Human aging alters social inference about others' changing intentions

Andrea M.F. Reiter^{1,2,3}, Andreea O. Diaconescu^{4,5,6}, Ben Eppinger^{1,7,8}, Shu-Chen Li^{1,9}

¹ Chair of Lifespan Developmental Neuroscience, Faculty of Psychology, Technische Universität Dresden, Germany

² Department of Neurology, Max-Planck-Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

³ Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London, UK.

⁴ Translational Neuromodeling Unit, University of Zurich & ETH Zurich, Switzerland

⁵ Department of Psychiatry, University of Basel, Switzerland

⁶ Krembil Centre for Neuroinformatics, Centre for Addiction and Mental Health (CAMH), University of Toronto, Canada

⁷ Department of Psychology, Concordia University, Canada

⁸ PERFORM Centre, Concordia University, Canada

⁹ CeTI – Centre for Tactile Internet with Human-in-the-Loop, Technische Universität Dresden, Germany

Keywords: Theory of Mind, Social Cognition, Social Learning, Aging, Lifespan Development, Bayesian Modelling

Corresponding author: Andrea M.F. Reiter, andrea.reiter@tu-dresden.de

Word count: Abstract: 244, Introduction: 630, Discussion: 1465

35 pages (including references), 1 Table, 5 Figures

Conflict of Interest:

The authors report no conflict of interest

Acknowledgements: The authors thank A. Calder and A. Berthold for their assistance with data collection and all participants for taking part in the study. We are grateful to Dimitrije Marković for fruitful discussions of data analysis. This work was supported by a grant from the German Ministry of Education and Research (BMBF EMOTISK 16SV7243) awarded to SCL and by a grant from the German Research Foundation (DFG) awarded to BE (SFB 940/2 B7). AMFR and SCL also acknowledge support by the DFG for a related project (DFG RE 4449/1-1, DFG LI 879/22-1). During the period of manuscript preparation SCL was also supported by DFG's Excellence Strategy (EXC 2050/1 – Project ID 390696704). AOD acknowledges support by the Swiss National Foundation (SNF PZ00P3_167952) and the Krembil Foundation.

39 **Highlights**

- 40 - We used an advice-taking paradigm together with biologically plausible
41 hierarchical Bayesian modelling to characterize effects and mechanisms of
42 aging on learning about others' time-varying intentions
- 43 - Both a person's age and his/her counterpart's age affect how we learn about
44 others in changing social contexts.
- 45 - Older adults expected their social partner to change intentions more frequently.
- 46 - Older adults showed higher confidence in their own beliefs and changed them
47 less.
- 48 - Together, this led older adults to update their predictions about the correctness
49 of their social partner's advice more quickly.
- 50 - Older advisers were perceived as more unstable, but also more faithful than
51 younger advisers.

52

Abstract (224)

Decoding others' intentions accurately in order to adapt one's own behavior is pivotal throughout life. Yet, it is a process that is imbued with uncertainty since others' intentions are not directly observable and may change over time. In this study, we asked the question of how younger and older adults deal with uncertainty in dynamic social environments. We used an advice-taking paradigm together with biologically plausible hierarchical Bayesian modelling to characterize effects and mechanisms of aging on learning about others' time-varying intentions. We observed age differences when comparing learning on two levels of social uncertainty: the fidelity of the adviser and the stability of intentions. We found that, prior to having any experience with the adviser, older adults expected the adviser to change his/her intentions more frequently. They also showed higher confidence in such beliefs and were less willing to change their beliefs over the course of the experiment. This led them to update their predictions about observable outcomes (i.e., advice correctness) more quickly. Potentially indicative of stereotype effects, we also observed that older advisers were perceived as more volatile, but at the same time, more faithful than younger advisers. Together these findings offer new insights into the behavioral and algorithmic mechanisms underlying adult age differences in response to social uncertainty, putatively driven by aging-related changes in neuromodulation to be tested in future studies.

1. Introduction

Throughout our lives we often consider the advice of others when making important decisions: Which insurances do we take? Is it better to buy or to rent a house? When relying on someone's advice in such scenarios, it is pivotal to decode others' intentions accurately. However, this inference process is imbued with various sorts of uncertainty: Others' intentions cannot be observed directly but can only be inferred from their behavior and the outcomes associated with their advices. Moreover, others' intentions may also change over time, such that we have to concurrently update our predictions about how stable and trustworthy another person is. Such social skills might even become more essential the more we rely on the support of others in older age. However, there is reason to hypothesize that social inference in dynamic social contexts changes over the course of the lifespan: Evidence from studies investigating learning in non-social environments indicates that older adults compute and represent uncertainty in changing environments differently than younger adults (Hämmerer et al., 2018; Nassar et al., 2016). It is well conceivable that these age-related differences play an important role also when older adults navigate social environments, especially in light of evidence showing compromised socio-cognitive skills in human aging (Henry, Phillips, Ruffman, & Bailey, 2013; Reiter, Kanske, Eppinger, & Li, 2017).

To investigate this question, we build on recent advances in computational modeling of social interaction in younger adults (Behrens, Hunt, Woolrich, & Rushworth, 2008; Diaconescu et al., 2014). This research demonstrated that a biologically plausible Bayesian framework, the Hierarchical Gaussian Filter (C. Mathys, Daunizeau, Friston, & Stephan, 2011; C. D. Mathys et al., 2014), can be used to formally capture the process of inferring the time-varying states of others (Diaconescu et al., 2014). Learning about the other is captured at hierarchically organized levels, where higher levels represent general features about the stability of social environment – i.e. intentions and fidelity – and lower levels encode beliefs

about the probability of outcomes. Importantly, these levels are connected: Knowing whether to disregard an unexpected outcome (e.g., one misleading piece of advice by someone who was helpful in the past), or whether to take it seriously (e.g., revise one's belief about the adviser's intension), depends on the *precision* of one's beliefs. When the precision about the adviser's fidelity is low, the agent is more sensitive to new outcomes. Thus, the agent will be more inclined to incorporate prediction errors (PEs) – surprising outcomes not in line with one's prior belief of the adviser – into future predictions about the adviser. In this case, the agent will also adapt to potential changes in intentionality. Different neuro-transmitter systems are suggested to be involved in signaling PEs on different levels: Whereas low-level PEs about advice accuracy were associated with dopaminergic neurotransmission, higher level PEs were found in the cholinergic septum (Diaconescu et al., 2017). This dissociation on the level of neuromodulation has also been observed in non-social tasks (Iglesias et al., 2013). Importantly, both the dopaminergic and the cholinergic system undergo marked change during ageing (Bäckman, Nyberg, Lindenberger, Li, & Farde, 2006; Grothe, Heinsen, & Teipel, 2012).

In this study, we tested younger and older adults in a dynamic advice-taking paradigm (Diaconescu et al., 2017), interacting either with a younger or an older adviser (manipulated in a within-subject manner). Specifically, we investigated whether age groups differ with regards to the computation of two main quantities: their beliefs about i) the probability of receiving accurate advice and ii) their belief precision about the adviser's fidelity. Estimating the parameters of a hierarchical Gaussian filter from participants' decisions allowed for the individual quantification of i) learning about the fidelity of an adviser (i.e., the *probability* of accurate advice) at the lower level, and ii) how this fidelity changes over time (i.e., the *volatility* of intentions) at the higher level. Furthermore, we were interested in whether such

121 computations were altered as a function of the social identity of the adviser (own age vs. other
122 age adviser).

123

2. Methods

2.1. Sample

The study sample consisted of a total of 82 participants (37 young adults (YA), age range 18-30 years; 46 older adults (OA), age range 65-80 years). Participants were excluded from analysis in case they met the following a priori exclusion criteria: dementia screening (Nasreddine et al., 2005) below cut-off ($n=4$), self-report of a current, diagnosed psychiatric condition (depression $n=2$, borderline personality disorder $n=1$), neurological diagnosis ($n=1$ stroke, $n=1$ Morbus Parkinson), reporting problems in discriminating stimulus colors during or after the experiment ($n=2$), technical or human errors during data acquisition ($n=5$). Thus, the final effective sample consisted of 33 OA (20 female, mean age=71.97 years, $SD=3.95$) and 34 YA (21 female, mean=24.35 years, $SD=3.11$). Participants provided written informed consent prior to study participation. They received a show-up fee of 8.50€/hour for participating and could win extra-money in the decision-making experiment. Ethical approval in accordance with the Helsinki declaration was granted by the TU Dresden ethics committee.

2.2. Experimental Task and Procedure

The advice-taking task was applied behaviorally in a similar manner as reported in Diaconescu et al. (2017). During 120 trials, participants played a simple binary lottery game where they had to decide whether to bet on blue or green, one of which would win at the end of the trial. To come to this decision, participants could rely on two sources of information: 1) a pie chart displaying the true probabilities of winning associated with the two colors (blue vs. green pie charts: 75:25, 65:35, 55:45, 45:55, 35:65, and 25:75), and 2) an advice given by another person, which was presented via a video (Figure 1).

The Participants were instructed that the adviser had been playing the game in a session before and had more insight into which color would win on each trial. Because of a different

incentive structure, the adviser's willingness to provide helpful advice might change over the course of the experiment. Participants were told that it is thus important to constantly monitor the adviser's intentions in giving accurate or false advice and that they could win extra money in case they reached a silver or a gold bar by making accurate predictions. The validity of the advice was predetermined based on the dominant strategy that previous advisers actually used in this task (see panel 1b and Diaconescu et al., 2014). In accord with previous studies (Behrens et al., 2008; Diaconescu et al., 2014; Diaconescu et al., 2017), the advice was manipulated to be stable and valid in a first phase of the experiment (40 trials), volatile in a second phase (40 trials) and stable and valid again in a last phase (40 trials) of the experiment. In order to investigate potential differences in social inference as a function of the adviser's age group (YA vs. OA), all participants came to the lab twice, once interacting with an older adviser, once interacting with a younger adviser (the order of the two within-subject sessions was counterbalanced across participants). Advisers were gender-matched to the participant. Participants were debriefed about the cover story after the second testing day. On average, the two testing sessions were $m=6.83 \pm 1.35$ days apart (range 3-14 days). The input structure was held constant across participants and simulated based on an optimal Bayesian learner. Two slightly different versions of the input structure, created by exchanging the two stable phases, were used for the two testing days of the within-subject assessment (counterbalanced in order).

Note that this experimental set-up strongly builds on a previous validation study in young adults (Diaconescu et al., 2014). Notably, Diaconescu et al. also included a control task, closely matched to the social task described above, in which any intentionality was removed from the adviser. In this control task, players were shown to rely significantly more on the non-social, and less on the social cue, lending evidence to the fact that the social intentionality manipulation in the social inference task adopted here is indeed effective.

Figure 1. Task and Raw Choice Data.

- A. An example trial sequence. In a simple, roulette-like decision-making game, participants decide between green and blue on each trial, one of which is the winning color. Before making a choice, participants see both (i) a pie chart indicating the true winning probability of the colors in the respective trial and (ii) the advice of another person, shown via video. After making a decision, participants receive feedback indicated by an asterisk appearing on the winning part of the pie chart. By observing the congruency of the advice with the outcome, participants can update their belief about how likely the adviser is to give accurate advice (i.e., the adviser's fidelity), and how much this changes over time (i.e., the adviser's volatility).*
- B. Choice data for younger and older adults. The figure shows the proportion of following the advice as a function of participants' age group and the phase of the experiment (stable, volatile, stable). We plot predicted values of a mixed model using the predict function in R.*

All data were analyzed using R, MATLAB and SPSS. Mean values of the decision whether to follow the advice or not were subjected to a repeated-measures ANOVA with two within-subject factors, phase (first, pre-volatile, stable phase, second, volatile phase, third post-volatile, stable phase) and adviser's age group, while the factor participants' age group was the between-subject factor. Trial-by-trial decisions on whether to follow an advice were subjected to a logistic mixed effects model (Bates, Mächler, Bolker, & Walker, 2014) with the following fixed effects: participants' age group, adviser's age group, phase and trial number within-phase (in order to mimic trial-by trial learning about the adviser's changing intentions). Testing session and version of the input structure were included as nuisance regressors. We included a maximal random effects structure as recommended by simulation studies (Barr,

Levy, Scheepers, & Tily, 2013). Degrees of freedom were approximated and p -values obtained using the R package *afex* (Singmann et al., 2016). Post-hoc tests were done using the package *lsmeans* (Lenth, 2016). None of our participants met our a-priori exclusion criterion of missing more than 10% (>12) of all trials.

2.3. *Hierarchical Bayesian Social Inference Model*

In accord with previous work (Diaconescu et al., 2014; Diaconescu et al., 2017), several cognitive models were applied to investigate age differences in updating beliefs about other people's (time-varying, uncertain) intentions (compare Figure 2a for an overview). The computational framework adopted in this study is based on Bayesian theories of brain function, which suggest that the brain maintains, and continuously updates, a model of the environment and uses this model to infer the hidden causes of the environmental inputs it receives (Friston, 2009). Note that the choice of computational models applied here was based on our previously published validation study in young adults, using a very similar task as well as the "Bayesian Brain" theory to explain social inference behavior in young adults (Diaconescu et al., 2014). Please see this model validation study for more details regarding the class of models described in the following.

Learning models. The Hierarchical Gaussian Filter (HGF), a hierarchical Bayesian model of learning under uncertainty and volatility (C. Mathys et al., 2011; C. D. Mathys et al., 2014) is one formal implementation of the Bayesian brain theory. According to the HGF, an agent uses environmental inputs (in our paradigm this is the congruency/incongruency of advice and outcome) to make inferences on hidden states of its environment (in our case, the time-varying intentions of the adviser). The HGF allows for the estimation of these hierarchically coupled states which characterize participants' learning about two types of environmental statistics—namely the probability and the volatility of the adviser giving correct advice.

222 The hidden states x_i on each level evolve as Gaussian random walks. An exception is the first
 223 level, $i = 1$, where x_1 represents advice accuracy in a categorical manner (a piece of advice
 224 in a trial k can either be accurate ($x_1^{(k)} = 1$) or inaccurate ($x_1^{(k)} = 0$) as input. Thus, the
 225 probability advice is accurate is denoted by Equation 1, and depends on the next higher
 226 (unbounded) state x_2 via the logistic sigmoid transformation $s(\cdot)$ (Equation 2).

$$227 \quad p(x_1|x_2) = s(x_2)^{x_1}(1 - s(x_2))^{1-x_1} = \text{Bernoulli}(x_1; s(x_2)) \quad (\text{Eq. 1})$$

228

229 where

$$230 \quad s(x) \stackrel{\text{def}}{=} \frac{1}{1 + \exp(-x)}. \quad (\text{Eq. 2})$$

231 $x_2^{(k)}$ can thus be interpreted as the adviser's tendency to deliver accurate advice (i.e., the
 232 adviser's fidelity). The step size of the Gaussian random walk of x_2 depends on the next higher
 233 level x_3 :

$$234 \quad p(x_2^{(k)} | x_2^{(k-1)}, x_3^{(k)}, \kappa, \omega) = \mathcal{N}(x_2^{(k)}; x_2^{(k-1)}, \exp(\kappa x_3^{(k)} + \omega)) \quad (\text{Eq. 3})$$

235

236

237 At this third (highest) level, the state, $x_3^{(k)}$, denotes the beliefs about the volatility of the adviser
 238 (i.e., the rate at which the adviser's intentions change).

239 $x_3^{(k)}$ evolves as a random walk (Eq. 4), with a constant variance (ϑ)

$$240 \quad p(x_3^k | x_3^{k-1}, \vartheta) = \mathcal{N}(x_3^k; x_3^{k-1}, \vartheta) \quad (\text{Eq. 4})$$

241 As defined in Eq. 3 and Eq. 4, the HGF learning model includes three free parameters, namely
 242 κ , ω , and ϑ , which determine how these states evolve as a function of the inputs the participant
 243 receives (pie chart, advice, trial outcome): The parameter κ represents the coupling between
 244 the second and the third level or how much does the learner utilizes the estimated adviser

volatility at the higher level when predicting the adviser's fidelity. The parameter ω is the constant or participant-specific component of log-volatility on the second level. It captures the degree to which the learner updates his/her beliefs about the adviser's fidelity $x_2^{(k)}$, independent of the volatility. The parameter ϑ represents the meta-volatility or the variance of the adviser's volatility.

Variational inversion of the HGF gives rise to update equations. They provide approximately Bayes-optimal rules for trial-by-trial updating of an agent's beliefs, given this agent's particular set of parameter values. The posterior means at each level i are proportional to the prediction error below, weighted by a precision ratio:

$$\mu_i^{(k)} \propto \frac{\hat{\pi}_{i-1}^{(k)}}{\pi_i^{(k)}} \delta_{i-1}^{(k)} \quad (\text{Eq. 5})$$

where precision $\pi_i^{(k)}$ is inverse variance $\pi_i^{(k)} = 1/\sigma_i^{(k)}$

The learning rate about the adviser's fidelity is transformed as below (as in Diaconescu et al. 2014; please also refer to the supplement of Iglesias et al., 2014 for more details):

$$\alpha_2 = s(\mu_2^{(k)}) \left(1 - s(\mu_2^{(k)})\right) \sigma_2^{(k)} \quad (\text{Eq. 6})$$

Figure 2b provides a summary and a graphical depiction of the HGF (for details on the algorithmic implementation, see (Diaconescu et al., 2014; C. D. Mathys et al., 2014). Here, we consider two versions of the HGF: the full, three-level version (as described above) as well as a reduced, two-level version, which lacks the above described third level, and thus does not postulate hierarchical learning about the adviser's volatility (Diaconescu et al., 2014). As a further (simpler) control model, we fit the Rescorla-Wagner (RW) model, which is a standard reinforcement learning model. In the widely used RW model, values of choice options are updated using a prediction error which is weighted by a learning rate. Other than the HGF, the RW model does not employ a hierarchy of hidden states, but a single state v (value of the advice) and one free parameter, the individual learning rate, α , which is constant across trials

(i.e., does not vary according to the environmental volatility). Bayesian Model Selection allows to compare these model classes and for testing whether the two age groups differ with regards to which model explains their behavior best (Rigoux, Stephan, Friston, & Daunizeau, 2014; Stephan, Penny, Daunizeau, Moran, & Friston, 2009).

Response models. Additionally, and in accord with previous work (Diaconescu et al., 2014; Diaconescu et al., 2017), we consider different versions of response models in order to map beliefs to decisions. First, to test whether participants rely only on i) the pie chart information, ii) the adviser's information, or iii) integrate both sources of information, we introduce a social weighing factor ζ which weighs beliefs on the first level of the HGF before they are translated into choices. We fit versions of the response model with $\zeta=0$ (i.e., agent exclusively relies on the pie chart probability), $\zeta=1$ (i.e., agent exclusively relies on the advice) and ζ as a free parameter (compare Eq.7). Fitting ζ as a free parameter corresponds to the assumption that an agent integrates both sources (social and non-social) of information, and that individuals differ to which degree they weigh either source.

Thus, participants generated a combined belief by weighting the posterior expectation of the advice accuracy with the pie chart probability (transformed into advice space, i.e. reflecting the color the adviser held up in this trial).

$$\mu_b^{(k)} = \zeta \hat{\mu}_1^{(k)} + (1 - \zeta) \tilde{c}^{(k)} \quad (Eq. 7)$$

Here, ζ was the degree of reliance on the social advice and served to balance $\hat{\mu}_1^{(k)}$, the participant's current belief that the adviser offers valid advice, against $\tilde{c}^{(k)}$, the probability (as signaled by the visual pie chart) of the recommended advice being correct. These beliefs were translated into choices (following the advice or not), using a sigmoid function:

$$p(y^{(k)} = 1 | \mu_b^{(k)}) = \frac{\mu_b^{(k)\beta}}{\mu_b^{(k)\beta} + (1 - \mu_b^{(k)})^\beta}, \quad (Eq. 8)$$

where β represents the inverse decision temperature (i.e., how tightly choices are mapped onto beliefs). In further variations of the response model, we included models with a participant-specific, but fixed β or varied β in time as a function of the estimated adviser's volatility. In the latter variant, the sigmoid function becomes less steep when the estimated volatility of the adviser increases according to $\exp(-\mu_3^{(k-1)})$.

Note that the latter variation is only meaningful in the full 3-level HGF, as the other perceptual models do not include volatility estimates (and thus, decision noise in the response models cannot be influenced by volatility). Again, we used Bayesian Model Selection to compare these response model classes and to compare whether the two age groups differ with regards to which information (social vs. non-social) influences their decisions more.

In sum, we varied 3 perceptual models (full 3-level-HGF, reduced 2-level-HGF, RW), with 3 variations of the social weighing factor ζ , and with, in case of the full HGF, 2 variations of the decision noise parameter β , totaling 12 models (Figure 2a gives an overview of the model space, and Table 1 lists priors of the free parameters). After model fitting, belief trajectories of all participants were visually inspected. In case of implausible trajectories, participants were excluded from all further modeling analyses ($n=3$).

Model Fitting. To fit the above described learning and response models to our participants' choice data, we used the freely available HGF toolbox 2.0 (C. Mathys et al., 2011; C. D. Mathys et al., 2014) as part of the Translational Algorithms for Psychiatry-Advancing Science tools (<https://www.tnu.ethz.ch/en/software/tapas.html>). Maximum a posteriori estimates of the parameters are obtained using variational Bayesian inversion (for details see Mathys et al., 2011). The update equations take the form of precision-weighted prediction errors following a form similar to an extended Kalman filter and are hence analytically tractable. All free parameters ($\vartheta, \omega, \kappa, \zeta, \beta$), the state variables μ_1 and μ_2 , and their prior values were estimated

for each participant by using a quasi-Newton minimization algorithm. For details of the update equations and the variational Bayesian inversion scheme, see Mathys et al. (2011), Daunizeau et al. (2010).

Bayesian Model Selection. We used the freely available VBA toolbox to perform random-effects Bayesian Model Selection (BMS) based on the negative variational free energy as an approximation to the log model evidence (Daunizeau, Adam, & Rigoux, 2014). BMS tests for and compares plausibility of the competing models in the comparison set, given the data. Exceedance probabilities give the probability that one model is more likely to have generated the data, relative to the other models in the model space (Stephan et al., 2009). BMS also accounts for the complexity of a model (i.e., the number of free parameters). We additionally employ family selection across the above-described perceptual as well as response model variations, by pooling models that share a common feature (Penny et al., 2010), i.e. across all response models and across all perceptual models, respectively (see Diaconescu *et al.*, 2014 for more details of this application in the context of social learning).

Parameter Recovery. We run a parameter recovery analysis for all free parameters of the winning model. To this end, we used the empirically derived (true) parameters to obtain estimated parameters for simulated agents based on the winning model. In other words, we simulated responses based on all participants' maximum *a posteriori* estimates of the parameters, and then fit the model to those simulated responses in order to test whether we could recover the same parameter estimates. We found significant correlations between true and estimated parameters for all parameters of the winning model (see S-Figure 1).

Analyzing Age Effects on Model Parameters. All free parameters of the winning model were extracted and analyzed as a function of i) participant's age group, ii) adviser's age group, iii) and the interaction of both, using a repeated-measures MANOVA on all parameters of the

model. This was followed up by post-hoc t-tests. Additionally, we extracted belief trajectories of i) the second level of the HGF, i.e. states reflecting the evolution of beliefs about the adviser's fidelity, and ii) the third level of the HGF, i.e. states representing how beliefs about the volatility of the adviser develop over the experiment. This analysis was done to gain more specific insight into the dynamics of belief evolution as a function of volatility, as well as the age differences therein. These trial-by-trial trajectories were subjected to a hierarchical mixed effects model with the following predictors: age group, volatility phase (stable, pre-volatile; volatile; stable post-volatile), trial number within phase, as well as all interactions. In a second set of regression models, we aimed to test the hypotheses that the adviser's age group alters the participants' beliefs about the adviser. Thus, the factor "age group of the adviser" was included as a regressor to the above-described models. Session and version of the input structure were included as nuisance regressors in all models. Participant-specific slopes and intercepts for the within-subject factors of interest, as well as their interactions, were estimated as random effects of the model. Degrees of freedom were approximated using Kenward-Roger's degrees of freedom method, as implemented in the R packages *afex* (Singmann et al., 2016) and *lmerTest* (Kuznetsova, Brockhoff, & Christensen, 2015). Post-hoc tests were performed using the package *lsmeans* (Lenth, 2016).

Code Accessibility. MATLAB code to fit the Hierarchical Gaussian filter and Reinforcement Learning Models is available at <https://github.com/translationalneuromodeling/tapas/tree/master/HGF>. Code to perform Bayesian Model Selection can be retrieved via <https://mbb-team.github.io/VBA-toolbox/>.

361

362 **Table 1. Prior means and variances for estimating HGF and RW model parameters. Note**
363 *that prior means and variances were chosen based on (Diaconescu et al., 2014). κ , ϑ , ζ ,*
364 *$\mu_2^{(k=0)}$, $\mu_3^{(k=0)}$ and α are estimated in logit space. The variances of σ and the decision noise*
365 *parameter β are estimated in log space. Whereas the prior variances for all parameters are*
366 *set to be rather broad, we selected a shrinkage prior mean for the decision noise parameter*
367 *β , such that behavior is explained more by variations in the rest of the parameters rather than*
368 *decision noise.*

369

	mean	variance
HGF		
mu2_0	0	1
sigma2_0	0	1
mu3_0	1	1
sigma3_0	1	1
kappa	0	1
omega	-2	1
theta	0	1
zeta	0	1
beta	48	1
RW		
v_0	0.5	1
alpha	0.4	1

370

B Full HGF with integrated response model

C Model stability

D Exceedance probability

Figure 2 Model Space and Model Selection.

- A. *Model Space.* The model space was set up based on a 3x2x2 structure. The first factor represents the perceptual/input model families (3-level HGF, reduced 2-level HGF, RW). Factor 2: In the case of the HGF, two response models were introduced: Decision noise when translating beliefs to decisions was either dependent on the trial-by-trial volatility ("volatility model"), or not ("decision noise model"). This cannot be formalized for the reduced 2-level HGF or RW, which does not describe computation of volatility. Factor 3: Different response models describe differential weighting of social (advice) and non-social (pie chart) information: These models propose that participants' decisions are based on: both advice and pie chart information ("integrated: Cue and advice"), only advice ("Reduced: advice") or only pie chart ("Reduced Cue"). Thus, at the bottom, the single models entering model selection (12 models in total) are shown.
- B. *Graphical depiction of the hierarchical Gaussian filter with an integrated response model (winning model).* Beliefs are organized as a hierarchy ($x_1 \dots x_3$). x_1 is the binary sensory input indicating the accuracy (true / false) of the current piece of advice, x_2 is the current probability that the advice is correct (the fidelity of the advice) and x_3 is the current volatility of the adviser's intentions. There are five ($\vartheta, \omega, \kappa, \zeta, \beta$) free parameters of the model. The parameter κ determines how strongly x_2 is influenced by x_3 , ω represents the tonic learning rate on the second level. The parameter ϑ is the meta-volatility parameter, governing updating speed on the third level. To map beliefs onto decisions y , the response model computes the probability of the outcome given both the social (advice) and the non-social (pie chart) information. Decisions y are influenced by the parameter ζ which determines the weight of the social cue relative to the non-social cue, and by β , which captures choice stochasticity.
- C. *Model stability across conditions.* As each participant completed the task twice (within-subject design), once interacting with an OA adviser, once interacting with an YA adviser, we tested for model stability across both conditions. Log model evidences do not differ as a function of condition (left panel). BMS reveals evidence of no difference as a function of adviser's age group condition (right panel). Thus, the strategy employed to solve the task does not differ when interacting with a YA vs. OA adviser.
- D. *Model Selection across the single models Using Family BMS,* we pooled perceptual (middle panel) and response models (right panel) of the same kind and compared their plausibility given the data. Exceedance probability indicated that evidence was highest for the 3-level-hgf with the volatility-influenced decision noise in the perceptual model set, and for the integrated response model using advice and cue information. 3level-hgf-vol: M1, M2, M3; 3level-hgf-dc: M4, M5, M6; 2level-hgf: M7, M8, M9; rw: M10, M11, M12. Integrated advice&cue: M1, M4, M7, M10; advice: M2, M5, M8, M11;

3. Results

3.1. Raw Choice Data

In the hierarchical logistic mixed effects model, we found a significant effect of phase, trial number and their interaction ($\chi^2 > 63.97$ all p s $< .001$, see Figure 1) on participants' trial-by-trial decisions of whether to follow an advice or not. This indicates that the social volatility manipulation alters the participants' tendency of whether or not to follow the adviser. Figure 1b indicates that, as expected, the tendency to follow the advice decreases in the middle, volatile phase. Indeed, in both age groups, the volatile phases differed from the stable phases in terms of following behavior ($z > 3.282$, $p < .001$). Thus, both YA and OA adapt their behavior according to the changing volatility in the social environment. With respect to age differences, we also observed a significant effect of participants' age group ($\chi^2 = 6.18$, $p = .013$) and an interaction effect of participants' own age group with phase ($\chi^2 = 12.60$, $p = .002$, Figure 1b). Thus, YA and OA seem to differ with regards to how they deal with the social volatility manipulation in our task. Post-hoc tests revealed significant age differences in the first, stable phase ($z = 3.57$, $p < .001$), and the post-volatile phase ($z = 2.161$, $p = .03$), such that OA followed the adviser less than YA in these two phases. No significant difference was observed in the middle, volatile phase ($z = 1.05$, $p = .291$), (compare Figure 1b).

3.1. Model-Based Analysis – Model Selection.

First, we tested for model stability across the two sessions (Rigoux et al., 2014). We found evidence for model stability across both advisers' age group conditions (exceedance probability, $XP = 1$, protected exceedance probability $PXP = 1$, figure 2c). We also found evidence for across-session (testing day 1 and 2) stability ($XP = 1$, $PXP = 1$). Thus, in a next step, we summed up the log model evidences and tested for age group differences regarding which model explained the data best. BMS revealed high evidence that both age groups were characterized best by the same model (posterior probability that the groups have the same

model frequencies=.99). The winning model for both groups together was the full HGF with volatility-influenced decision noise including the social weighing parameter (M_1 , $XP=1$, $PXP=1$, see Figure 2b and d). Family selection (i.e., pooling across models which share a common feature) across all perceptual model classes revealed that the perceptual model family the full HGF with volatility-influenced decision noise had highest model evidence ($XP=0.99$, Figure 2d). Family selection across all response model classes showed that the family comprising models with the social weighing factor ζ as a free parameter, outperformed the other response models ($XP=1$, Figure 2d). This indicates that participants integrate social and non-social information instead of relying exclusively on one of the two sources. When running model selection for both groups separately, we obtained the same results. These results replicate previous studies in YA (Diaconescu et al., 2014; Diaconescu et al., 2017). BMS also answers questions about age differences regarding the strategy used to solve task by quantifying evidence for both the null and alternative hypothesis. More specifically, BMS results demonstrate that both age groups infer two quantities at the same time: the accuracy of the advice, and the volatility of the adviser's intentions, with the latter being used in the decision process. BMS also shows that both age groups integrate social and non-social information to come to their decisions.

3.2. Model Bases Analysis – Parameter Comparison

We extracted and subjected all seven free parameters of the winning model (i.e., of M_1 , the full HGF with volatility-influenced decision noise in the response model) to a repeated-measures MANOVA to test for the effect of age group, adviser's age and their interaction. This revealed a significant main effect of one's own age group on the HGF parameters ($F=6.40$, $p<.001$). Post-hoc tests indicated that this difference was driven by group differences on the initial values of μ_2 ($F=23.83$, $p<.001$, Figure 3) and μ_3 ($F=6.90$, $p=0.01$, Figure 3). OA also had lower a priori beliefs about the fidelity of the adviser. A t-test against zero showed

that this was due to YA being optimistic about the fidelity of the adviser (μ_2 significantly positive in both adviser conditions, $t=>3.73$, $ps<.001$), whereas OA seem to hold a more neutral prior regarding the adviser's fidelity (μ_2 not significantly different from zero in both adviser conditions, $t<=1.71$, $ps>.09$). OA also had higher a priori beliefs about the volatility of the adviser (Figure 3). Interestingly, we further found a significant age group difference in θ , the meta volatility parameter that governs the speed of volatility updating ($F=6.69$, $p=.012$). OA had lower values on the meta-volatility parameter (see Figure 3), suggesting that they update volatility more slowly than YA. None of the other learning and response parameters differed significantly between age groups ($F<= 1.92$, $p>.171$).

Figure 3 Parameter comparisons. Parameter comparisons between age groups showed significant effects of age group on the priors of μ_2 and μ_3 . Older adults had lower a priori beliefs about the fidelity of the adviser (μ_2) and higher a priori beliefs about the volatility of the adviser (μ_3). Additionally, the meta-volatility parameter (theta) differed as a function of age group. Older adults updated volatility less than did younger adults. No significant difference was observed with regards to parameters of the response model. In the cat-eye plots, dots represent medians, asterisk the mean and vertical lines the inter quartile range.

3.3. *Model-Based Analysis – Age Differences in Belief Trajectories*

After we had established age differences in prior beliefs as well as in the meta-volatility parameter, we were next interested to see how these differences influenced belief evolution on the different levels of the learning hierarchy. Whereas the analysis of the free parameters reported above provides insights in terms of summary statistic of the psychological processes of interest, in this analysis we were specifically interested in how beliefs and age differences therein changed dynamically over the experiment, as a function of volatility.

Thus, in a next analysis step, we subjected participant-specific trial-by-trial trajectories of the evolution of beliefs on different levels of the hierarchy (μ_3 , π_3 , μ_2), to a mixed model with the following fixed effects: one's own age group, volatility phase, trial number within phase.

Trial-by trial beliefs about the adviser's volatility and their precision. As per our volatility manipulation, the model predicting the evolution of volatility beliefs (μ_3) revealed significant main effects of trial number within phase and phase, as well as their interaction (all $X^2 > 121.64$, all $p < .001$). These proof-of-concept findings confirm that volatility estimates are updated according to our experimental manipulation throughout the course of the experiment in both age groups. Regarding age differences, we found a significant main effect of age group ($X^2 = 6.1706$, $p = .013$), with higher volatility estimates in OA as compared to YA. This is in line with the finding of older adults' higher prior on volatility. Note that the main effect of age, without significant interaction effects with trial or phase, suggests that the a priori stronger beliefs about volatility of the adviser in OA persist over the course of the task, even in the face of reduced environmental volatility (as per our manipulation in phase 1 or 3).

In OA, the reduced meta-volatility (ϑ) described above (compare Figure 3) leads to volatility beliefs being adapted less, even in the face of environmental stability. The HGF posits that belief updating is governed by the precision of beliefs: If precision is low, new input should be

given more weight and belief updating should be upregulated; if precision is high, more weight is assigned to one's prior beliefs. Thus, reduced volatility updating might be related to OA's higher precision on the third level. To test this hypothesis, we subjected trial-by trial precision estimates (π_3) to the same regression model as described above. As a proof of-concept finding, this revealed significant main effects of trial number within phase and phase, as well as their interaction, (all $X^2 > 292.34$, all $p < .001$). These findings suggest varying certainty about one's beliefs in phases of different volatility, throughout the course of the experiment. Most interestingly, however, we also observed a main effect of age group ($X^2 = 5.31$, $p = .021$) which was due to OA displaying *higher* precision in their beliefs about volatility than YA (Figures 4a and 4c). Moreover, a significant interaction of trial number within phase * age group was observed ($X^2 = 5.84$, $p = .016$, Figure 4a). This suggests age differences in the evolution of precision within the volatility phase. Examining this further by comparing the slopes of trial number on trial-by-trial precision estimates per group, we observed that the association was more negative in OA than in YA (z ratio = 2.42, $p = .016$) showing preserved modifiability of precision estimates within phase, as a function of changing volatility in the environment.

519

Figure 4 Age Differences in the Evolution of Beliefs, Precision and Learning Rate on different levels of the hierarchy

A. Left: Evolution of volatility estimates (ie., μ_3 of the HGF) through the course of the experiment as a function of participants' own age group. Right: Evolution of precision about volatility.

B. Left: Evolution of fidelity estimates (ie., μ_2 of the HGF) over the course of the experiment as a function of participants' own age group. Right: Learning rate of the second level as a function of participants' own age group.

C. Mean values about the precision on volatility (third level) and the learning rate on fidelity (second level).

530

3.4. Beliefs about the adviser's fidelity

In a Bayesian framework, as formalized in the HGF, how we estimate volatility impacts on learning on the lower level (in our case: learning about the fidelity of the adviser, see figure 2a and methods). If the world is perceived as more instable, we should update the lower level beliefs about probability more quickly than when it is perceived as stable. Given the

significantly higher estimates of volatility in OA, we predicted that this should also influence learning about fidelity in OA vs. YA on the lower level.

The model predicting trial-by trial beliefs about the adviser's fidelity, (μ_2) again showed a significant effect of phase, trial number within phase (i.e., learning over the course of one phase), and their interaction (all $X^2 > 166.22$, all $p < .001$). Regarding age group differences, in line with our raw choice data analysis, we observed a significant interaction of age group and phase ($X^2 = 6.81$, $p = .033$). Post-hoc tests showed that the only significant difference between age groups was in the last phase, where OA had higher estimates of the adviser's fidelity than YA (z ratio = 2.29, $p = .022$). We also observed a significant interaction effect of age group and trial number ($X^2 = 7.39$, $p = .006$). Comparing the slopes of the effect of trial number between groups revealed that OA changed beliefs more on a trial-by-trial basis (z ratio = 2.80, $p = .005$). The computational quantity that directly captures this speed of updating on a trial-by-trial level is the learning rate on the second level, α_2 . Thus, the finding that OA change beliefs more on a trial-by-trial basis gives rise to the hypothesis of differences in the second-level learning rate. Repeating the above-described model to predict trial by trial fluctuations of the learning rate α_2 , we found a higher learning rate in OA compared to YA ($X^2 = 3.98$, $p = .046$). Additionally, we observed a trial number * age group interaction ($X^2 = 5.36$, $p = .021$). Interestingly, the learning rate co-varied more with trial number in OA (z ratio = 2.36, $p = .02$) as compared to YA. The latter was also moderated by phase ($X^2 = 6.22$, $p = .045$). As can be seen from Figure 4, learning rate fluctuations differed most strongly between groups in the first, and the second, unstable phase. There were no age-related significant differences on the precision on this level.

In sum, OA started into the experiment with lower expectations about the adviser's probability to deliver accurate advice. However, due to their pronounced and stable beliefs about intentions being instable (i.e. higher μ_3 estimates throughout, see figure 4a), along with

a reduced tendency to update them (i.e., reduced meta volatility parameter θ), OA were also well prepared to adjust these beliefs over the course of the experiment (i.e., higher learning rates regarding the fidelity, μ_2).

3.5. *Effects of the adviser's age group on volatility and fidelity estimation*

In a next analysis step, we were interested in whether the age group of the adviser also influenced beliefs about the adviser on the different levels of the learning hierarchy. Thus, we included the factor age group of the adviser into the above-described models to predict i) the adviser's fidelity, ii) the adviser's volatility. Please note that the input structure was predetermined and, in reality, advisers did not differ as a function of their age group (YA vs OA). Despite this fact, the adviser's age group did have an influence on how participants estimated volatility throughout the experiment: We observed a significant interaction of one's own age group and adviser's age group moderated by phase ($X^2=7.11$, $p=.029$). Post-hoc tests showed that significant adviser age differences in volatility we described above were only present when comparing the adviser condition in OAs ($X^2=2.24$, $p=.03$), but not when comparing the adviser conditions in YA (z ratio=2.23, $p=0.26$). This was due to YA displaying higher volatility estimates when interacting with OA as compared to YA, particularly in middle, volatile ($z=2.61$, $p=.009$), and the last phase ($z=2.18$, $p=.029$).

Figure 5 Differences in Belief Evolution as a Function of the Adviser's Age Group.

A. Beliefs about the adviser's volatility as a function of participants' age, advisers' age and phase. B. YA assigned significantly more fidelity to OA advisers in the first, stable phase than to YA advisers. No significant adviser's age effect on the belief about fidelity was found in OA.

The model also revealed a significant interaction of adviser's age group and trial ($\chi^2=6.3047$, $p=.012$) and of one's own age group, adviser's age group, phase and trial ($\chi^2=6.7406$, $p=.035$). Figure 5A shows that whilst YA updated their volatility beliefs more when interacting with OA in all phases, OA showed this tendency only in the last phase of the task.

The adviser's age group also had an effect on how fidelity was estimated (Figure 5B). We again observed an interaction of one's own age group x adviser's age group x phase on beliefs about fidelity of the adviser (μ_2 , $\chi^2=6.492$, $p=.039$). Interestingly, this interaction effect was driven by adviser age group effects in YA who assigned significantly more fidelity to OA advisers in the first, stable phase ($z=3.006$, $p=.003$, see Figure 5). Contrary to this effect in YA, no significant difference in μ_2 as a function of adviser's age was found in OA.

4. Discussion

In this study, we employed a hierarchical Bayesian modelling framework to describe social learning about others' intentions in younger and older adults. We found that both our own age as well as the counterpart's age affect how we learn about others in dynamically changing social contexts. In general, model selection suggests that both age groups use the same hierarchical Bayesian model to infer the changing intentions of an adviser, and that both younger and older adults integrate social and non-social information to decide on what to choose. Interestingly, differential age group effects emerged when comparing learning about different sources of social uncertainty. Older adults were more confident about their appraisals of environmental volatility, namely that the adviser's intentions are instable. At the same time, they were more uncertain and thus updated their beliefs more quickly when making predictions about the adviser's fidelity (i.e., his probability of offering accurate recommendations). We find that, before learning about the adviser, older adults *a priori* expected the adviser to be more volatile, and thus changed his/her intentions more over the course of the task than younger adults. Older adults showed higher confidence (larger higher-level belief precision) in these beliefs and were less willing to change them over the course of learning, as indexed by the significantly lower meta-volatility parameter ϑ values. From a Bayesian perspective, when confronted with an unexpected outcome (e.g., a misleading piece of advice from someone you trusted before), knowing whether to disregard it (e.g. to consider it a sign of incomplete information) or whether to take it seriously, as evidence that the adviser's intentions have changed, depends on the precision of one's higher-level beliefs. More specifically, when the precision about the adviser's changing intentions is high, as in the case of older adults here, agents expect the adviser's strategy for providing helpful suggestions to change more, and thus are sensitive to incoming prediction errors, which could reveal a contextual change. This does not necessarily mean that they react overly strongly to

any new prediction errors, but rather because they expect the adviser to change, adapt more to global rather than local changes in advice validity and thus adjust their beliefs about contextual changes in the adviser's fidelity more. Indeed, in line with this interpretation, we found that older adults, who were equipped with a more precise belief that the adviser's strategy will change, adapted their learning rate and with it their beliefs about the probability of the advice validity more quickly.

Interestingly, these results are in accord with previous studies in the non-social domain and extend them to social inference. Our finding of older adults' stronger prior beliefs about volatility is consistent with very recent findings of age-related differences in learning in uncertain environments (Hämmerer et al., 2018). In this study, OA were similarly found to overestimate reversal probabilities (and thus volatility) in a reversal learning paradigm. Our observation of a *reduced* meta-volatility parameter and thus more precise higher-level beliefs about the instability of adviser's intentions is in line with the proposal of reduced Bayesian belief updating and the maintenance of a "more robust model of the world" in healthy aging (Moran, Symmonds, Dolan, & Friston, 2014). Similar to non-social contexts, we find that higher-level beliefs about intentions are more precise and therefore less susceptible to change in older compared to younger adults. This may equip older adults with increased adaptability to changes in the other's trustworthiness but also increased rigidity about how quickly that may change in time. Interestingly, our findings also resonate with a very recent study applying the same Bayesian framework as here to demonstrate that impressions about morally bad agents are more volatile and are more rapidly updated than impressions about good agents (Siegel, Mathys, Rutledge, & Crockett, 2018). Older adults in this study had a less positive (albeit more realistic) prior view on the adviser than younger adults and were more inclined to update this view than younger adults. An evolutionary adaptive interpretation of instable

bad and more stable good impressions is that this promotes cognitive flexibility in the service of building richer models of potentially bad others and thus helps forgiving (Siegel et al., 2018).

Parallel to age differences in non-social learning tasks, the findings here might indirectly speak to the notion that social inference is realized by means of at least partially overlapping processes with learning about other (non-social) features of the environment (Behrens et al., 2008). However, recent evidence also indicates domain-specificity in social development, namely that social and non-social cognition age differentially (Reiter et al., 2017). Future studies could include a direct non-social comparison task (compare e.g. (Garvert, Moutoussis, Kurth-Nelson, Behrens, & Dolan, 2015) to investigate the specificity of the social effects. A recent study used the same Bayesian modeling approach to unravel neural correlates of social learning (Diaconescu et al., 2017). The authors showed that learning at different levels of the hierarchy mapped onto distinct neuro-modulatory systems: Learning about the stability (or instability) of the others' intentions (as reduced in older adults in the study at hand) was associated with activation in the cholinergic basal forebrain. Learning about the adviser's fidelity (as pronounced in the older adults in the study at hand) activated the dopaminergic midbrain. These differences between younger and older adults may be related to aging related changes in neuromodulation. Of note, both the cholinergic system and the dopaminergic system undergo marked change over the course of the lifespan (Bäckman et al., 2006; Dreher, Meyer-Lindenberg, Kohn, & Berman, 2008; Grothe et al., 2012; Schliebs, 2006 #5084). Evidence from earlier computational studies of aging neuromodulation (Li et al., 2001; Li, von Oertzen & Lindenberger, 2006) suggest that older adult's deficient neuromodulation (particularly dopaminergic modulation) would result in noisier neural information processing, with the computational consequences of reduced precision of neurocognitive representations of learning experiences and intraindividual cognitive variability (cf. empirical data from MacDonald et al., 2012). These effects in turn could contribute to age differences in making

cognitive and social inference under uncertainty. Future pharmacoinaging studies are needed to investigate the relationship between aging-related changes in neuromodulation with changes in dynamic social inference learning.

The current findings also add to the emerging literature on adult development of social cognition. Accumulating evidence points to the fact that social cognitive function differs between younger and older adults (Bolenz, Reiter, & Eppinger, 2017; Henry et al., 2013; Reiter et al., 2017). A limitation of many of the previously applied paradigms might be that they did not explicitly test social inference evolving over the course of time. That is, these paradigms did not require participants to constantly update beliefs about others, but asked for “one-shot” ratings about others’ minds (compare e.g. (Baron-Cohen, Leslie, & Frith, 1985; Kanske, Bockler, Trautwein, & Singer, 2015; Reiter et al., 2017). Our findings thus extend the existing adult developmental literature towards inference in a *dynamically changing* social environment. Our results demonstrate that this dynamic nature is indeed dealt with differently by younger and older adults – we show higher, more precise beliefs about volatility and more quickly changing learning rates about the fidelity of the other in older adults. It appears thus important for future age comparative studies to more explicitly take this dynamic nature into account, given that natural social interactions are indeed characterized by the demand to continuously update our models about the social environment in order to simulate and predict others’ behavior accurately and to react to it appropriately (Siegel et al., 2018; Suzuki et al., 2012)

Besides age differences on the observer side, we also observed age differences as a function of the adviser’s age. Given that validity of the advice was pre-determined in a standardized manner, and thus held constant between older and younger advisers, these findings might be interpreted as stereotype effects. Our results showed that younger adults expected older adults to change their intentions more frequently but also held a belief about

higher fidelity in older advisers compared to younger advisers. Social psychology research has established such stereotype effects towards older people (“ageism”, e.g. (Cuddy & Fiske, 2002) before. Previous qualitative and cluster analyses showed that “doddering but dear”, i.e., incompetent but warm, types of stereotypes exist about elderly in the general population. It is interesting to speculate that our quantitative-computational findings (maybe most tentatively interpreted as “fickle but dear” elderly) might in part mirror these qualitative findings. While more research is needed to clearly map the different approaches onto each other, our study might serve as an example on how to use computational approaches for stereotype research (Siegel et al., 2018; Tamir & Thornton, 2018).

5. Conclusions

In summary, adopting a Bayesian modeling account, we show age differences between older and younger adults when dynamically updating beliefs about the intentions of younger vs. older counterparts. Older adults were characterized by an enhanced learning rate about the adviser’s fidelity due to their high confidence in their higher and more precise volatility estimates. We also observed computational evidence for the fact that older and younger advisers are perceived differently during social interactions: young adults assigned more volatility and more fidelity to older advisers than to younger advisers. In sum, the study offers insight into the behavioral and computational mechanisms that underlie age differences when dealing with changes in dynamic social interactions. Taken together, these findings showcase the utility of quantitative models to construct individual fingerprints in order to test socio-developmental hypotheses.

References

- Bäckman, L., Nyberg, L., Lindenberger, U., Li, S.-C., & Farde, L. (2006). The correlative triad among aging, dopamine, and cognition: current status and future prospects. *Neuroscience & Biobehavioral Reviews*, 30(6), 791-807.
- Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a "theory of mind"? *Cognition*, 21(1), 37-46.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, 68(3), 255-278.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*.
- Behrens, T. E., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. (2008). Associative learning of social value. *Nature*, 456(7219), 245-249. doi:nature07538 [pii]
- 10.1038/nature07538
- Bolenz, F., Reiter, A. M., & Eppinger, B. (2017). Developmental changes in learning: Computational mechanisms and social influences. *Frontiers in Psychology*, 8, 2048.
- Cuddy, A. J., & Fiske, S. T. (2002). Doddering but dear: Process, content, and function in stereotyping of older persons. *Ageism: Stereotyping and prejudice against older persons*, 3, 26.
- Daunizeau, J., Adam, V., & Rigoux, L. (2014). VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioural data. *PLoS computational biology*, 10(1), e1003441.
- Daunizeau, J., Den Ouden, H. E., Pessiglione, M., Kiebel, S. J., Stephan, K. E., & Friston, K. J. (2010). Observing the observer (I): meta-bayesian models of learning and decision-making. *PLoS One*, 5(12), e15554.
- Diaconescu, A. O., Mathys, C., Weber, L. A., Daunizeau, J., Kasper, L., Lomakina, E. I., . . . Stephan, K. E. (2014). Inferring on the intentions of others by hierarchical Bayesian learning.
- Diaconescu, A. O., Mathys, C., Weber, L. A., Kasper, L., Mauer, J., & Stephan, K. E. (2017). Hierarchical prediction errors in midbrain and septum during social learning. *Social cognitive and affective neuroscience*, 12(4), 618-634.
- Dreher, J.-C., Meyer-Lindenberg, A., Kohn, P., & Berman, K. F. (2008). Age-related changes in midbrain dopaminergic regulation of the human reward system. *Proceedings of the National Academy of Sciences*, 105(39), 15106-15111.
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends Cogn Sci*, 13(7), 293-301. doi:S1364-6613(09)00117-X [pii]
- 10.1016/j.tics.2009.04.005
- Garvert, M. M., Moutoussis, M., Kurth-Nelson, Z., Behrens, T. E., & Dolan, R. J. (2015). Learning-induced plasticity in medial prefrontal cortex predicts preference malleability. *Neuron*, 85(2), 418-428.
- Grothe, M., Heinsen, H., & Teipel, S. J. (2012). Atrophy of the cholinergic basal forebrain over the adult age range and in early stages of Alzheimer's disease. *Biological psychiatry*, 71(9), 805-813.
- Hämmerer, D., Schwartenbeck, P., Gallagher, M., FitzGerald, T., Düzel, E., & Dolan, R. J. (2018). Older adults fail to form stable task representations during model-based reversal inference. *Neurobiology of Aging*.
- Henry, J. D., Phillips, L. H., Ruffman, T., & Bailey, P. E. (2013). A meta-analytic review of age differences in theory of mind. *Psychology and aging*, 28(3), 826.
- Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E., & Stephan, K. E. (2013). Hierarchical Prediction Errors in Midbrain and Basal Forebrain during Sensory Learning. *Neuron*, 80(2), 519-530. doi:S0896-6273(13)00807-6 [pii]
- 10.1016/j.neuron.2013.09.009
- Kanske, P., Bockler, A., Trautwein, F. M., & Singer, T. (2015). Dissecting the social brain: Introducing the EmpaToM to reveal distinct neural networks and brain-behavior relations for empathy and Theory of Mind. *Neuroimage*, 122, 6-19. doi:10.1016/j.neuroimage.2015.07.082

- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2015). Package 'lmerTest'. *R package version*, 2(0).
- Lenth, R. V. (2016). Least-squares means: the R package lsmeans. *J Stat Softw*, 69(1), 1-33.
- Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A bayesian foundation for individual learning under uncertainty. *Front Hum Neurosci*, 5, 39. doi:10.3389/fnhum.2011.00039
- Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the Hierarchical Gaussian Filter. *Front Hum Neurosci*, 8, 825. doi:10.3389/fnhum.2014.00825
- Moran, R. J., Symmonds, M., Dolan, R. J., & Friston, K. J. (2014). The brain ages optimally to model its environment: evidence from sensory learning over the adult lifespan. *PLoS Comput Biol*, 10(1), e1003422. doi:10.1371/journal.pcbi.1003422
- Nasreddine, Z. S., Phillips, N. A., Bédirian, V., Charbonneau, S., Whitehead, V., Collin, I., . . . Chertkow, H. (2005). The Montreal Cognitive Assessment, MoCA: a brief screening tool for mild cognitive impairment. *Journal of the American Geriatrics Society*, 53(4), 695-699.
- Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature communications*, 7, 11609.
- Penny, W. D., Stephan, K. E., Daunizeau, J., Rosa, M. J., Friston, K. J., Schofield, T. M., & Leff, A. P. (2010). Comparing families of dynamic causal models. *PLoS Comput Biol*, 6(3), e1000709. doi:10.1371/journal.pcbi.1000709 [doi]
- Reiter, A. M., Kanske, P., Eppinger, B., & Li, S.-C. (2017). The Aging of the Social Mind-Differential Effects on Components of Social Understanding. *Scientific Reports*, 7(1), 11046.
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies—revisited. *NeuroImage*, 84, 971-985.
- Siegel, J. Z., Mathys, C., Rutledge, R. B., & Crockett, M. J. (2018). Beliefs about bad people are volatile. *Nature Human Behaviour*, 2(10), 750-756. doi:10.1038/s41562-018-0425-1
- Singmann, H., Bolker, B., Westfall, J., Aust, F., Højsgaard, S., Fox, J., . . . Mertens, U. (2016). afex: analysis of factorial experiments. R package version 0.16-1.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, 46(4), 1004-1017. doi:S1053-8119(09)00263-8 [pii] 10.1016/j.neuroimage.2009.03.025 [doi]
- Suzuki, S., Harasawa, N., Ueno, K., Gardner, J. L., Ichinohe, N., Haruno, M., . . . Nakahara, H. (2012). Learning to simulate others' decisions. *Neuron*, 74(6), 1125-1137.
- Tamir, D. I., & Thornton, M. A. (2018). Modeling the predictive social mind. *Trends in cognitive sciences*.

805

806 Supplementary Information

807 Figure S1: Parameter recovery for all free parameters of the model. To examine the recovery
 808 of the estimated parameters we simulated responses using the empirical maximum *a*
 809 *posteriori* parameter estimates from all participants' data (x-axis), and then re-estimated all
 810 parameters (y-axis) based on the simulated data. The correlation coefficients (with
 811 corresponding p values) as well as regression slopes are included to quantify the parameter
 812 recovery results.

813