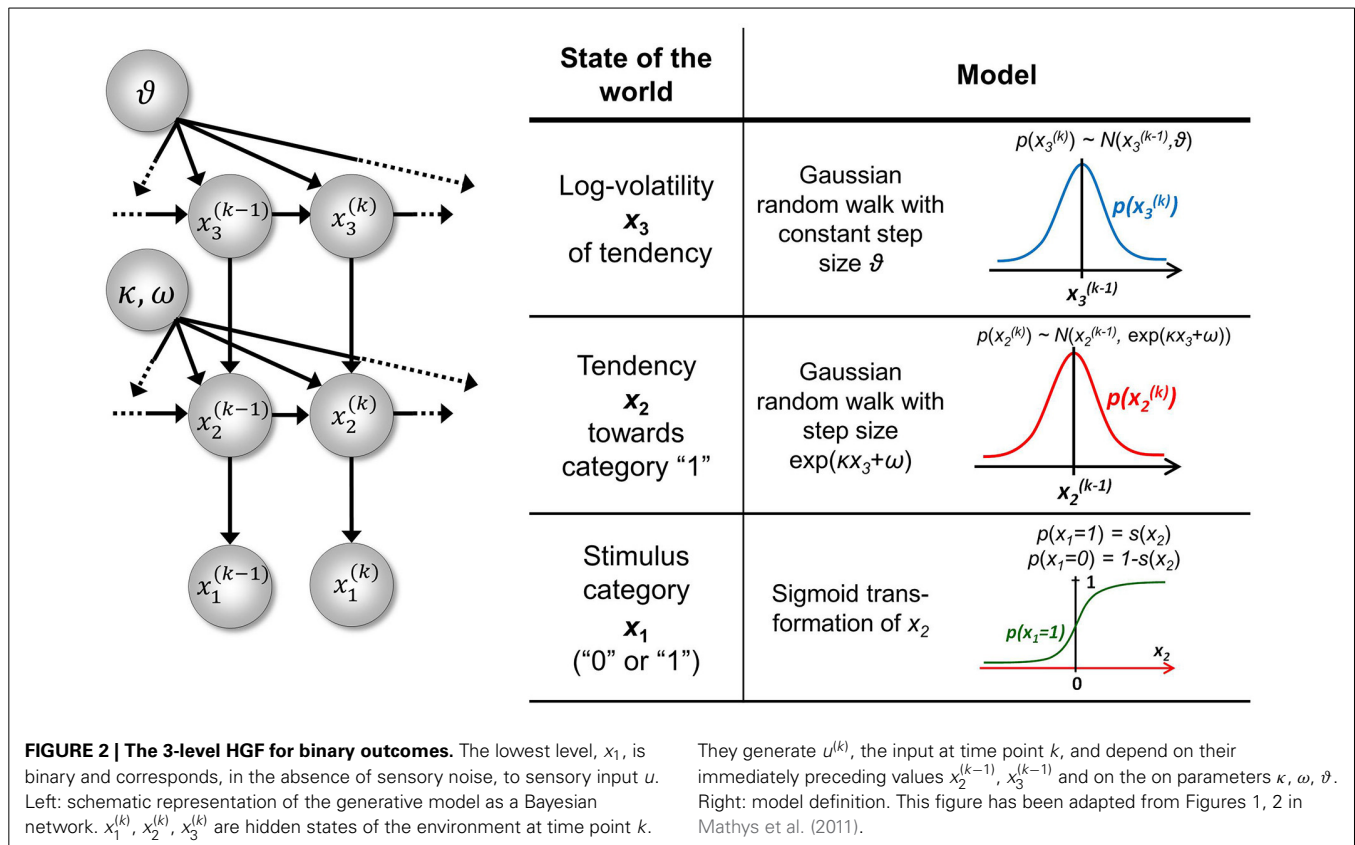
The second level, x_2 , of the model then describes the tendency of x_1 toward state 1:

$$p(x_1 | x_2) = s(x_2)^{x_1} (1 - s(x_2))^{1-x_1} = \text{Bernoulli}(x_1; s(x_2)) \quad (31)$$

where $s(x_2) \stackrel{\text{def}}{=} \frac{1}{1 + \exp(-x_2)}$ is the logistic sigmoid function. Similarly, when x_1 represents $d > 2$ outcomes, the probability of each can be represented by its own x_2 , performing (at most) $d - 1$ independent random walks. $p(x_1 = 0)$ simply is $1 - p(x_1 = 1)$.

In the three-level HGF for binary outcomes (Mathys et al., 2011) the third level, x_3 is at the top, with constant step variance ϑ . The only level with a coupling of the form of Equation 5 is therefore the second level; this allows us to write $\kappa_2 \equiv \kappa$ and $\omega_2 \equiv \omega$. We can allow for sensory uncertainty by including an additional level at the bottom of the hierarchy that predicts sensory input u from the state x_1 . In the absence of sensory uncertainty, knowledge of the state x_1 enables accurate prediction of input u and vice versa; we may then simply set $u \equiv x_1$ and treat x_1 as if it were directly observed. A graphical overview of this model is given in **Figure 2**.

For the particular case of the three-level HGF for binary outcomes, the general update equations in Equations 9–14 (with $t^{(k)} = 1$ for all k) take the following specific form (as previously derived in (Mathys et al., 2011), with additional detail in Appendix D):



$$\mu_2^{(k)} = \mu_2^{(k-1)} + \frac{1}{\pi_2^{(k)}} \delta_1^{(k)} \quad (32)$$

$$\pi_2^{(k)} = \hat{\pi}_2^{(k)} + \frac{1}{\hat{\pi}_1^{(k)}}. \quad (33)$$

with

$$\hat{\mu}_1^{(k)} \stackrel{\text{def}}{=} s(\mu_2^{(k-1)}) \quad (34)$$

$$\delta_1^{(k)} \stackrel{\text{def}}{=} \mu_1^{(k)} - \hat{\mu}_1^{(k)} \quad (35)$$

$$\hat{\pi}_1^{(k)} \stackrel{\text{def}}{=} \frac{1}{\hat{\mu}_1^{(k)} (1 - \hat{\mu}_1^{(k)})} \quad (36)$$

$$\hat{\pi}_2^{(k)} \stackrel{\text{def}}{=} \frac{1}{\sigma_2^{(k-1)} + t^{(k)} e^{\kappa \mu_3^{(k-1)} + \omega}} \quad (37)$$

Details of the derivation of these update equations are given in Appendix D. The update equations for binary outcomes differ from those given in Equations 9 and 10 only at the second level. On all higher levels, they are the generic HGF updates from Equations 9 and 10. This difference is entirely due to the sigmoid transformation that links the first and second level, enabling the filtering of binary outcomes. Note that in the binary case, the second level corresponds to the first level of the continuous case in the sense that they are the lowest levels where a Gaussian random walk takes place.

To illustrate of how the HGF can deal with the simplest kind of informational uncertainty, sensory uncertainty, we simulate two agents, one with high and the other with low sensory uncertainty, who are otherwise equal. Sensory uncertainty is captured by the following relation between the binary state x_1 and sensory input u :

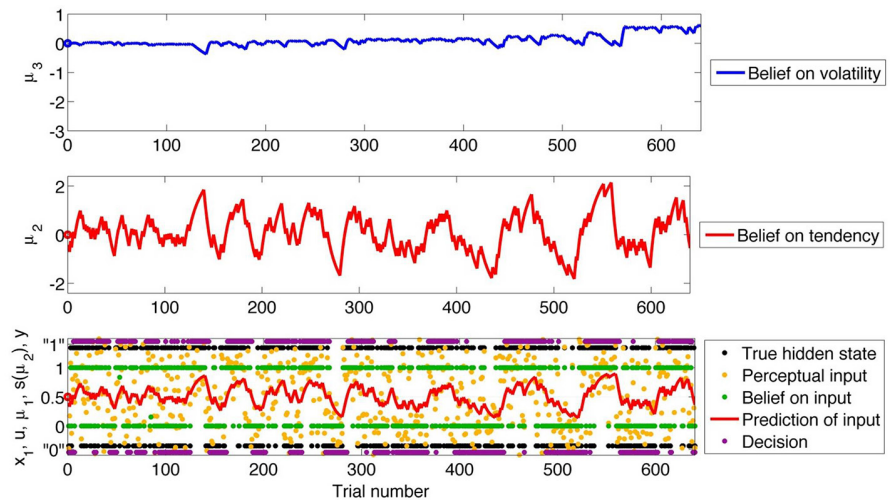
$$u \sim \begin{cases} \mathcal{N}(\eta_1, \hat{\pi}_u^{-1}) & \text{if } x_1 = 1 \\ \mathcal{N}(\eta_0, \hat{\pi}_u^{-1}) & \text{if } x_1 = 0 \end{cases} \quad (38)$$

This means that the probability of u is Gaussian with precision $\hat{\pi}_u$ around a mean of η_1 for $x_1 = 1$ and η_0 for $x_1 = 0$. In this case (cf. Mathys et al., 2011, Equation 47), the update equation for μ_1 is

$$\mu_1^{(k)} = \frac{\hat{\mu}_1^{(k)} \exp\left(-\frac{\hat{\pi}_u}{2} (u^{(k)} - \eta_1)^2\right)}{\hat{\mu}_1^{(k)} \exp\left(-\frac{\hat{\pi}_u}{2} (u^{(k)} - \eta_1)^2\right) + (1 - \hat{\mu}_1^{(k)}) \exp\left(-\frac{\hat{\pi}_u}{2} (u^{(k)} - \eta_0)^2\right)} \quad (39)$$

Figure 3 shows our simulation. We first chose a sequence of true hidden states $x_1^{(k)}$, $k = 1, \dots, 640$. We then drew a sequence of inputs $u^{(k)} \sim \mathcal{N}(x_1^{(k)}, 0.1)$, corresponding to a setting of $\eta_1 = 1$, $\eta_0 = 0$, and $\hat{\pi}_u = 10$ (cf. Equation 38). These inputs were fed into two three-level HGFs that differed only in the amount of sensory uncertainty they assumed: $\hat{\pi}_u^{-1} =$

Low sensory
uncertainty



High sensory
uncertainty




FIGURE 3 | The consequences of sensory uncertainty. Simulation of inference on a binary hidden state x_1 (black dots) using a three-level HGF under low ($\hat{\pi}_u = 1000$, top panel) and high ($\hat{\pi}_u = 10$, bottom panel) sensory uncertainty.

Trajectories were simulated using the same input and parameters (except $\hat{\pi}_u$) in both cases: $\mu_2^{(0)} = \mu_3^{(0)} = 0$, $\sigma_2^{(0)} = \sigma_3^{(0)} = 1$, $\kappa = 1$, $\omega = -3$, and $\vartheta = 0.7$. Decisions were simulated using a unit-square sigmoid model with $\zeta = 8$.

0.001 (low) and $\hat{\pi}_u^{-1} = 0.1$ (high). Clearly, higher input precision leads to greater responsiveness to fluctuations in input, as reflected in the trajectory of belief on tendency x_2 and in a higher volatility estimate (belief on x_3). In the case of low input precision, the volatility estimate keeps declining because most of the variation in input is attributed to noise instead of fluctuations in the underlying tendency toward one outcome category x_1 or the other. Decisions (purple dots) were simulated using a unit-square sigmoid response model with $\zeta = 8$ in both cases (cf. Equation 18). The consequences of high sensory uncertainty for decision-making in this scenario are apparent: the agent with higher sensory uncertainty is less consistent in favoring one option over the other at any given time. This accords well with recent accounts of psychopathological symptoms as a failure of sensory attenuation (Adams et al., 2013).

We now turn to the third and fourth elements defining our Bayesian agent: decision models based on loss functions.

DECISION MODEL FOR A SIMPLE BINARY LOSS FUNCTION

One of the simplest decision situations for an agent is having to choose between two options, only one of which will be rewarded, but both of which offer the same gain (i.e., negative loss), if rewarded. In the three-level version of the HGF from **Figure 2**, we may code one such binary outcome as $x_1 = 1$ and the other as $x_1 = 0$. This allows us to define a quadratic loss function ℓ where making the wrong choice $y \in \{0, 1\}$ leads to a loss of 1 while the right choice leads to a loss of 0:

$$\ell(x_1, y) = (x_1 - y)^2 \quad (40)$$

The expected loss \mathcal{Q} of decision y , given the agent's representations λ , is then the expectation of ℓ under the distributions q described by λ :

$$\mathcal{Q}(y; \lambda) = \langle \ell(x_1, y) \rangle_{q(x_1; \lambda)} = \sum_{x_1 \in \{0, 1\}} \ell(x_1, y) \cdot q(x_1; \lambda) \quad (41)$$

To evaluate this, we must remember that the agent has to rely on its beliefs deriving from time $k-1$ to make decision $y^{(k)}$ at time k . In the above equation, elements of λ therefore have time index $k-1$, while x_1 and y have time index k . Specifically, the belief on the outcome probability at the first level is $\hat{\mu}_1^{(k)} = s(\mu_2^{(k-1)})$. With $q(x_1) = (\mu_1)^{x_1}(1-\mu_1)^{1-x_1}$ (Mathys et al., 2011, Equation 12), we then have

$$Q(y=0; \lambda) = \hat{\mu}_1; \quad Q(y=1; \lambda) = 1 - \hat{\mu}_1 \quad (42)$$

The optimal decision y^* is the one that minimizes expected loss Q :

$$y^* \stackrel{\text{def}}{=} \arg \min_y Q(y; \lambda) = \begin{cases} 1 & \text{if } \hat{\mu}_1 > \frac{1}{2} \\ 0 & \text{if } \hat{\mu}_1 < \frac{1}{2} \\ \{0, 1\} & \text{if } \hat{\mu}_1 = \frac{1}{2} \end{cases}$$

$$= \lim_{\zeta \rightarrow \infty} \frac{\hat{\mu}_1^\zeta}{\hat{\mu}_1^\zeta + (1 - \hat{\mu}_1)^\zeta} \quad (43)$$

This simply means that to minimize its losses, the agent should choose the option it believes more likely to be rewarded. It may seem superfluous to go to such lengths to derive such an obvious result, but the purpose of the above is also to give an illustration of the principled way a decision rule can be derived by combining the HGF with a loss function.


It is, however, unreasonable to assume that human agents will always choose the option that minimizes their expected loss in the current trial, for two reasons. First, if there is more than one trial and the probabilities of the different options are independent, there is an exploration/exploitation tradeoff that makes it worth the agent's while (in the long run) sometimes to choose an option that is not expected to minimize loss in the current trial (Macready and Wolpert, 1998; Daw et al., 2006). Second, biological agents exhibit decision noise (Faisal et al., 2008), e.g., owing to implementation constraints at the molecular, synaptic or circuit level. To allow for exploration and noise, we use a decision model that corresponds to the right-hand side of Equation 43, without taking the limit, instead leaving ζ as a parameter to be estimated from the data (cf. Equation 18):

$$p(y|\lambda, \zeta) = \left(\frac{\hat{\mu}_1^\zeta}{\hat{\mu}_1^\zeta + (1 - \hat{\mu}_1)^\zeta} \right)^y \cdot \left(\frac{(1 - \hat{\mu}_1)^\zeta}{\hat{\mu}_1^\zeta + (1 - \hat{\mu}_1)^\zeta} \right)^{1-y} \quad (44)$$

Figure 4 contains a graph of this function for $p(y=1)$ where ζ plays the role of the noise (or exploration) parameter. This decision model was the basis for the simulations we conducted to assess the accuracy of parameter estimations (results below).

INVERSION EXAMPLE

To illustrate how real datasets can be inverted and different response models compared, we take the data of one subject from Iglesias et al. (2013). This consists of 320 inputs u and responses y . Our perceptual model is the three-level HGF for binary outcomes without sensory noise, and a first choice of decision model is the unit-square sigmoid of Equation 44. Using the HGF Toolbox



(<http://www.translationalneuromodeling.com/tapas>), we specify the following priors (mean, variance) in appropriately transformed spaces: $\mu^{(0)}_2 : (0, 1)$, $\sigma_2 : (0, 1)$, $\mu^{(0)}_3 : (1, 0)$, $\sigma_3^{(0)} : (0, 1)$, $\kappa : (0, 2)$, $\omega : (-4, 0)$, $\vartheta : (0, 2)$, and $\zeta : (48, 1)$. The variance of 0 on $\mu^{(0)}_3$ and ω fixes these parameters to 1 and -4 , respectively. The spaces for $\sigma_i^{(0)}$ and ζ were log-transformed while κ was estimated in a logit-transformed space with upper bound 6, and ϑ was estimated in logit-transformed space with upper bound 0.005.

We now modify our response model so that it no longer has a constant free parameter (ζ) as its inverse decision temperature, but the inverse volatility estimate $\exp(-\mu_3)$:

$$p(y|\lambda) = \left(\frac{\hat{\mu}_1^{\exp(-\mu_3)}}{\hat{\mu}_1^{\exp(-\mu_3)} + (1 - \hat{\mu}_1)^{\exp(-\mu_3)}} \right)^y \cdot \left(\frac{(1 - \hat{\mu}_1)^{\exp(-\mu_3)}}{\hat{\mu}_1^{\exp(-\mu_3)} + (1 - \hat{\mu}_1)^{\exp(-\mu_3)}} \right)^{1-y} \quad (45)$$

This means that the agent will behave the less deterministically the more volatile it believes its environment to be. Since this is now a decision model that contains μ_3 , it permits us to estimate all parameters including $\mu_3^{(0)}$ and ω . Accordingly, we increase the variance of their priors to 4. The result of this inversion is shown in **Figure 5**. This figure illustrates how the HGF deals with perceptual uncertainty by updating beliefs throughout its hierarchy on the basis of precision-weighted prediction errors. The learning rates ψ_i (definitions see figure caption) are adjusted continually at each level separately, which provides the flexibility needed to adapt to changes in outcome tendency and volatility (i.e., to perceptual uncertainty).

Under the Laplace assumption (Friston et al., 2007), the negative variational free energy, an approximation to the log-model evidence, is -196.19 for the first response model and -188.84 for the second. This corresponds to a Bayes factor of $\exp(-188.84 - (-196.19)) = 1556$, giving the second model a decisive advantage despite the fact that it contains an additional




FIGURE 5 | Model inversion. Maximum-a-posteriori parameter estimates are $\mu_2^{(0)} = 0.87$, $\sigma_2^{(0)} = 1.20$, $\mu_3^{(0)} = -0.65$, $\sigma_3^{(0)} = 0.88$, $\kappa = 1.32$, $\omega = -0.71$, and $\nu = 0.0023$. These parameter values correspond to the following trajectories: **(A)** Posterior expectation μ_3 of log-volatility x_3 . **(B)** Precision weight $\psi_3 \stackrel{\text{def}}{=} \frac{\pi_2}{\pi_3}$ which modulates the impact of prediction error δ_2 on log-volatility updates. **(C)** Volatility prediction error δ_2 . **(D)** Posterior expectation μ_2 of tendency μ_2 . **(E)** Precision weight $\psi_2 \stackrel{\text{def}}{=} \pi_2^{-1}$ (in green) which modulates the impact of input prediction error δ_1 on μ_2 . Since μ_2

is in logit space, the function of σ_2 as a dynamic learning rate is more easily visible after transformation to x_1 -space. This results in the red line labeled $q(\psi_2) \stackrel{\text{def}}{=} s'(\mu_2) \cdot \psi_2$. **(F)** Prediction error δ_1 about input u . (In Iglesias et al., 2013, Figures S1 and S2, δ_1 is defined as an outcome prediction error, which corresponds to the absolute value of δ_1 as defined here). **(G)** Black: true probability of input 1. Red: posterior expectation of input $u = 1$, $\hat{\mu}_1$; this corresponds to a sigmoid transformation of μ_2 in **(E)**. Green: sensory input. Orange: subject's observed decisions.

free parameter. In this example, including a measure of higher-level uncertainty has clearly improved our model of a subject's learning and decision-making.

Another possible choice for the inverse decision temperature is $\hat{\pi}_2$ (cf. Paliwal et al., 2014). This choice is interesting because it is similar to the hypothesis (Friston et al., 2013; Schwartenbeck et al., 2013) that precision serves as the inverse decision temperature in active inference. With a negative variational free energy of -189.63 , this model performs similarly to the one with $\exp(-\mu_3)$ as the inverse decision temperature. This similarity in performance is not surprising since $\hat{\pi}_2$ is (inversely) driven to a large extent by μ_3 (cf. Equations 11 and 13).

DECISION MODEL FOR A ONE-ARMED BANDIT

As an additional example, we discuss a more complex binary decision task that we used to collect data from human subjects (Cole et al., in preparation). In this variant of a one-armed bandit experiment, subjects were asked to play a series of gambles with the goal of maximizing their overall score (Figure 6). On each trial, subjects chose between two options represented by the same two

fractals, which had different and time-varying reward probabilities. At any point in time, these probabilities summed to unity, implying that exactly one of the two options would be rewarded. Although subjects knew that probabilities varied throughout the course of the experiment, they were not told the schedule that governed these changes. The schedule included both a period of low volatility and a period of high volatility (Figure 6), similar to the task used by Behrens et al. (2007).

In order to encourage subjects to switch options above and beyond normal exploration behavior (Macready and Wolpert, 1998), the two cards were associated with varying reward magnitudes. On each trial, magnitudes were drawn from a discrete uniform distribution $\mathcal{U}(1, 9)$ (i.e., rewards would take values from the range $\{1, 2, \dots, 9\}$ with equal probability).

Subjects began the experiment with an initial score of 0 points. Once a card had been chosen, if that card was rewarded, the associated reward would be added to the current score. The final score at the end of the experiment was translated into monetary reimbursement. The experiment consisted of 160 trials.

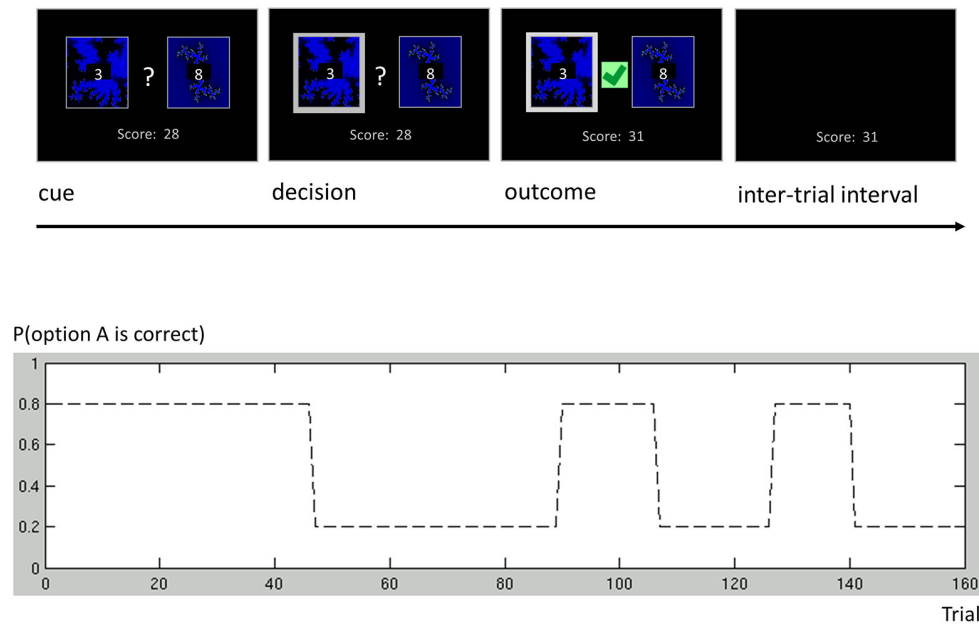


FIGURE 6 | One-armed bandit task. Participants were engaged in a simple decision-making task. Each trial consisted of four phases. (i) Cue phase. Two cards and their costs were displayed. (ii) Decision phase. Once the subject had made a decision, the chosen card was highlighted. (iii) Outcome phase. The outcome of a decision was displayed, and added to the score bar only if

the chosen card was rewarded. (iv) Inter-trial interval (ITI). The screen only showed the score bar, until the beginning of the next trial. Our experimental paradigm consisted of a number of phases with different reward structures. Different phase lengths induced both a phase of low volatility (trials 1 through 90) and a phase of high volatility (trials 91 through 160).

Calling the two fractals A and B, we parameterize the agent's response by

$$y = \begin{cases} 0 & \text{for choice A} \\ 1 & \text{for choice B} \end{cases} \quad (46)$$

Correspondingly, the state x_1 is

$$x_1 = \begin{cases} 0 & \text{if A rewarded} \\ 1 & \text{if B rewarded} \end{cases} \quad (47)$$

Taking r_A and r_B to be the rewards for A and B, respectively, we introduce the quadratic loss function

$$\begin{aligned} \ell(x, y) &= -(1 - x_1)(y + x_1 - 1)^2 \cdot r_A - x_1(y + x_1 - 1)^2 \cdot r_B \\ &= -(y + x_1 - 1)^2((r_B - r_A)x_1 + r_A) \end{aligned} \quad (48)$$

This corresponds to the following loss table:

	x_1	
	0	1
y	0	1
	$-r_A$	0
	1	0
		$-r_B$

(49)

Following the same procedure as above, we get:

$$\begin{aligned} Q(y; \lambda) &= \sum_{x_1 \in \{0,1\}} \ell(x_1, y) \cdot q(x_1; \lambda) \\ &= -y^2 r_B \hat{\mu}_1 - (1 - y)^2 r_A (1 - \hat{\mu}_1) \\ &= \begin{cases} -r_B \hat{\mu}_1 & \text{for } y = 1, \\ -r_A (1 - \hat{\mu}_1) & \text{for } y = 0. \end{cases} \end{aligned} \quad (50)$$

With the expected loss from each option on a continuous scale, a simple but powerful decision model is the softmax rule (Sutton and Barto, 1998; Daw et al., 2006)

$$p(y_i | \lambda, \zeta) = \frac{\exp(\zeta \cdot Q(y_i; \lambda))}{\sum_j \exp(\zeta \cdot Q(y_j; \lambda))}, \quad (51)$$

where y_i is one particular option and the sum runs over all options. This means that the decision probabilities are Boltzmann-distributed according to their expected rewards (i.e., their expected negative losses) with the parameter ζ serving as the analog of inverse temperature. In our binary case, this evaluates to

$$\begin{aligned} p(y = 1 | \lambda, \zeta) &= \frac{\exp(\zeta(r_B \hat{\mu}_1 - r_A(1 - \hat{\mu}_1)))}{\exp(\zeta(r_B \hat{\mu}_1 - r_A(1 - \hat{\mu}_1))) + \exp(\zeta(r_A(1 - \hat{\mu}_1) - r_B \hat{\mu}_1))}, \\ p(y = 0 | \lambda, \zeta) &= \frac{\exp(\zeta(r_A(1 - \hat{\mu}_1) - r_B \hat{\mu}_1))}{\exp(\zeta(r_B \hat{\mu}_1 - r_A(1 - \hat{\mu}_1))) + \exp(\zeta(r_A(1 - \hat{\mu}_1) - r_B \hat{\mu}_1))}. \end{aligned} \quad (52)$$

This is a logistic sigmoid function of the difference $r_B \hat{\mu}_1 - r_A(1 - \hat{\mu}_1)$ of expected reward for choice B minus expected

reward for choice A. If the expected reward of choice B exceeds that of choice A, the likelihood of choice B is greater than half and vice versa.

SIMULATION STUDY

Given the nontrivial nature of our model, it is important to verify the robustness of the inversion scheme. This, in turn, may depend on the numerical optimization method employed. To assess our ability to estimate the parameters under different optimization schemes, we conducted a systematic simulation study based on the 3-level HGF for binary outcomes. This model is shown graphically in **Figure 2** and was the basis for the studies of Vossel et al. (2013) and Iglesias et al. (2013). κ was chosen as the perceptual parameter to vary because of the interesting effects it has on the nature of the inferential process (cf. Mathys et al., 2011). The response parameter ζ was chosen as the second parameter to vary because it represents inverse response noise (cf. Equation 20), i.e., for lower values of ζ the mapping from beliefs to responses becomes less deterministic, which renders it more difficult to estimate the perceptual parameters.

Simulations took place in four steps:

1. We chose a particular sequence of 320 binary input $u = \{u^{(1)}, \dots, u^{(320)}\}$; this was the input sequence in a recent study using the HGF (Iglesias et al., 2013).
2. We chose a particular set of values for the parameters ξ .
3. We generated 320 binary responses $y = \{y^{(1)}, \dots, y^{(320)}\}$ by drawing from the response distribution given by Equation 44 below.
4. We estimated ξ^* according to Equation 21.

Step 1 was only performed once, so that u was the same in all simulations. The values of ξ in step 2 were constant for all parameters except κ and ζ . The values of κ and ζ were taken from a two-dimensional grid in which the κ dimension took the values $\{0.5, 1, 1.5, \dots, 3.5\}$ while the ζ dimension took the values $\{0.5, 1, 6, 24\}$. Steps 3 and 4 were then repeated 1'000 times for each value pair on the $\{\kappa, \zeta\}$ grid (for MCMC, owing to its computational burden, only 100 estimations were performed). The ζ values on the grid were chosen such that they covered the whole range from very low ($\zeta = 24$) to very high response noise ($\zeta = 0.5$, cf. **Figure 4**). The κ values were chosen to cover the range observed in an empirical behavioral study using the same inputs u (Iglesias et al., 2013). The remaining model parameters were held constant ($\omega = -4$, $\vartheta = 0.0025$). In total, six parameters were estimated. These were (with the space they were estimated in and prior mean and variance in that space in brackets): $\mu_2^{(0)}$ (native, 0, 1), $\sigma_2^{(0)}$ (log, 0, 1), $\sigma_3^{(0)}$ (log, 0, 1), κ (logit with upper bound at 6, 0, 9), ϑ (logit with upper bound at 0.005, 0, 9), and ζ (log, 48, 1). The prior mean of the response variable ζ was chosen relatively high to provide shrinkage on the estimation of decision noise.

This procedure was repeated for four different optimization methods which are commonly used but possess different properties with regard to computational efficiency and robustness to getting trapped in local extrema:

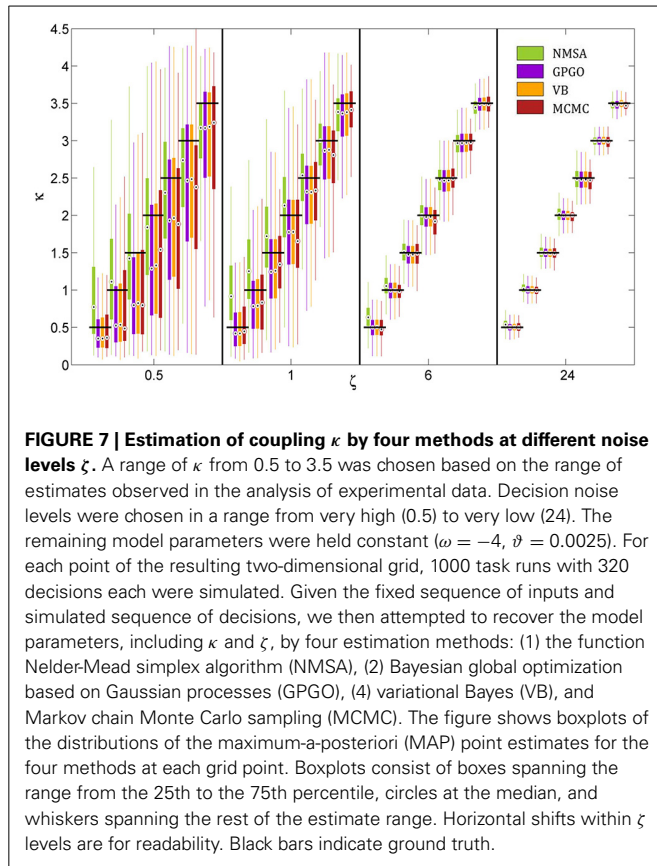
1. Nelder-Mead simplex algorithm (NMSA),
2. Gaussian process-based global optimization (GPGO),
3. Variational Bayes (VB),
4. Markov Chain Monte Carlo estimation (MCMC).

In brief, NMSA (Nelder and Mead, 1965) is a popular local optimization algorithm which is implemented, for example, in the `fminsearch` function of Matlab. VB also optimizes locally (by gradient descent); for details see Bishop (2006, p. 461ff). For our simulation study, we used VB as implemented in the DAVB toolbox, available at <http://goo.gl/As8p7> (Daunizeau et al., 2009, 2014). In contrast, GPGO (Rasmussen and Williams, 2006; Lomakina et al., 2012) provides a global optimum of the objective function and is thus potentially more robust than NMSA and VB albeit computationally more expensive. The final method was MCMC (Gelman et al., 2003, p. 283ff) which served as a “gold standard” against which we compared the other methods. Specifically, we used Gibbs sampling with a one-dimensional Metropolis step for each of the parameters (cf. Gelman et al., 2003, p. 292). For each of the 100 simulation runs (at each point on our parameter grid) we used one chain with a length of 500'000 samples and a burn-in period of 25'000 samples. In summary, our simulations thus consider two algorithms (NMSA and VB) which are computationally very efficient but provide a local optimum only, in comparison to another two algorithms (GPGO and MCMC) which are computationally more expensive but are capable of finding global optima.

All optimization methods could reliably distinguish different values of κ at low or moderate decision noise (**Figure 7**). At higher noise levels, estimates became less reliable. With GP, VB, and MCMC, they then exhibited a tendency to underestimate κ , while NMSA tended to mid-range values. Nonetheless, substantial differences in κ within the range tested could be detected by all four methods even at high levels of noise.

The noise level itself could also be determined by all four methods (**Figure 8**). The methods did not differ appreciably in their performance. They all tended to underestimate the noise level owing to a mild shrinkage due to the prior on ζ . Errors are smaller for moderate noise levels, increasing for both high and low noise (cf. **Figure 9A2**).

Figures 9A1,A2 shows the root mean squared error in κ and $\log(\zeta)$, jointly for all values of κ . The results show that the noise level could best be estimated at moderate levels, where in fact most estimates of experimental data are found. Again, the methods perform comparably well, with NMSA best at high noise. **Figures 9B,B2** contrasts the performance of VB and MCMC and displays the accuracy of the confidence with which VB and MCMC make their estimates. To this end, it uses the fact that VB and MCMC estimate the whole posterior distribution. Parameter estimates can therefore not only be summarized as point estimates, but also as posterior central intervals (PCIs; the 95% PCI is the interval that excludes 2.5% of the posterior probability mass on either side). If an estimation method were neither over- nor underconfident, 95 of 95% PCIs would contain the true parameter value. If the proportion is less than 0.95, this indicates overconfidence; if it is greater than 0.95, underconfidence. Both methods were realistically confident about their inference



on κ across noise levels, with a slight tendency toward overconfidence with higher noise. This tendency was more pronounced with estimates of ζ .

DISCUSSION

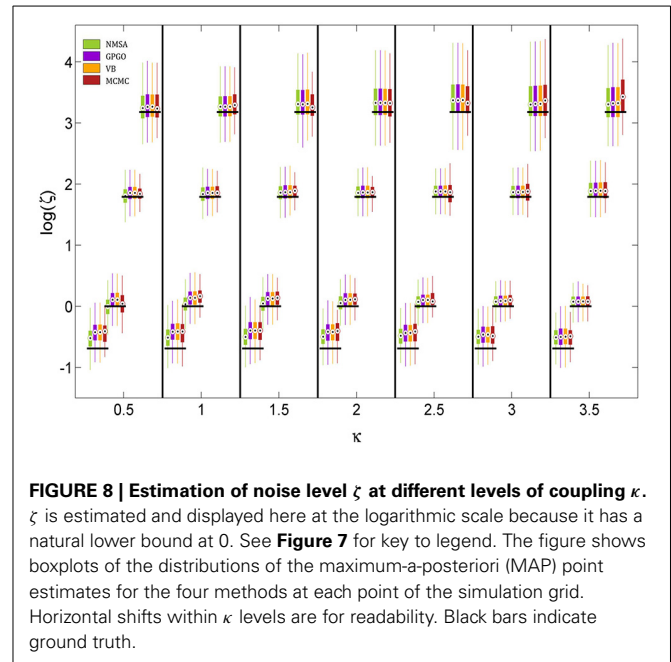
In this paper, we have shown that the hierarchical Bayesian model of Mathys et al. (2011) can be extended in several ways, resulting in a general framework referred to as the HGF. Furthermore, we have demonstrated how the HGF can be combined with decision models to allow for parameter estimation from empirical data. We start by discussing the nature of the HGF updates in the context of Bayesian inference.

A crucial feature of the HGF's update equations is emphasized by the notation used in Equation 9: the updates of the means are *precision-weighted prediction errors*. For a full understanding of their role, we will first discuss Bayesian updates in the simplest possible case, where they can be calculated exactly. In this simplest case, there is only one hidden state $x \in \mathbb{R}$ that is the target of our inference, and there is a Gaussian prior on x :

$$x \sim \mathcal{N}(\mu_x, \pi_x^{-1}), \quad (53)$$

where μ_x is the mean and π_x the precision. The likelihood of x (i.e., the probability of observing the datum $u \in \mathbb{R}$ given x) is also Gaussian, with precision (inverse observation noise) π_u :

$$u \sim \mathcal{N}(x, \pi_u^{-1}). \quad (54)$$



According to Bayes' theorem, the posterior is now also Gaussian:

$$x | u \sim \mathcal{N}(\mu_{x|u}, \pi_{x|u}^{-1}). \quad (55)$$

The posterior precision $\pi_{x|u}$ and mean $\mu_{x|u}$ can be written as the following analytical and exact one-step updates:

$$\mu_{x|u} = \mu_x + \frac{\pi_u}{\pi_{x|u}} (u - \mu_x) \quad (56)$$

$$\pi_{x|u} = \pi_x + \pi_u \quad (57)$$

The update in the mean is a precision-weighted prediction error. The prediction error $u - \mu_x$ is weighted proportionally to the observation precision π_u , reflecting the fact that the more observation noise there is, the less weight should be assigned to the prediction error. On the other hand, prediction error is weighted inversely proportionally to the posterior precision $\pi_{x|u}$; that is, with higher certainty about x , the impact of any new information on its estimate becomes smaller.

The same precision-weighting of prediction errors appears in the update of the means μ_i of the states x_i in the inversion of the general HGF (Figure 10, Equation 9):

$$\mu_i^{(k)} - \hat{\mu}_i^{(k)} = \frac{1}{2} \kappa_{i-1} v_{i-1}^{(k)} \frac{\hat{\pi}_{i-1}^{(k)}}{\pi_i^{(k)}} \delta_{i-1}^{(k)}, \quad (58)$$

or, in more compact notation,

$$\Delta \mu_i \propto \frac{\hat{\pi}_{i-1}}{\pi_i} \delta_{i-1}. \quad (59)$$

Owing to the hierarchical nature of the HGF, the place of the likelihood precision π_u in Equation 56 is here taken by the

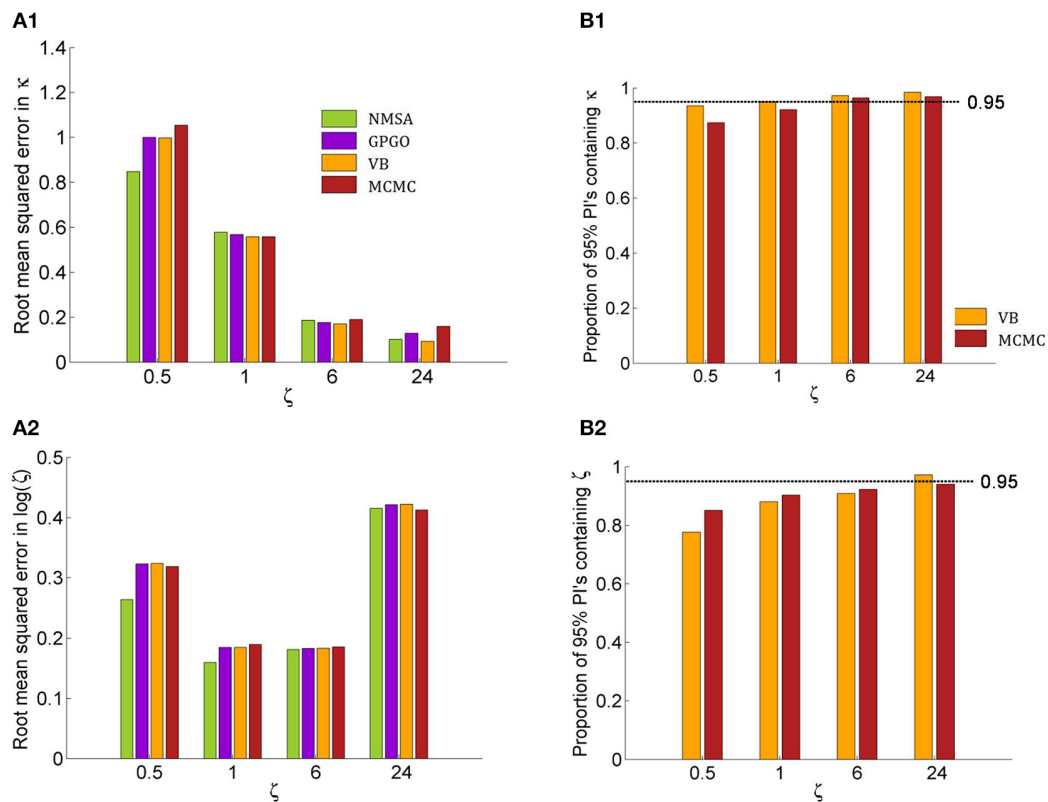


FIGURE 9 | Quantitative assessment of parameter estimation. (A) Root mean squared error of MAP estimates by noise level ζ for all four estimation methods (see **Figure 7** for key to legends). **(A1)** Estimates for κ improve with decreasing noise and do not exhibit substantially significant differences between methods although NMSA is somewhat better at very high noise.

(A2) As in **Figure 8**, estimates for ζ were assessed at the logarithmic scale. **(B)** Confidence of VB and MCMC. **(B1)** Both methods are realistically confident about their inference on κ across noise levels, with a slight tendency toward overconfidence with higher noise. **(B2)** This tendency is more pronounced with estimates of ζ .

precision of the prediction on the level below, $\hat{\pi}_{i-1}$, while the posterior precision π_i in the HGF corresponds exactly to the posterior precision $\pi_{x|u}$ in Equation 56. The precision ratio in these updates is the *learning rate* with which the prediction error is weighted. Prediction errors weighted by a learning rate are a defining feature of many reinforcement learning models (e.g., Rescorla and Wagner, 1972). The HGF furnishes a Bayesian foundation for these heuristically derived models in that it provides learning rates that are optimal given a particular agent's parameter setting. The numerator of the precision ratio in Equation 59 contains the precision of the prediction onto the level below. This relation makes sense because the higher this precision, the more meaning a given prediction error has. The denominator of the ratio contains the precision of the belief about the level being updated. Again, it makes sense that the update should be antiproportional to this since the more certain the agent is that it knows the true value of x_i , the less inclined it should be to change it. What sets the HGF apart from other models with adaptive learning rates (e.g., Sutton, 1992; Nassar et al., 2010; Payzan-LeNestour and Bossaerts, 2011; Wilson et al., 2013) is that its update equations are derived to optimize a clearly defined objective function, variational free energy, thereby minimizing surprise. Furthermore, while the Kalman filter (Kalman, 1960) is

optimal for data generated by linear dynamical systems, the HGF has the advantage that it can deal with nonlinear systems because it adapts its volatility estimate as the data come in. This adaptive adjustment of learning rates corresponds to an optimal “forgetting” algorithm that prevents learning rates from becoming too low.

Notably, the prediction error δ_{i-1} is a volatility prediction error (VOPE) in the HGF while the prediction errors in the single-level Gaussian updates Equation 56, like the first-level updates in the HGF Equation 15, refer to value prediction errors (VAPES). While a VAPE captures the error about the magnitude of a hidden state, a VOPE captures the error about the amount of change in a hidden state. The crucial point here is that the levels of the HGF are linked via the variance (or, equivalently, precision) of the prediction of the next lower level. Consequently, the inversion proceeds by updating the higher level based on the variance (or volatility) of the lower level. This becomes apparent in Equation 14. The denominator of the fraction contains predicted uncertainty about the level below, while the numerator contains observed uncertainty. These can again be broken down into informational and environmental uncertainty (see below). Whenever observed uncertainty exceeds predicted, the fraction is greater than one and the VOPE is positive. Conversely,

Update equation for posterior means

$$\Delta\mu_i^{(k)} = \frac{1}{2} \kappa_{i-1} v_{i-1}^{(k)} \frac{\hat{\pi}_{i-1}^{(k)}}{\pi_i^{(k)}} \delta_{i-1}^{(k)}$$

with

$$\Rightarrow \Delta\mu_i \propto \frac{\hat{\pi}_{i-1}}{\pi_i} \delta_{i-1}$$

$\hat{\pi}_{i-1}$	Precision of the prediction onto the level below
π_i	Posterior precision at the current level
δ_{i-1}	Prediction error for the volatility of the level below

FIGURE 10 | Posterior mean update equation. Updates are precision-weighted prediction errors. This general feature of Bayesian updating is concretized by the HGF for volatility predictions in a hierarchical setting.

when observed uncertainty is less than predicted, the VOPE is negative.

The two sources of uncertainty, informational and environmental, are clearly visible in the precision of the predictions Equation 13 and in the VOPEs Equation 14. In Equations 13 and 14, $\sigma_i^{(k-1)}$ is the informational posterior uncertainty about x_i , while $v_i^{(k)} = \tau^{(k)} \exp(\kappa_i \mu_{i+1}^{(k-1)} + \omega_i)$ is the environmental uncertainty, the magnitude of which is determined by a combination of two kinds of volatility: phasic ($\mu_{i+1}^{(k-1)}$) and tonic (ω_i). The less we know about x_i , the greater the informational uncertainty $\sigma_i^{(k-1)}$; by contrast, the more volatile the environment is, the greater the environmental uncertainty $v_i^{(k)}$.

The relation between uncertainty (informational and environmental, expected and unexpected) and volatility (phasic and tonic) can be summarized as follows: informational uncertainty could be seen as a form of expected uncertainty which, however, differs from Yu and Dayan (2005) in that it is defined in terms of posterior variance instead of estimated deviation from certainty. By contrast, environmental uncertainty can be linked to “unexpected” uncertainty and is the result of phasic and tonic volatility. We use the term “environmental” instead of “unexpected” because, in the context of the HGF, unexpected uncertainty is incorporated into the precision of predictions (cf. Equation 13), i.e., there is always some degree of belief that the environment might be changing.

In a review of the literature on different kinds of uncertainty in human decision-making, Bland and Schaefer (2012) argue that unexpected uncertainty and volatility are often not sufficiently differentiated while Payzan-LeNestour and Bossaerts (2011) make a further subdistinction of unexpected uncertainty: they differentiate between stochastic volatility and a narrower concept of unexpected uncertainty. This distinction maps exactly onto the difference between tonic and phasic volatility in the HGF. While the Kalman Filter deals optimally with tonic/stochastic volatility, the HGF can also accommodate sudden environmental changes via phasic volatility. An illustration of this ability can be found

in Mathys et al. (2011), where the U.S. Dollar to Swiss Franc exchange rate time series from the first half of the year 2010 is filtered using the HGF (their Figure 11). In addition to tonic volatility, this time series also reflects a clear change point, i.e., the markets’ realization that Greece was insolvent. The latter is captured by the HGF in phasic volatility shooting up almost vertically.

Environmental uncertainty is updated by adjusting μ_{i+1} , the estimate of the next higher level. This is done in the VOPE by comparing predicted total uncertainty (informational plus environmental, $\sigma_i^{(k-1)} + v_i^{(k)}$) to observed total uncertainty ($\sigma_i^{(k)} + (\mu_i^{(k)} - \hat{\mu}_i^{(k)})^2$). In this way, environmental uncertainty estimates are dynamically adapted to changes in the environment, leading to changes in learning rates that reflect an optimal (with respect to avoiding surprise) balance between informational and environmental uncertainty estimates.

At first glance, the precision-weighting of prediction errors may seem different in the “classical” 3-level HGF (Figure 2) with categorical outcomes, where the update for μ_2 (Equation 32) is:

$$\mu_2^{(k)} = \mu_2^{(k-1)} + \sigma_2^{(k)} \delta_1^{(k)} \quad (60)$$

At first, this simply looks like an uncertainty-weighted (not precision-weighted) update. However, if we unpack σ_2 according to Equation 33 and do a Taylor expansion in powers of $\hat{\pi}_1$, we see that it is again proportional to the precision of the prediction on the level below:

$$\sigma_2^{(k)} = \frac{\hat{\pi}_1^{(k)}}{\hat{\pi}_2^{(k)} \hat{\pi}_1^{(k)} + 1} = \hat{\pi}_1^{(k)} - \hat{\pi}_2^{(k)} (\hat{\pi}_1^{(k)})^2 + (\hat{\pi}_2^{(k)})^2 (\hat{\pi}_1^{(k)})^3 + O(4). \quad (61)$$

We have further shown a principled way how to define decision models based on perceptual HGF inferences, namely by deriving them from a loss function. Based on such decision models,

it is possible to infer on model parameters and state trajectories from observed decisions. In the simulation study we have reported here, we could show that even with considerable decision noise, we can reliably infer model parameters based on a few hundred data points for binary decisions. Several recent studies have done this in practice, estimating subject-specific HGF parameters from behavioral data. For example, (Vossel et al., 2013) used the HGF to model learning in human subjects performing a Posner task with varying outcome contingencies. This study sought to compare different possible explanations for measured eye movements (saccadic reaction speeds), using a factorial model space comprising three alternative perceptual and three different response models. Model comparison showed that the 3-level HGF had greater model evidence than simpler versions of itself and Rescorla–Wagner learning (Rescorla and Wagner, 1972). This indicates that humans are capable of hierarchically structured learning, exploiting volatility estimates to adapt their learning rate dynamically. The same conclusion emerged from the study of Iglesias et al. (2013) who used the HGF to analyze human learning of auditory-visual associations which varied unpredictably in time. This study subsequently used the trial-wise estimates of precision-weighted prediction errors (i.e., $\frac{\hat{\pi}_{i-1}}{\pi_i} \delta_{i-1}$) in fMRI analyses, demonstrating activation of the dopaminergic midbrain with first-level (i.e., sensory outcome) precision-weighted prediction errors, and activation of the cholinergic basal forebrain with second-level (i.e., probability) precision-weighted prediction errors. These findings resonate with recent proposals that an important aspect of neuromodulatory function is the encoding of precision (Friston, 2009).

The HGF can, in principle, accommodate any form of loss function in the decision model. This choice will depend on the particular question addressed and the assumptions of the application domain (e.g., rationality assumptions). In the examples shown in this paper, we employ loss functions that are quadratic. This reflects the fact that squared losses imply a Gaussian distribution of errors, which is the appropriate choice where little is known about the true distribution because the Gaussian has maximum entropy (i.e., the least arbitrary assumptions) for a given mean and variance. This means that using a quadratic loss function is the most conservative choice in the absence of additional prior knowledge about the error distribution, where the term error refers to the agent's failure to make the choice that minimizes its expected loss.

Since all four methods of inverting the decision model performed well in our simulation study, we may focus on secondary criteria in choosing a method for practical applications. The most important of these criteria are the computational burden imposed and the amount of information contained in the estimate. The best performer in these respects is currently variational Bayes because it is efficient and provides an estimate of the whole posterior distribution for all parameters in addition to an approximation to the free energy bound on the log-model evidence, enabling model comparison. MCMC offers the same in principle, but at a considerably higher computational cost. GPGO is computationally more expensive than VB but may be a strong contender for future cases with multimodal posterior distributions. The weakest contender is NMSA because it is not much

more efficient than VB but only offers a point estimate of the MAP parameter values.

In summary, the HGF provides a general and powerful framework for inferring on belief updating processes and learning styles of individual subjects in a volatile environment. This makes it a generic tool for studying perception in a Helmholtzian sense. The simple nature of the HGF updates in the form of precision-weighted prediction errors do not only enhance their biological interpretability and plausibility (cf. Friston, 2009) but are also crucial for practical applications. The ability of the HGF to infer learning styles of individual subjects from behavioral data and its support of Bayesian model comparison offer interesting opportunities for studying individual differences and particularly for clinical studies on psychopathology. To facilitate such practical applications, we have developed a software toolbox based on Matlab that is freely available for downloading as part of the TAPAS collection at <http://www.translationalneuromodeling.org/tapas/>. The HGF toolbox is specifically tailored to the implementations of discrete time filtering models, as opposed to the DAVB toolbox, which is mostly aimed at inverting dynamic models in continuous time. The HGF toolbox implements most of the models described in this article, plus some additional ones, and will be the focus of a forthcoming article.

ACKNOWLEDGMENTS

The authors gratefully acknowledge support by the UCL-MPS Initiative on Computational Psychiatry and Ageing Research (Christoph D. Mathys), the René and Susanne Braginsky Foundation (Klaas E. Stephan), the Clinical Research Priority Program “Multiple Sclerosis” (Klaas E. Stephan), and the Wellcome Trust (Karl J. Friston).

REFERENCES

- Adams, R. A., Stephan, K. E., Brown, H. R., and Friston, K. J. (2013). The computational anatomy of psychosis. *Front. Schizophr.* 4:47. doi: 10.3389/fpsyt.2013.00047
- Behrens, T. E. J., Hunt, L. T., Woolrich, M. W., and Rushworth, M. F. S. (2008). Associative learning of social value. *Nature* 456, 245–249. doi: 10.1038/nature07538
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221. doi: 10.1038/nn1954
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York, NY: Springer.
- Bland, A. R., and Schaefer, A. (2012). Different varieties of uncertainty in human decision-making. *Front. Neurosci.* 6:85. doi: 10.3389/fnins.2012.00085
- Daunizeau, J., Adam, V., and Rigoux, L. (2014). VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioural data. *PLoS Comput. Biol.* 10:e1003441. doi: 10.1371/journal.pcbi.1003441
- Daunizeau, J., den Ouden, H. E. M., Pessiglione, M., Kiebel, S. J., Friston, K. J., and Stephan, K. E. (2010a). Observing the Observer (II): deciding when to decide. *PLoS ONE* 5:e15555. doi: 10.1371/journal.pone.0015555
- Daunizeau, J., den Ouden, H. E. M., Pessiglione, M., Kiebel, S. J., Stephan, K. E., and Friston, K. J. (2010b). Observing the Observer (I): meta-bayesian models of learning and decision-making. *PLoS ONE* 5:e15554. doi: 10.1371/journal.pone.0015554
- Daunizeau, J., Friston, K. J., and Kiebel, S. J. (2009). Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Phys. Nonlinear Phenom.* 238, 2089–2118. doi: 10.1016/j.physd.2009.08.002
- Daw, N. D., and Doya, K. (2006). The computational neurobiology of learning and reward. *Curr. Opin. Neurobiol.* 16, 199–204. doi: 10.1016/j.conb.2006.03.006

- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., and Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature* 441, 876–879. doi: 10.1038/nature04766
- Dayan, P., Hinton, G. E., Neal, R. M., and Zemel, R. S. (1995). The helmholtz machine. *Neural Comput.* 7, 889–904. doi: 10.1162/neco.1995.7.5.889
- Doya, K., Ishii, S., Pouget, A., and Rao, R. P. N. (2011). *Bayesian Brain: Probabilistic Approaches to Neural Coding*. Cambridge, MA: MIT Press.
- Evans, L. C. (2010). *Partial Differential Equations*. Providence, RI: American Mathematical Society.
- Faisal, A. A., Selen, L. P. J., and Wolpert, D. M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303. doi: 10.1038/nrn2258
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends Cogn. Sci.* 13, 293–301. doi: 10.1016/j.tics.2009.04.005
- Friston, K. J., and Dolan, R. J. (2010). Computational and dynamic models in neuroimaging. *Neuroimage* 52, 752–765. doi: 10.1016/j.neuroimage.2009.12.068
- Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., and Penny, W. (2007). Variational free energy and the Laplace approximation. *Neuroimage* 34, 220–234. doi: 10.1016/j.neuroimage.2006.08.035
- Friston, K. J., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., and Dolan, R. J. (2013). The anatomy of choice: active inference and agency. *Front. Hum. Neurosci.* 7:598. doi: 10.3389/fnhum.2013.00598
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2003). *Bayesian Data Analysis*. Boca Raton, FL: Chapman & Hall/CRC.
- Helmholtz, H. (1860). *Handbuch der Physiologischen Optik*. English translation (1962): J. P. C. Southall. New York, NY: Dover.
- Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E. M., et al. (2013). Hierarchical prediction errors in mid-brain and basal forebrain during sensory learning. *Neuron* 80, 519–530. doi: 10.1016/j.neuron.2013.09.009
- Joffily, M., and Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLoS Comput. Biol.* 9:e1003094. doi: 10.1371/journal.pcbi.1003094
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *J. Basic Eng.* 82, 35–45.
- Knill, D. C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719. doi: 10.1016/j.tins.2004.10.007
- Körding, K. P., and Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* 10, 319–326. doi: 10.1016/j.tics.2006.05.003
- Lomakina, E., Vezhnevets, A., Mathys, C., Brodersen, K. H., Stephan, K. E., and Buhmann, J. (2012). *Bayesian Global Optimization for Model-Based Neuroimaging*. HBM E-Poster. Available online at: <https://www.aievolution.com/hbm1201/index.cfm?do=abs.viewAbs&abs=6320>
- Macready, W. G., and Wolpert, D. H., I. (1998). Bandit problems and the exploration/exploitation tradeoff. *Evol. Comput. IEEE Trans.* 2, 2–22. doi: 10.1109/4235.728210
- Mathys, C., Daunizeau, J., Friston, K. J., and Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* 5:39. doi: 10.3389/fnhum.2011.00039
- Nassar, M. R., Wilson, R. C., Heasly, B., and Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* 30, 12366–12378. doi: 10.1523/JNEUROSCI.0822-10.2010
- Nelder, J. A., and Mead, R. (1965). A simplex method for function minimization. *Comput. J.* 7, 308–313. doi: 10.1093/comjnl/7.4.308
- O'Doherty, J. P., Hampton, A., and Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Ann. N.Y. Acad. Sci.* 1104, 35–53. doi: 10.1196/annals.1390.022
- Paliwal, S., Petzschner, F., Schmitz, A. K., Tittgemeyer, M., and Stephan, K. E. (2014). A model-based analysis of impulsivity using a slot-machine gambling paradigm. *Front. Hum. Neurosci.* 8:428. doi: 10.3389/fnhum.2014.00428
- Payzan-LeNestour, E., and Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Comput. Biol.* 7:e1001048. doi: 10.1371/journal.pcbi.1001048
- Rasmussen, C. E., and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press.
- Rescorla, R. A., and Wagner, A. R. (1972). "A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement," in *Classical Conditioning II: Current Research and Theory*, eds A. H. Black and W. F. Prokasy (New York, NY: Appleton-Century-Crofts), 64–99.
- Schwartenbeck, P., FitzGerald, T., Dolan, R. J., and Friston, K. (2013). Exploration, novelty, surprise, and free energy minimization. *Front. Cogn. Sci.* 4:710. doi: 10.3389/fpsyg.2013.00710
- Stephan, K. E., Tittgemeyer, M., Knösche, T. R., Moran, R. J., and Friston, K. J. (2009). Tractography-based priors for dynamic causal models. *Neuroimage* 47, 1628–1638. doi: 10.1016/j.neuroimage.2009.05.096
- Sutton, R. (1992). "Gain adaptation beats least squares?," in *Proceedings of the 7th Yale Workshop on Adaptive and Learning Systems* (New Haven, CT), 161–166.
- Sutton, R. S., and Barto, A. G. (1998). *Reinforcement Learning: an Introduction*. Cambridge, MA: MIT Press.
- Vossel, S., Mathys, C., Daunizeau, J., Bauer, M., Driver, J., Friston, K. J., et al. (2013). Spatial attention, precision, and Bayesian inference: a study of saccadic response speed. *Cereb. Cortex* 24, 1436–1450. doi: 10.1093/cercor/bhs418
- Wilson, R. C., Nassar, M. R., and Gold, J. I. (2013). A mixture of delta-rules approximation to Bayesian inference in change-point problems. *PLoS Comput. Biol.* 9:e1003150. doi: 10.1371/journal.pcbi.1003150
- Yu, A., and Dayan, P. (2003). *Expected and Unexpected Uncertainty: ACh and NE in the Neocortex*. Available online at: <http://books.nips.cc/nips15.html> (Accessed: July 29, 2013).
- Yu, A. J., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692. doi: 10.1016/j.neuron.2005.04.026

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 22 February 2014; accepted: 27 September 2014; published online: 19 November 2014.

Citation: Mathys CD, Lomakina EI, Daunizeau J, Iglesias S, Brodersen KH, Friston KJ and Stephan KE (2014) Uncertainty in perception and the Hierarchical Gaussian Filter. *Front. Hum. Neurosci.* 8:825. doi: 10.3389/fnhum.2014.00825

This article was submitted to the journal *Frontiers in Human Neuroscience*.

Copyright © 2014 Mathys, Lomakina, Daunizeau, Iglesias, Brodersen, Friston and Stephan. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

APPENDICES

(A) COUPLING BETWEEN LEVELS

Since $f(x)$ is a positive function, there must be a function $g(x)$ whose exponential it is.

$$f(x) > 0 \forall x \implies \exists g : f(x) = \exp(g(x)) \forall x \quad (\text{A1})$$

Expanding $g(x)$ then amounts to expanding the logarithm of $f(x)$.

$$\begin{aligned} g(x) &= g(a) + g'(a) \cdot (x - a) + O(2) = \log f(x) = \\ &= \log f(a) + \frac{f'(a)}{f(a)} \cdot (x - a) + O(2) = \\ &= \underbrace{\frac{f'(a)}{f(a)}}_{\stackrel{\text{def}}{=} \kappa} \cdot x + \underbrace{\log f(a) - a \cdot \frac{f'(a)}{f(a)}}_{\stackrel{\text{def}}{=} \omega} + O(2) = \end{aligned} \quad (\text{A2})$$

$$\begin{aligned} &= \kappa x + \omega + O(2) \\ \implies f(x) &\approx \exp(\kappa x + \omega) \end{aligned} \quad (\text{A3})$$

Given $f(x)$, κ and ω are only unique with respect to the choice of a particular expansion point a , except when $f(x)$ is the exponential of a first-order polynomial, in which case the above approximation is exact. The greater the weight of higher-order terms in $g(x)$, the less accurate the approximation will be far from the expansion point a . In ignoring second and higher order terms, we effectively restrict the HGF to first-order coupling functions (i.e., to coupling functions that are the exponentials of first-order polynomials). However, as this derivation shows, locally (i.e., within small variations of x), this restriction does not matter.

(B) VARIATIONAL INVERSION OF THE HGF

Variational inversion provides closed-form one-step update equations for the sufficient statistics that describe the agent's belief about the state of its environment. These update equations are derived in two steps. First, for a particular level i , $1 < i < n$, of the hierarchy, a mean field approximation is introduced, where the distributions $q(x_j)$ for all the other levels $j \neq i$ are assumed to be known and Gaussian, fully described by the known sufficient statistics $\{\mu_j, \sigma_j\}_{j \neq i}$:

$$q(x_j) = \mathcal{N}(x_j; \mu_j, \sigma_j). \quad (\text{B1})$$

In the absence of sensory noise, x_1 is directly observed; this means that $\mu_1^{(k)} = x_1^{(k)}$ and $\sigma_1^{(k)} = 0$ for all k . The approximate posterior $\hat{q}(x_i^{(k)})$ for level i at time k , given sensory input $\mu_1^{(1 \dots k)} = \{\mu_1^{(1)}, \mu_1^{(2)}, \dots, \mu_1^{(k)}\}$, is then

$$\hat{q}(x_i^{(k)}) = \frac{1}{\mathcal{Z}_i^{(k)}} \exp(I(x_i^{(k)})). \quad (\text{B2})$$

with variational energy I ,

$$\begin{aligned} I(x_i^{(k)}) &= \int q(x_{\setminus i}^{(k)}) \ln p(x_i^{(k)}, x_{\setminus i}^{(k)}, \chi | \mu_1^{(1 \dots k)}) dx_{\setminus i}^{(k)}, \\ q(x_{\setminus i}^{(k)}) &= \prod_{j \neq i} q(x_j), \\ x_{\setminus i} &= \{x_j\}_{j \neq i}, \\ \chi &= \{\kappa_i, \omega_i, \vartheta\}_{1 \leq i \leq n}. \end{aligned} \quad (\text{B3})$$

We further unpack this to get

$$\begin{aligned}
 p\left(x_i^{(k)}, x_{\setminus i}^{(k)}, \chi | \mu_1^{(1\dots k)}\right) &= \int \prod_{j=1}^n p\left(x_j^{(k)} | x_j^{(k-1)}, x_{\setminus j}^{(k)}, \chi\right) p\left(x_j^{(k-1)} | \mu_1^{(1\dots k-1)}\right) dx_j^{(k-1)} \\
 p\left(x_j^{(k)} | x_j^{(k-1)}, x_{\setminus j}^{(k)}, \chi\right) &= \mathcal{N}\left(x_j^{(k)}; x_j^{(k-1)}, t^{(k)} \exp\left(\kappa_j x_{j+1}^{(k)} + \omega_j\right)\right) \\
 p\left(x_j^{(k-1)} | \mu_1^{(1\dots k-1)}\right) &= \mathcal{N}\left(x_j^{(k-1)}; \mu_j^{(k-1)}, \sigma_j^{(k-1)}\right).
 \end{aligned} \tag{B4}$$

Performing this last integral, we obtain

$$\begin{aligned}
 p\left(x_i^{(k)}, x_{\setminus i}^{(k)}, \chi | \mu_1^{(1\dots k)}\right) &= \prod_{j=1}^n \mathcal{N}\left(x_j^{(k)}; \mu_j^{(k-1)}, \sigma_j^{(k-1)} + t^{(k)} \exp\left(\kappa_j x_{j+1}^{(k)} + \omega_j\right)\right), \\
 \exp\left(\kappa_n x_{n+1}^{(k)} + \omega_n\right) &\equiv \vartheta.
 \end{aligned} \tag{B5}$$

This constitutes the “prediction step” of the update. Specifically, we let the random walk do its work by integrating out all states from the previous time point, thus ensuring that the mean field approximation only applies to current values of states. This is akin to the prediction step in the Kalman filter since it gives us the predictive densities for $x_i^{(k)}$ given inputs up to time $k - 1$.

This now enables us to solve the integral in Equation B3, yielding (see Appendix C)

$$\begin{aligned}
 I\left(x_i^{(k)}\right) &= -\frac{1}{2} \ln\left(\sigma_{i-1}^{(k-1)} + t^{(k)} \exp\left(\kappa_{i-1} x_i^{(k)} + \omega_{i-1}\right)\right) \\
 &\quad - \frac{1}{2} \frac{\sigma_{i-1}^{(k)} + \left(\mu_{i-1}^{(k)} - \mu_{i-1}^{(k-1)}\right)^2}{\sigma_{i-1}^{(k-1)} + t^{(k)} \exp\left(\kappa_{i-1} x_i^{(k)} + \omega_{i-1}\right)} \\
 &\quad - \frac{1}{2} \frac{1}{\sigma_i^{(k-1)} + t^{(k)} \exp\left(\kappa_i \mu_{i+1}^{(k-1)} + \omega_i\right)} \left(x_i^{(k)} - \mu_i^{(k-1)}\right)^2.
 \end{aligned} \tag{B6}$$

Following the procedure of Mathys et al. (2011), we can now calculate the mean and precision of the Gaussian posterior for $x_i^{(k)}$:

$$\pi_i^{(k)} = -I''\left(\hat{\mu}_i^{(k)}\right), \tag{B7}$$

$$\mu_i^{(k)} = \hat{\mu}_i^{(k)} + \frac{1}{\pi_i^{(k)}} I'\left(\hat{\mu}_i^{(k)}\right), \tag{B8}$$

where a double and single prime denote second and first derivative, respectively. This gives us Equations 9 and 10.

(C) CALCULATION OF THE VARIATIONAL ENERGY

By substituting Equation B5 into Equation B3, we get

$$\begin{aligned}
 I\left(x_i^{(k)}\right) &= \int q\left(x_{\setminus i}^{(k)}\right) \ln p\left(x_i^{(k)}, x_{\setminus i}^{(k)}, \chi | \mu_1^{(1\dots k)}\right) dx_{\setminus i}^{(k)} \\
 &= \int q\left(x_{\setminus i}^{(k)}\right) \ln \left(\prod_{j=1}^n \mathcal{N}\left(x_j^{(k)}; \mu_j^{(k-1)}, \sigma_j^{(k-1)} + \exp\left(\kappa_j x_{j+1}^{(k)} + \omega_j\right)\right)\right) dx_{\setminus i}^{(k)} \\
 &= \sum_{j=1}^n \int q\left(x_{\setminus i}^{(k)}\right) \ln \mathcal{N}\left(x_j^{(k)}; \mu_j^{(k-1)}, \sigma_j^{(k-1)} + \exp\left(\kappa_j x_{j+1}^{(k)} + \omega_j\right)\right) dx_{\setminus i}^{(k)} \\
 &= \int q\left(x_{\setminus i}^{(k)}\right) \ln \mathcal{N}\left(x_{i-1}^{(k)}; \mu_{i-1}^{(k-1)}, \sigma_{i-1}^{(k-1)} + \exp\left(\kappa_{i-1} x_i^{(k)} + \omega_{i-1}\right)\right) dx_{\setminus i}^{(k)}
 \end{aligned} \tag{C1}$$

$$\begin{aligned}
& + \int q(x_i^{(k)}) \ln \mathcal{N}(x_i^{(k)}; \mu_i^{(k-1)}, \sigma_i^{(k-1)} + \exp(\kappa_i x_{i+1}^{(k)} + \omega_i)) dx_i^{(k)} + \text{const.} \\
& = \int q(x_{i-1}^{(k)}) \ln \mathcal{N}(x_{i-1}^{(k)}; \mu_{i-1}^{(k-1)}, \sigma_{i-1}^{(k-1)} + \exp(\kappa_{i-1} x_i^{(k)} + \omega_{i-1})) dx_{i-1}^{(k)} + \\
& + \int q(x_{i+1}^{(k)}) \ln \mathcal{N}(x_i^{(k)}; \mu_i^{(k-1)}, \sigma_i^{(k-1)} + \exp(\kappa_i x_{i+1}^{(k)} + \omega_i)) dx_{i+1}^{(k)} + \text{const.}
\end{aligned}$$

We now have to solve these last two integrals. The first one can be solved analytically:

$$\begin{aligned}
& \int q(x_{i-1}^{(k)}) \ln \mathcal{N}(x_{i-1}^{(k)}; \mu_{i-1}^{(k-1)}, \sigma_{i-1}^{(k-1)} + \exp(\kappa_{i-1} x_i^{(k)} + \omega_{i-1})) dx_{i-1}^{(k)} \\
& = -\frac{1}{2} \ln(\sigma_{i-1}^{(k-1)} + \exp(\kappa_{i-1} x_i^{(k)} + \omega_{i-1})) - \frac{1}{2} \frac{\sigma_{i-1}^{(k)} + (\mu_{i-1}^{(k)} - \mu_{i-1}^{(k-1)})^2}{\sigma_{i-1}^{(k-1)} + \exp(\kappa_{i-1} x_i^{(k)} + \omega_{i-1})} \quad (\text{C2})
\end{aligned}$$

To solve the second integral, we take

$$\begin{aligned}
& \int q(x_{i+1}^{(k)}) \ln \mathcal{N}(x_i^{(k)}; \mu_i^{(k-1)}, \sigma_i^{(k-1)} + \exp(\kappa_i x_{i+1}^{(k)} + \omega_i)) dx_{i+1}^{(k)} \\
& \approx -\frac{1}{2} \frac{1}{\sigma_i^{(k-1)} + \exp(\kappa_i \mu_{i+1}^{(k-1)} + \omega_i)} (x_i^{(k)} - \mu_i^{(k-1)})^2 + \text{const.} \quad (\text{C3})
\end{aligned}$$

The sum of Equations C2 and C3 then is the variational energy of Equation B6, up to a constant term that can be absorbed into the normalization constant \mathcal{Z}_i (cf. Equation B2).

(D) VARIATIONAL ENERGIES FOR CATEGORICAL OUTCOMES

Using the notation

$$\langle f(x) \rangle_{q_j} \stackrel{\text{def}}{=} \int q(x_j) f(x) dx_j \quad (\text{D1})$$

for the expectation of $f(x)$ under $q(x_j)$ together with the definition of the model described graphically **Figure 2**, we can rewrite Equation B3 as a sum of expectations

$$\begin{aligned}
I(x_i^{(k)}) & = \langle \ln p(u^{(k)}, x^{(k)} | \chi) \rangle_{q_i} \\
& = \langle \ln p(u^{(k)} | x_1^{(k)}) \rangle_{q_i} + \langle \ln p(x_1^{(k)} | x_2^{(k)}) \rangle_{q_i} \\
& + \langle \ln p(x_2^{(k)} | x_3^{(k)}, \kappa, \omega) \rangle_{q_i} + \langle \ln p(x_3^{(k)} | \vartheta) \rangle_{q_i}
\end{aligned} \quad (\text{D2})$$

The term $\ln p(u^{(k)} | x_1^{(k)})$ is included here to cover also models with sensory uncertainty as discussed in Mathys et al. (2011). In cases without such uncertainty (sc. $x_1 = u$), the term vanishes. $p(x_1^{(k)} | x_2^{(k)})$ can be taken directly from Equation 5, while

$$\begin{aligned}
p(x_2^{(k)} | x_3^{(k)}, \kappa, \omega) & = \int p(x_2^{(k)} | x_2^{(k-1)}, x_3^{(k)}, \kappa, \omega) p(x_2^{(k-1)}) dx_2^{(k-1)} \\
& \approx \int p(x_2^{(k)} | x_2^{(k-1)}, x_3^{(k)}, \kappa, \omega) q(x_2^{(k-1)}) dx_2^{(k-1)} \\
& = \int \mathcal{N}(x_2^{(k)}; x_2^{(k-1)}, t^{(k)} \exp(\kappa x_3^{(k)} + \omega)) \mathcal{N}(x_2^{(k-1)}; \mu_2^{(k-1)}, \sigma_2^{(k-1)}) dx_2^{(k-1)} \\
& = \mathcal{N}(x_2^{(k)}; \mu_2^{(k-1)}, \sigma_2^{(k-1)} + t^{(k)} \exp(\kappa x_3^{(k)} + \omega)) \quad (\text{D3})
\end{aligned}$$

and

$$\begin{aligned}
 p\left(x_3^{(k)}|\vartheta\right) &= \int p\left(x_3^{(k)}|x_3^{(k-1)}, \vartheta\right) p\left(x_3^{(k-1)}\right) dx_3^{(k-1)} \\
 &\approx \int p\left(x_3^{(k)}|x_3^{(k-1)}, \vartheta\right) q\left(x_3^{(k-1)}\right) dx_3^{(k-1)} \\
 &= \int \mathcal{N}\left(x_3^{(k)}; x_3^{(k-1)}, t^{(k)}\vartheta\right) \mathcal{N}\left(x_3^{(k-1)}; \mu_3^{(k-1)}, \sigma_3^{(k-1)}\right) dx_3^{(k-1)} \\
 &= \mathcal{N}\left(x_3^{(k)}; \mu_3^{(k-1)}, \sigma_3^{(k-1)} + t^{(k)}\vartheta\right)
 \end{aligned} \tag{D4}$$

This is the “prediction step” (cf. Appendix B).

We only need to determine the $I\left(x_i^{(k)}\right)$ up to a constant because any constant term can always be absorbed into \mathcal{Z}_i when forming $\hat{q}\left(x_i^{(k)}\right)$ according to Equation B2. For the three levels of our example model, this means

$$I\left(x_1^{(k)}\right) = \left\langle \ln p\left(x_1^{(k)}|x_1^{(k)}\right) \right\rangle_{q_{\backslash 1}} + \left\langle \ln p\left(x_1^{(k)}|x_2^{(k)}\right) \right\rangle_{q_{\backslash 1}} + \text{const.} \tag{D5}$$

$$I\left(x_2^{(k)}\right) = \left\langle \ln p\left(x_1^{(k)}|x_2^{(k)}\right) \right\rangle_{q_{\backslash 2}} + \left\langle \ln p\left(x_2^{(k)}|x_3^{(k)}, \kappa, \omega\right) \right\rangle_{q_{\backslash 2}} + \text{const.} \tag{D6}$$

$$I\left(x_3^{(k)}\right) = \left\langle \ln p\left(x_2^{(k)}|x_3^{(k)}, \kappa, \omega\right) \right\rangle_{q_{\backslash 3}} + \left\langle \ln p\left(x_3^{(k)}|\vartheta\right) \right\rangle_{q_{\backslash 3}} + \text{const.} \tag{D7}$$

With two exceptions, all integrals on the right-hand sides above can be solved analytically in all cases considered here, including sensory uncertainty and inference on continuous-valued states.

The two exceptions are the following: first, to solve $\left\langle \ln p\left(x_1^{(k)}|x_2^{(k)}\right) \right\rangle_{q_{\backslash 1}}$, we expand $\ln s\left(x_2^{(k)}\right)$ to second order around the prior expectation $\mu_2^{(k-1)}$ of $x_2^{(k)}$:

$$\begin{aligned}
 \ln s\left(x_2^{(k)}\right) &\approx \ln s\left(\mu_2^{(k-1)}\right) + \left(1 - s\left(\mu_2^{(k-1)}\right)\right)\left(x_2^{(k)} - \mu_2^{(k-1)}\right) \\
 &\quad + \frac{1}{2}\left(s\left(\mu_2^{(k-1)}\right)^2 - s\left(\mu_2^{(k-1)}\right)\right)\left(x_2^{(k)} - \mu_2^{(k-1)}\right)^2
 \end{aligned} \tag{D8}$$

Second, to solve $\left\langle \ln p\left(x_2^{(k)}|x_3^{(k)}, \kappa, \omega\right) \right\rangle_{q_{\backslash 2}}$, we take

$$\left\langle \frac{1}{\sigma_2^{(k-1)} + t^{(k)} \exp\left(\kappa \cdot x_3^{(k)} + \omega\right)} \right\rangle_{q_{\backslash 2}} \approx \frac{1}{\sigma_2^{(k-1)} + t^{(k)} \exp\left(\kappa \cdot \mu_3^{(k-1)} + \omega\right)} \tag{D9}$$

The result of doing the integrals in Equations D5–D7 is:

$$I\left(x_1^{(k)}\right) = x_1^{(k)} \ln s\left(\mu_2^{(k-1)}\right) + \left(1 - x_1^{(k)}\right) \ln \left(1 - s\left(\mu_2^{(k-1)}\right)\right) \tag{D10}$$

$$I\left(x_2^{(k)}\right) = \ln s\left(x_2^{(k)}\right) + x_2^{(k)}\left(\mu_1^{(k)} - 1\right) - \frac{1}{2\left(\sigma_2^{(k-1)} + e^{\kappa \mu_3^{(k-1)} + \omega}\right)}\left(x_2^{(k)} - \mu_2^{(k-1)}\right)^2 \tag{D11}$$

$$\begin{aligned}
 I\left(x_3^{(k)}\right) &= -\frac{1}{2} \ln \left(\sigma_2^{(k-1)} + e^{\kappa x_3^{(k)} + \omega}\right) - \frac{1}{2} \frac{\sigma_2^{(k)} + \left(\mu_2^{(k)} - \mu_2^{(k-1)}\right)^2}{\sigma_2^{(k-1)} + e^{\kappa x_3^{(k)} + \omega}} \\
 &\quad - \frac{1}{2\left(\sigma_3^{(k-1)} + \vartheta\right)}\left(x_3^{(k)} - \mu_3^{(k-1)}\right)^2.
 \end{aligned} \tag{D12}$$

At the first level, q can be calculated directly. By inserting Equation D10 into Equation B2 and noting that, at the first level, $q = \hat{q}$ since \hat{q} already conforms to the distributional assumptions (Bernoulli) that q is subject to, we find

$$q(x_1^{(k)}) = \begin{cases} 1 & \text{for } x_1^{(k)} = u^{(k)}, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{D13})$$

This result is due to the fact that, in the absence of sensory uncertainty, $I(x_1)$ is only defined for $x_1 = u$ while taking on a value of “minus infinity” otherwise (cf. Equation D5).

At the second level, applying Equations B7 and B8 to Equation D11 yields Equations 32 and 33. At the third level, we recover the standard HGF variational energy for the top level given in Equation B6, yielding the standard HGF updates of Equations 9 and 10.

(E) DERIVATION OF ξ^*

ξ^* can be unpacked in the following way:

$$\begin{aligned} \xi^* &= \arg \max_{\xi} p(\xi | y, u) = \arg \max_{\xi} \frac{p(\xi, y | u)}{p(y)} \\ &= \arg \max_{\xi} p(\xi, y | u) = \arg \max_{\xi} \ln p(\xi, y | u) \\ &= \arg \max_{\xi} \ln(p(y | \xi, u) p(\xi)) \\ &= \arg \max_{\xi} (\ln p(y | \xi, u) + \ln p(\xi)) \\ &= \arg \max_{\xi} \left(\sum_k \ln p(y^{(k)} | \xi, u) + \ln p(\xi) \right) \\ &= \arg \max_{\xi} \left(\sum_k \ln p(y^{(k)} | \lambda^{(k)}(\chi, \lambda^{(0)}, u), \zeta) + \ln p(\xi) \right) \end{aligned} \quad (\text{E1})$$

(F) COORDINATE CHOICE ON HIGHER LEVELS

This appendix deals with the consequences of choosing a particular scale and origin for x_i on higher levels (i.e., where i is greater than 1 for x_1 continuous and greater than 2 for x_1 discrete). These consequences are important for parameter identifiability. Whenever, there is an ambiguity between parameter values and a coordinate transformation (i.e., shifting and rescaling) of x_i , one or several parameters will not be identifiable. An example will show what we mean by this.

In a three-level HGF with binary x_1 , the state x_3 at the third level of the model represents the log-volatility of x_2 . We now make use of the fact that any change in the initial value $\mu_3^{(0)}$ of μ_3 can be neutralized by corresponding changes in κ and ω . This means that by adjusting $\mu_3^{(0)}$ appropriately, we can set $\kappa = 1$, thereby making it seemingly disappear from the model. As an example, here are parameter estimates from a behavioral study where estimation was performed with $\mu_3^{(0)}$ set to 1:

$$\begin{aligned} \kappa &= 2.49 \\ \sigma_3^{(0)} &= 0.988 \\ \vartheta &= 0.000592 \end{aligned} \quad (\text{F1})$$

We can now make κ disappear by setting it to 1. This change should be neutral to the model's predictions of input, which can only be achieved by the following compensatory substitutions:

$$\begin{aligned} \mu_3^{(0)} = 1 &\rightarrow \mu_3^{(0)} = 2.49 \\ \sigma_3^{(0)} = 0.988 &\rightarrow \sigma_3^{(0)} = 0.988 \cdot 2.49^2 \\ \vartheta = 0.000592 &\rightarrow \vartheta = 0.000592 \cdot 2.49^2 \end{aligned} \quad (\text{F2})$$

At the first two levels, nothing has changed (the trajectory of μ_2 and therefore the input predictions $\hat{\mu}_1$ are the same); however, at the third level x_3 has been rescaled by the inverse of the factor with which κ has been rescaled (i.e., $1/2.49$). This means that the term

$$v^{(k)} \stackrel{\text{def}}{=} t^{(k)} \exp\left(\kappa \mu_3^{(k-1)} + \omega\right) \quad (\text{F3})$$

is invariant under the above transformation for all time points k , leading to the same trajectory in μ_2 and σ_2 according to Equations 37 and 38 as before.

Effectively, because of the transformation to a fixed $\kappa = 1$, our estimate of κ has been reinterpreted as an estimate of $\mu_3^{(0)}$. This is a consequence of our freedom to choose coordinates on x_3 . To take an analogy from geometry, the distance between Zurich and London is an objective geometric quantity that does not change whether it is measured in miles or kilometers; we may even introduce new units where this distance is 1. Likewise, we may always rescale x_3 individually for each agent (and estimation) such that the coupling κ between the second and third level has value 1. Note, however, that this does not prevent the coupling from differing between agents, just as the actual geometric lengths of a mile and a kilometer are different even though they both have value 1 in a mile-based and kilometer-based coordinate system, respectively.

Whenever the representations of x_3 (i.e., μ_3 and σ_3) are part of the observation model, this gives us a direct handle on x_3 and measures of its representations can immediately be compared between agents, provided we use always the same coordinates (which we would automatically do without even thinking about it; individually rescaling κ to 1 would then obviously be unwise). However, in cases where we have no measure of x_3 through the observation model, there is a fundamental ambiguity between individual differences in coupling (i.e., κ) and individual differences in priors for x_3 (i.e., $\mu_3^{(0)}$).

Nonetheless, we have to make some choice of coordinates. In setting $\mu_3^{(0)} = 1$, we choose to take the belief on environmental volatility that an agent begins inference with as the benchmark. Observed differences in learning are then attributed to, inter alia, differences in coupling. This is one way to obtain comparable measures of belief on x_3 (and consequently, ?) between agents, since one may be able to influence their priors while there is usually no way to equalize their coupling levels.

Just as it is possible to set κ to an arbitrary non-zero value while keeping v invariant by compensatory substitutions, one can set ω to an arbitrary value with invariant v using another set of compensatory substitutions. In particular, ω can be set to zero, thereby making it seemingly disappear from the model, just as κ seems to disappear when set to 1.

We have seen that any change of scale in x_3 is expressed in a corresponding change of κ . However, there is an additional degree of freedom in choosing coordinates on x_3 : the choice of origin. Changes in this are expressed in a corresponding change of ω :

$$x'_3 \stackrel{\text{def}}{=} x_3 + a \quad (\text{F4})$$

then

$$\exp(\kappa x_3 + \omega) = \exp(\kappa (x'_3 - a) + \omega) = \exp(\kappa x'_3 + \omega') \quad (\text{F5})$$

with

$$\omega' \stackrel{\text{def}}{=} -\kappa a + \omega \quad (\text{F6})$$

In our example, we have reinterpreted the estimate of κ as one of $\mu_3^{(0)}$. We may now go on to reinterpret it another time, this time as an estimate of ω (up to now fixed to -4) by shifting the origin on x_3 such that $\mu_3^{(0)}$ is again fixed to 1:

$$\begin{aligned} \mu_3^{(0)} = 2.49 &\rightarrow \mu_3^{(0)} = 1 \\ \omega = -4 &\rightarrow \omega = -\kappa (1 - 2.49) + (-4) = -2.51 \end{aligned} \quad (\text{F7})$$

Again, the trajectories at the first two levels have not changed. It is now apparent that by rescaling x_3 and shifting its origin, we can choose arbitrary values for two out of the three parameters $\mu_3^{(0)}$, κ , and ω . However, we repeat that this is only possible as long as we do not measure x_3 on any objective scale by including its representations in the response model.

Equivalence classes (with equivalence defined as leading to invariant v) of parameter values are defined by the following *conservation laws*:

$$\begin{aligned} \kappa' \mu_3'^{(0)} + \omega' &= \kappa \mu_3^{(0)} + \omega \\ \kappa'^2 \vartheta' &= \kappa^2 \vartheta \\ \kappa'^2 \sigma_3'^{(0)} &= \kappa^2 \sigma_3^{(0)} \end{aligned} \quad (\text{F8})$$

Using these equations, we can now make κ and ω seemingly disappear by setting them to 1 and 0, respectively. Note that this does not amount to a simplification of the model: it is only a coordinate choice. In our example this means:

$$\begin{aligned}\omega &= -2.51 \rightarrow \omega = 0 \\ \mu_3^{(0)} &= 1 \rightarrow \mu_3^{(0)} = -1.51\end{aligned}\tag{F9}$$

This disappearance of κ and ω (and in the general case, all κ_i and ω_i) from the model may seem convenient. However, a danger here is to confuse coordinates with the underlying reality they describe. Crucially, it is impossible to discuss the choice of coordinate choice on higher levels in a model that lacks κ_i and ω_i . Yet this choice does not cease to exist only because one makes it implicitly instead of explicitly.

In summary, in order to avoid an overparameterized model when μ_i is not included in the response model, it is advisable to be explicitly arbitrary by fixing two out of $\mu_i^{(0)}$, κ_i , and ω_i , reflecting a choice of origin and scale on x_i . By contrast, when μ_i is included in the response model, all of $\mu_i^{(0)}$, κ_i , and ω_i can be estimated.