Neurobiology of Aging 103 (2021) 98–108

Contents lists available at ScienceDirect

Neurobiology of Aging

journal homepage: www.elsevier.com/locate/neuaging.org

Human aging alters social inference about others' changing intentions

Andrea M.F. Reiter[a,b,c,d,*], Andreea O. Diaconescu[e,f,g], Ben Eppinger[a,h,i], Shu-Chen Li[a,j]

[a] Lifespan Developmental Neuroscience, Faculty of Psychology, Technische Universität Dresden, Germany
[b] Department of Neurology, Max-Planck-Institute for Human Cognitive and Brain Sciences, Leipzig, Germany
[c] Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London, UK
[d] Department of Child and Adolescent Psychiatry, Psychosomatics and Psychotherapy, Center of Mental Health, University of Würzburg, Würzburg, Germany
[e] Translational Neuromodeling Unit, University of Zurich & ETH Zurich, Switzerland
[f] Department of Psychiatry, University of Basel, Switzerland
[g] Krembil Centre for Neuroinformatics, Centre for Addiction and Mental Health (CAMH), University of Toronto, Canada
[h] Department of Psychology, Concordia University, Canada
[i] PERFORM Centre, Concordia University, Canada
[j] CeTI – Centre for Tactile Internet With Human-in-the-Loop, Technische Universität Dresden, Germany

ARTICLE INFO

Article history:
Received 4 February 2020
Revised 8 January 2021
Accepted 9 January 2021
Available online 20 February 2021

Keywords:
Theory of mind
Social cognition
Social learning
Aging
Lifespan development
Bayesian cognitive modeling

ABSTRACT

Decoding others' intentions accurately in order to adapt one's own behavior is pivotal throughout life. In this study, we asked how younger and older adults deal with uncertainty in dynamic social environments. We used an advice-taking paradigm together with Bayesian modeling to characterize effects 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 volatility of intentions. Older adults expected the adviser to change his/her intentions more frequently (i.e., a higher volatility of the adviser). They also showed higher confidence (i.e., precision) in their volatility beliefs and were less willing to change their beliefs about volatility over the course of the experiment. This led them to update their predictions about the fidelity of the adviser more quickly. Potentially indicative of stereotype effects, we observed that older advisers were perceived as more volatile, but also more faithful than younger advisers. This offers new insights into adult age differences in response to social uncertainty.

1. Introduction

Throughout our lives we often consider the advice of others when making important decisions. Which insurance should 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 the other's 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 advice. 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 easily conceivable that these age-related differences in dealing with uncertainty also play an important role when older adults navigate social environments. Previous studies investigating the effects of aging on the visual processing of social stimuli (Goh et al., 2010), on social evaluation (Cassidy et al, 2012), as well as on trust and trustworthiness (Bailey et al., 2012, 2015; Suzuki et al., 2019), have revealed differences between younger and older adults. Moreover, impairments regarding higher-order socio-cognitive skills like Theory of Mind or mentalizing have been demonstrated in older as compared to younger adults (Henry et al., 2013; Reiter et al., 2017). Studies regarding adult age differences on trust, however, have showed mixed results (Bailey et al., 2012, 2015; Bond Jr and DePaulo, 2008; Pak et al., 2017; Sutter and

* Corresponding author at: Lifespan Developmental Neuroscience, Faculty of Psychology, Technische Universität Dresden, Zellescher Weg 17, Germany Tel.:+49 (0)351 463 32065.
E-mail address: andrea.reiter@tu-dresden.de (A.M.F. Reiter).

0197-4580/$ – see front matter © 2021 Elsevier Inc. All rights reserved.
https://doi.org/10.1016/j.neurobiolaging.2021.01.034

Kocher, 2007; Suzuki et al., 2019, see Bailey and Leon (2019) for a very recent meta-analysis), which was attributed to the different methodological approaches used to study trust (e.g., self-report vs. behavior, financial vs. nonfinancial expressions of trust). According to recent meta-analytical evidence, a consistent finding across studies was that older adults were more trusting of the untrustworthy than younger adults (Bailey and Leon, 2019). However, notably, most of these studies did not use paradigms that required participants to learn about, and adapt to, changing social input over time.

To investigate the question of how older adults deal with a dynamically changing social environment, and the uncertainties associated with it, we build on recent advances in computational modeling of social interaction in younger adults (Behrens et al., 2008; Diaconescu et al., 2014). This research demonstrated that a biologically plausible Bayesian framework, the Hierarchical Gaussian Filter (Mathys et al., 2011, 2014), can be used to formally capture the process of inferring the time-varying states of others (Diaconescu et al., 2014). The authors used an advice taking paradigm to study how participants inferred, and learned about, the fidelity of another person over time. In the Bayesian framework, such social learning is captured at hierarchically organized levels. The highest level (third) represents the volatility of the social environment, that is, the rate at which intentions change. One level below (second level) represents the adviser's fidelity, that is, the tendency of the adviser to give accurate advice. The lowest level (first) represents the observation of whether the advice received was indeed accurate or not. Importantly, these levels are connected. Knowing whether to disregard an unexpected outcome (e.g., ignore 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 fidelity), depends on the *precision* of one's beliefs. When the precision of beliefs about the adviser's fidelity is low, the agent is more sensitive to newly observed outcomes. Thus, the agent will be more inclined to incorporate prediction errors – that is, surprising outcomes not in line with one's prior belief of the adviser – into future predictions about the adviser and adapt her behavior accordingly. Different neurotransmitter systems are suggested to be involved in signaling prediction errors (PEs) on different levels. Whereas lower-level PEs concerning the adviser's fidelity were associated with dopaminergic neurotransmission, higher level PEs about the adviser's volatility were found in the cholinergic septum (Diaconescu et al., 2017). This dissociation on the level of neuro-modulation has also been observed in nonsocial tasks (Iglesias et al., 2013). Importantly, both the dopaminergic and the cholinergic system undergo marked change during aging (Bäckman et al., 2006; Grothe et al., 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 2 main quantities: (1) their beliefs about the *probability* of receiving accurate advice and (2) 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 (1) learning about the fidelity of an adviser (i.e., the *probability* of receiving accurate advice) at the lower level, and (2) how this fidelity changes over time (i.e., the *volatility* of intentions) at the higher level. Based on findings from the non-social domain (Hämmerer et al., 2018; Nassar et al., 2016), we expected to find higher volatility estimates in older adults than younger adults. Furthermore, we were interested in whether such computations were altered as a function of the social identity of the adviser (own age vs. other age adviser).

2. Methods

2.1. Sample

The study sample consisted of a total of 82 participants (37 young adults [YA], age range 18–30 years and 46 older adults [OA], age range 65–80 years). Participants were excluded from analysis if 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 ($n = 2$ depression, $n = 1$ borderline personality disorder), neurological diagnosis ($n = 1$ stroke, $n = 1$ Parkinson's disease), reporting problems in discriminating stimulus colors during or after the experiment ($n = 2$). $N = 5$ participants had to be excluded due to technical or human errors during data acquisition. Thus, the final sample consisted of 33 OA (20 female, mean age = 71.97 years, SD = 3.95) and 34 YA (21 female, mean age = 24.35 years, SD = 3.11), a sample size which was similar to recent studies detecting age differences in decision-making under uncertainty (Hämmerer et al., 2018), or defining inter-individual differences on a related advice-taking paradigm (Henco et al., 2020; Sevgi et al., 2020). Participants provided written informed consent prior to study participation. They received 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). Participants played 120 trials of 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 2 sources of information: (1) a pie chart displaying the true probabilities of winning associated with the 2 colors (blue vs. green pie charts: 75:25, 65:35, 55:45, 45:55, 35:65, and 25:75), and (2) a piece of advice given by another person (the 'advisor'), which was presented via a video (Fig. 1). To provide advice the advisor held up a card in the color that he/she recommended to pick. The accuracy of the advice was randomized to the majority color of the pie chart; hence, the accuracy of the advice was independent of the pie chart probabilities.

The participants were instructed that the adviser had been playing the game in a session before and had been provided more insight by the experimenter than themselves concerning which color would win on each trial. They were instructed that because of a different incentive structure, the adviser's motivation to provide helpful advice might change over the course of the experiment – it could be that an adviser is misleading at certain points of the experiment, but helpful at other points. Participants were told that it is thus important to constantly monitor the adviser's intentions in giving accurate or false advice. They were instructed that their accurate predictions of the winning color would be financially incentivized. This was visualized by a progress bar which increased every time participants bet on the correct color, and decreased when participants missed a trial or predicted the wrong color. By reaching certain thresholds (visualized by silver and gold bars in the experiment) with their progress bar, they would gain extra money on top of their basic study reimbursement. The validity of the advice was predetermined based on the dominant strategy that previous real human advisers actually used in this task (Diaconescu et al., 2014). This input structure (i.e. the trial-by-trial sequence of congruency of the advice with the outcome) was held

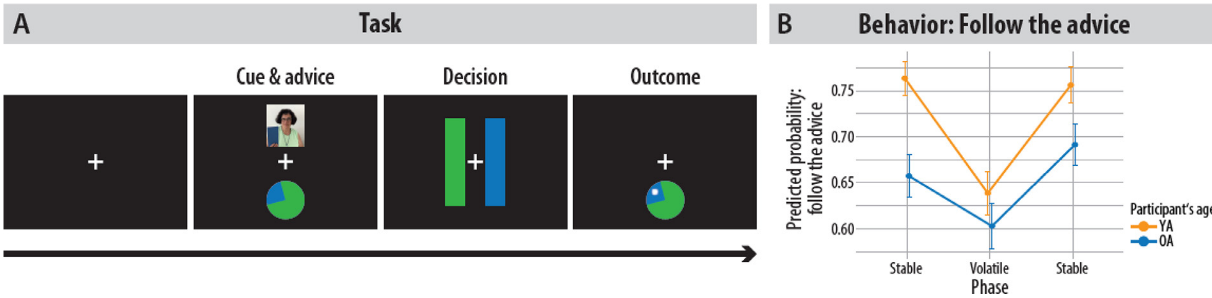

Fig. 1. Task and Raw Choice Data. (A) An example trial sequence. In a simple, roulette-like decision-making game, participants decided between green and blue on each trial, one of which was the winning color. Before making a choice, participants saw both (i) a pie chart indicating the true winning probability of the colors in the respective trial and (ii) the advice of another person, shown via video. After making a decision, participants received 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 advice following as a function of participants' age group and the phase of the experiment (stable, volatile, stable). Predicted values of a mixed model were plotted using the predict function in R.

constant across participants and was simulated based on an optimal Bayesian learner.

In accord with previous studies (Behrens et al., 2008; Diaconescu et al., 2014, 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 and once interacting with a younger adviser (the order of the 2 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 2 testing sessions were m = 6.83 ±1.35 days apart (range 3–14 days). Two slightly different versions of the input structure, created by exchanging the 2 stable phases, were used for the 2 testing days of the within-subject assessment (counter-balanced 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 cue, and less on the social cue, lending evidence to the notion that the social intentionality manipulation adopted here is indeed effective.

We used MATLAB for computational modeling, R for mixed models and SPSS for repeated-measures MANOVAs. Trial-by-trial decisions on whether to follow a piece of advice were subjected to a logistic mixed effects model (Bates et al., 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 within a certain phase). 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 et al., 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; 2017), several cognitive models were applied to investigate age differences in updating beliefs about other people's (time-varying, uncertain) intentions (compare Fig. 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 that used a very similar task as well as the "Bayesian Brain" theory to explain social inference behavior in young adults (Diaconescu et al., 2014). See this model validation study for more details regarding the class of models described in the following.

2.3.1. Learning models

The Hierarchical Gaussian Filter (HGF), a hierarchical Bayesian model of learning under uncertainty and volatility (Mathys et al., 2011, 2014), is one formal implementation of the Bayesian brain theory. According to the HGF, an agent uses observable 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 that characterize participants' learning about two types of environmental statistics—namely the probability and the volatility of the adviser giving correct advice (see Fig. 2 B for a schema of this hierarchy). The HGF learning model includes 3 free parameters (see Eq. 3 and Eq. 4 (Supplementary Material) for a definition, Fig. 2B for a graphical depiction, Table 1 for a conceptual description). κ, ω, and ϑ, determine how these states evolve as a function of the inputs the participant receives (pie chart, advice, trial outcome). The parameter κ represents the coupling between the second and the third level, or how much 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.

Here, we consider 2 versions of the HGF: the full, 3-level version (see Supplementary Methods) as well as a reduced, 2-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, val-

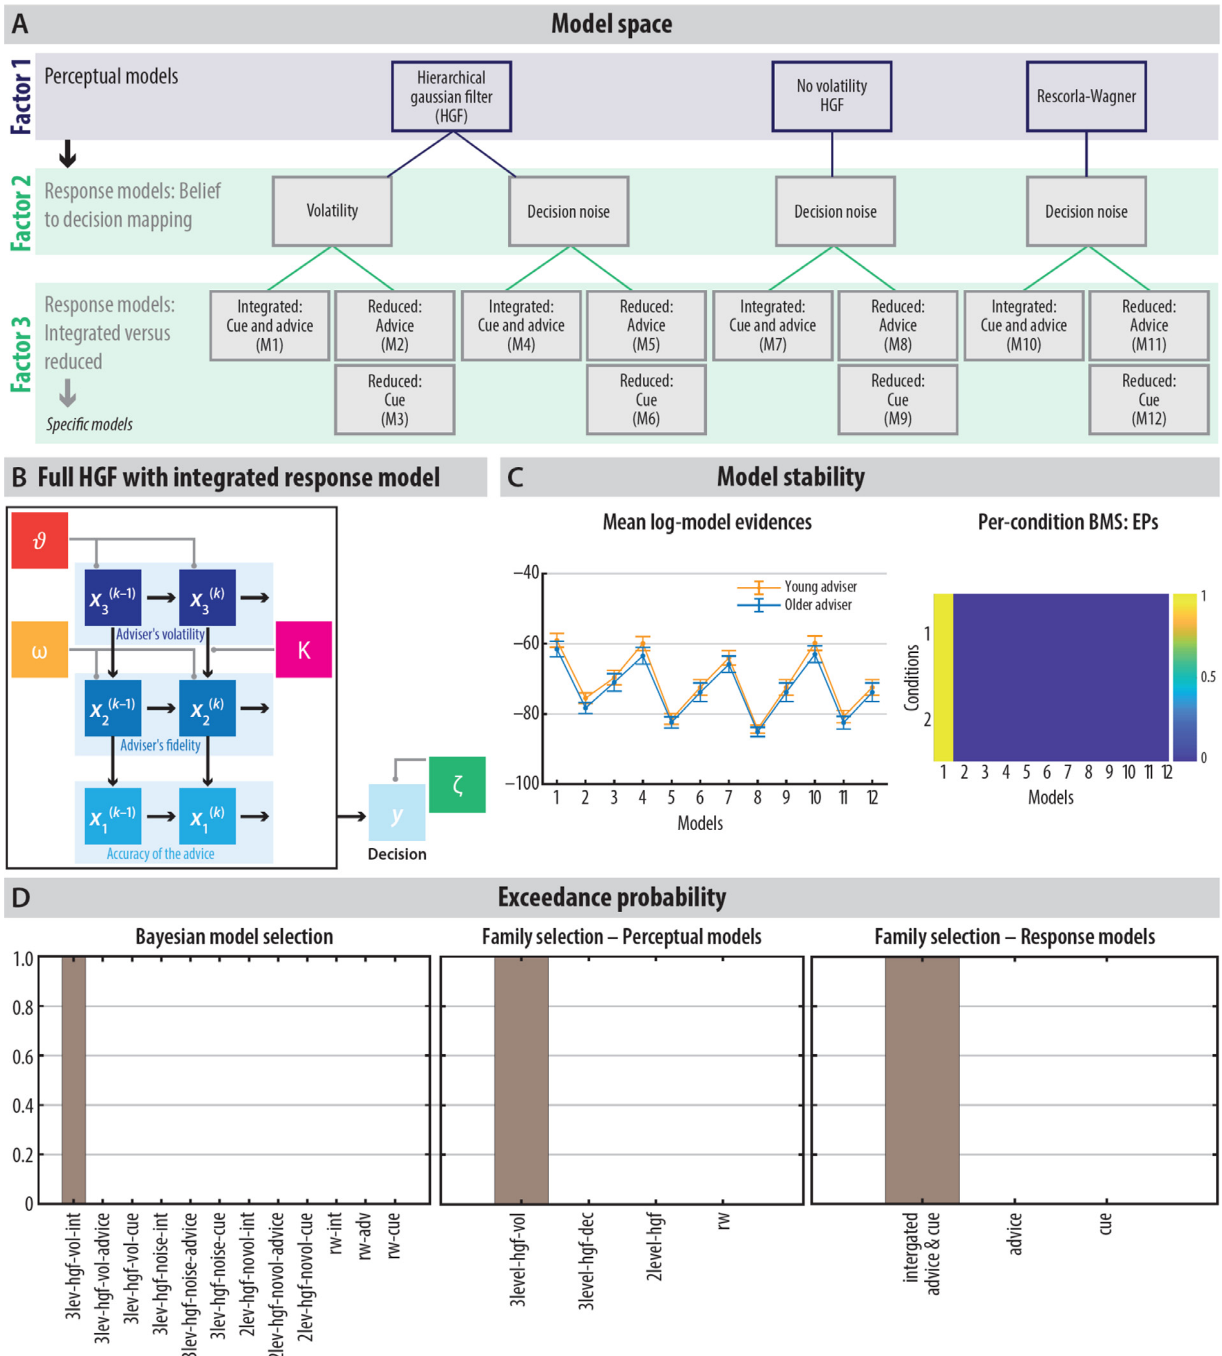

Fig. 2. Model space and model selection. (A) Model Space. The model space was set up based on a 3 × 2 × 2 structure. Factor 1: The first factor represents the perceptual/input model families (3-level HGF, reduced 2-level HGF, RW). Factor 2: In the case of the HGF, 2 response models were introduced. Decision 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 nonsocial (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 (x1…x3). x1 is the binary sensory input indicating the accuracy (true/false) of the current piece of advice, x2 is the current probability that the advice is correct (the fidelity of the advice) and x3 is the current volatility of the adviser's intentions. There are five (ϑ, ω, κ, ζ, β) free parameters of the model. The parameter κ determines how strongly x2 is influenced by x3, 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 nonsocial (pie chart) information. Decisions y are influenced by the parameter ζ, which determines the weight of the social cue relative to the nonsocial cue, and by β, which captures choice stochasticity. (C) Model stability across conditions. As each participant completed the task twice (within-subject design), interacting with an OA adviser and a 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. an OA adviser. See Panel A for (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.

Table 1
Free parameters of the computational models with a conceptual description of how they influence learning about the adviser

Parameter	Conceptual description
HGF	
mu2_0	A priori fidelity of the adviser
sigma2_0	A priori variance of the adviser's fidelity
mu3_0	A priori volatility of the adviser's volatility
sigma3_0	A priori variance of the adviser
kappa	Coupling between second and third level: How much does the agent use volatility when updating fidelity?
omega	Degree to which the agent updates beliefs about fidelity, independent of volatility
theta	Meta-volatility: How fast is volatility updated?
zeta	Weighing Parameter of social and non-social information: How much does the agent rely on social information?
beta	Decision Noise: How tightly are choices mapped onto beliefs?
RW	
v_0	A priori value of the advice
alpha	Learning rate (constant over trials): impact of prediction errors on the value of the advice

Table 2
Prior means and variances for estimating HGF and RW model parameters

	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

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

ues of choice options are updated using a prediction error that is weighted by a learning rate. See Supplementary Methods for the algorithmic implementation of the models.

2.3.2. Response models

Additionally, and in accord with previous work (Diaconescu et al., 2014; 2017), we consider different versions of response models in order to map beliefs to decisions. First, to test whether participants rely only on (1) the pie chart information, (2) the adviser's information, or (3) integrate both sources of information, we introduce a social weighing factor ζ that 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 Supplement, Eq. 7). Fitting ζ as a free parameter corresponds to the assumption that an agent integrates both sources (social and nonsocial) of information, and that individuals differ in the degree to which they weigh either source.

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

2.3.3. 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 (Mathys et al., 2011; 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 (ϑ, ω, κ, ζ, β), 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) and Daunizeau et al. (2010).

2.3.4. 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 of the log model evidence (Daunizeau et al., 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 employed family selection across the above-described learning as well as response model variations (see Fig. 2A and B), 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).

2.3.5. Parameter Recovery

We ran 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-Fig. 1).

2.3.6. Analyzing Age Effects on Model Parameters

All free parameters of the winning model were extracted and analyzed as a function of (1) participant's age group, (2) adviser's age group, (3) and their interaction, 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 (1) the second level of the HGF, that is, states reflecting the evolution of beliefs about the adviser's fidelity, and (2) the third level of the HGF, that is, 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 prevolatile; volatile; stable postvolatile), 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 as implemented in the R packages *afex* (Singmann et al., 2016) and *lmerTest* (Kuznetsova et al., 2015). Post-hoc tests were performed using the package *lsmeans* (Lenth, 2016).

2.3.7. Code Accessibility

MATLAB code to fit the Hierarchical Guassian 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/.

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 (χ^2 > 63.97, all *ps* < 0.001, see Fig. 1) on participants' trial-by-trial decisions of whether to follow a piece of advice or not. This indicates that the social volatility manipulation alters the participants' tendency of whether or not to follow the adviser. Fig. 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 ($z > 3.28$, $p < 0.001$). 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 ($\chi^2 = 6.18$, $p = 0.013$) and an interaction effect of participants' own age group with phase ($\chi^2 = 12.60$, $p = 0.002$, Fig. 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 ($z = 3.57$, $p < 0.001$), and the post-volatile phase ($z = 2.16$, $p = 0.03$) such that OA followed the adviser less than YA in these 2 phases. No significant difference was observed in the middle volatile phase ($z = 1.05$, $p = 0.291$) (compare Fig. 1b).

3.2. Model-Based Analysis – Model Selection

First, we tested for model stability across the 2 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, Fig. 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 evidence and tested for age group differences regarding which model explained the data best. BMS revealed clear evidence that both age groups were characterized best by the same model (posterior probability that the groups have the same model frequencies = 0.99). The winning model for the groups together was the full HGF with volatility-influenced decision noise including the social weighing parameter (M1, XP = 1, PXP = 1, see Fig. 2c and d). Family selection (i.e., pooling across models that share a common feature) across all perceptual model classes revealed that within the perceptual model family, the full HGF with volatility-influenced decision noise had highest model evidence (XP = 0.99, Fig. 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, Fig. 2d). This indicates that participants integrate social and nonsocial information instead of relying exclusively on one of the 2 sources. When running model selection for each group separately, we obtained the same results. These results replicate previous studies in YA (Diaconescu et al., 2014, 2017). BMS also answers questions about age differences regarding the strategy used to solve the task by quantifying evidence for both the null and alternative hypotheses. More specifically, BMS results demonstrate that both age groups infer 2 quantities at the same time: the fidelity of the adviser 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 nonsocial information to come to their decisions.

3.3. 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, compare Fig. 2) 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 ($F = 6.40$, $p < 0.001$). Post-hoc tests indicated that this difference was driven by group differences on the initial values of μ_2 ($F = 23.83$, $p < 0.001$, Fig. 3) and μ_3 ($F = 6.90$, $p = 0.01$, Fig. 3). Thus, OA 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 (initial μ_2 significantly positive in both adviser conditions, $t \geq 3.73$, *ps* < 0.001), whereas OA seem to hold a more neutral prior regarding the adviser's fidelity (initial μ_2 not significantly different from zero in both adviser conditions, $t \leq 1.71$, *ps* > 0.09). OA also had higher *a priori* beliefs about the volatility of the adviser (initial μ_3, Fig. 3). Interestingly, we further found a significant age group difference in ϑ, the-meta volatility parameter that governs the speed of volatility updating ($F = 6.69$, $p = 0.012$). OA had lower values on the meta-volatility parameter (see Fig. 3), suggesting that they update volatility more slowly than YA. None of the other learning and response parameters differed significantly between age groups ($F \leq 1.92$, $p > 0.171$).

3.4. 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 a 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 the 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.

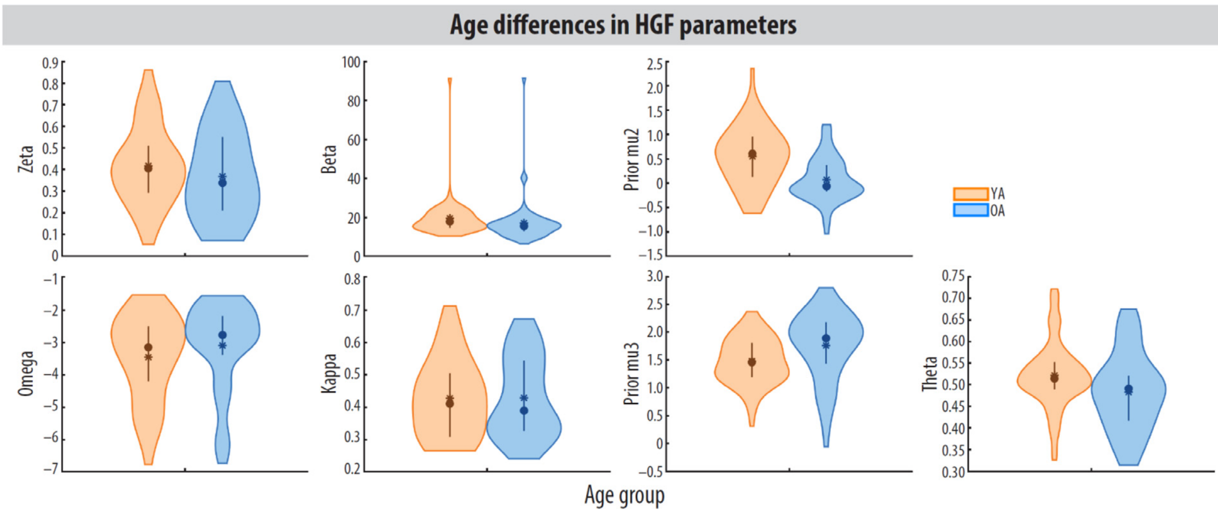

Fig. 3. Parameter comparisons. Parameter comparisons between age groups showed significant effects of age group on the priors of μ_2 and μ_3. Older adults had lower a priori beliefs about the fidelity of the adviser (μ_2) and higher a priori beliefs about the volatility of the adviser (μ_3). Additionally, the meta-volatility parameter (ϑ) 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.

3.4.1. 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 phase and trial number within phase, as well as their interaction (all $\chi^2 > 121.64$, all $p < 0.001$). 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 ($\chi^2 = 6.17$, $p = 0.013$), 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 stronger *a priori* 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 Fig. 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 phase and trial number within phase, as well as their interaction, (all $\chi^2 > 292.34$, all $p < 0.001$). These findings suggest certainty about one's beliefs vary in phases of different volatility, throughout the course of the experiment. Most interestingly, however, we also observed a main effect of age group ($\chi^2 = 5.31$, $p = 0.021$), which was due to OA displaying *higher* precision in their beliefs about volatility than YA (Fig. 4a and c). Moreover, a significant interaction of trial number within phase and age group was observed ($\chi^2 = 5.84$, $p = 0.016$, Fig. 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, $p = 0.016$), showing preserved modifiability of precision estimates within phase, as a function of changing volatility in the environment.

3.5. Beliefs about the adviser's fidelity

In a Bayesian framework, as formalized in the HGF, how we estimate volatility impacts learning on the lower level (in our case: learning about the fidelity of the adviser, see Fig. 2 B and methods). If the world is perceived as more volatile, 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 versus 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 $\chi^2 > 166.22$, all $p < 0.001$). Regarding age group differences, in line with our raw choice data analysis, we observed a significant interaction of age group and phase ($\chi^2 = 6.81$, $p = 0.033$). 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, $p = 0.022$). We also observed a significant interaction effect of age group and trial number ($\chi^2 = 7.39$, $p = 0.006$). 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, $p = 0.005$). 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 ($\chi^2 = 3.98$, $p = 0.046$, Fig. 4c). Additionally, we observed a trial number by age group interaction ($\chi^2 = 5.36$, $p = 0.021$). Interestingly, the learning rate co-varied more with trial number in OA (*z ratio* = 2.36, $p = 0.02$) as compared to YA. The latter was also moderated by phase ($\chi^2 = 6.22$, $p = 0.045$). As can be seen

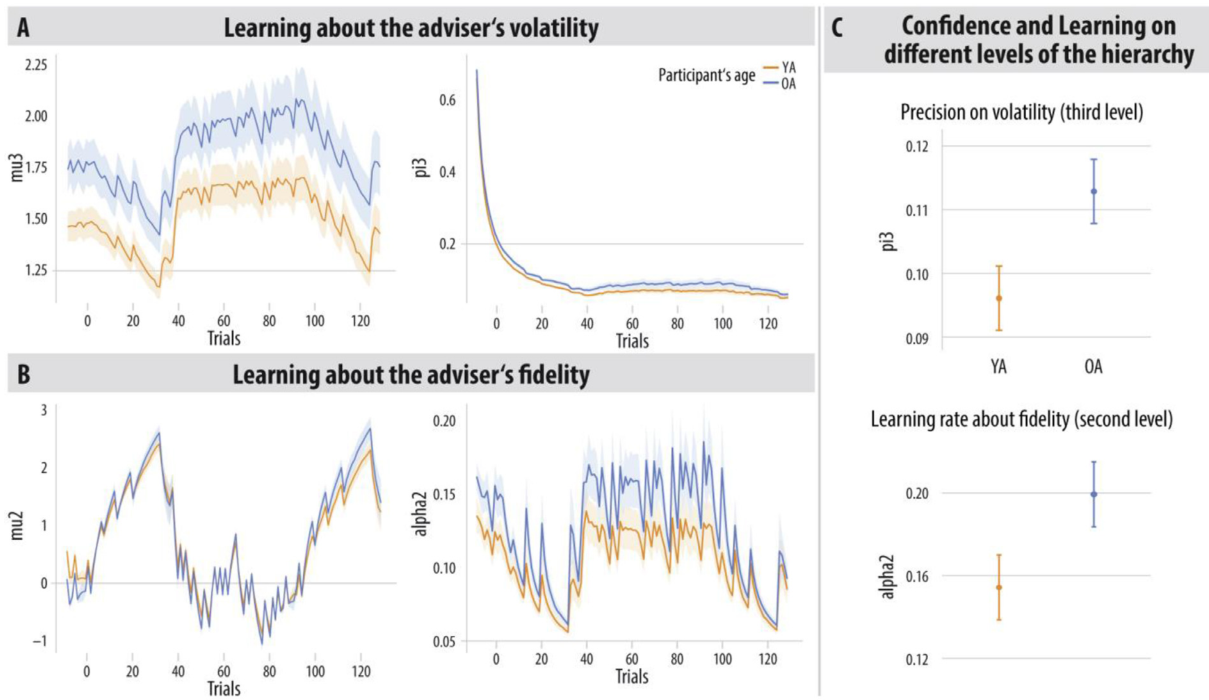

Fig. 4. Age differences in the evolution of beliefs, precision, and learning rate on different levels of the hierarchy. (A) Left: Evolution of volatility estimates (i.e., μ_3 of the HGF) through the course of the experiment as a function of participants' own age group. Right: Evolution of precision about volatility. (B) Left: Evolution of fidelity estimates (i.e., 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 of belief-precision concerning volatility (third level) and the learning rate on fidelity (second level).

from Fig. 4b, learning rate fluctuations differed most strongly between groups in the first, and the second, unstable phase. There were no age-related significant differences on precision on this level.

In sum, OA started the experiment with lower expectations about the adviser's fidelity. However, due to their pronounced and stable beliefs about intentions being volatile (i.e., higher μ_3 estimates throughout, see Fig. 4a), along with a reduced tendency to update them (i.e., reduced meta volatility parameter theta), OA were also well prepared to adjust their fidelity beliefs over the course of the experiment (i.e., higher learning rates regarding the fidelity, $\mu_{2,}$ see Fig. 4b).

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

In the 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 (1) the adviser's fidelity, (2) 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, see methods). 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 the participants' own age group and the adviser's age group, moderated by phase ($\chi^2 = 7.11$, $p = 0.029$). Post-hoc tests showed that the significant adviser age differences in volatility we described above were only present when comparing the adviser condition in YAs ($\chi^2 = 2.24$, $p = 0.03$), but not when comparing the adviser conditions in OA O OA (z $ratio = 0.32$, $p = 0.75$). This was due to YA displaying higher volatility estimates when interacting with OA as compared to YA, particularly in the middle volatile (z $ratio = 2.61$, $p = 0.009$) and the last phase ($z = 2.18$, $p = 0.029$).

The model also revealed significant interactions of the adviser's age group and trial ($\chi^2 = 6.30$, $p = 0.012$) as well as of one's own age group, adviser's age group, phase, and trial ($\chi^2 = 6.74$, $p = 0.035$). Fig. 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 (Fig. 5B). We again observed an interaction of one's own age group, adviser's age group and phase on beliefs about fidelity of the adviser ($\mu_{2,}$ $\chi^2 = 6.49$, $p = 0.039$). 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 $ratio = 3.00$, $p = 0.003$, see Fig. 5). Contrary to this effect in YA, no significant difference in $\mu_{2,}$ as a function of adviser's age, was found in OA.

4. Discussion

In this study, we employed a hierarchical Bayesian modeling framework to describe social learning about others' intentions in younger and older adults. We found that both our own age as well as our 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 mak-

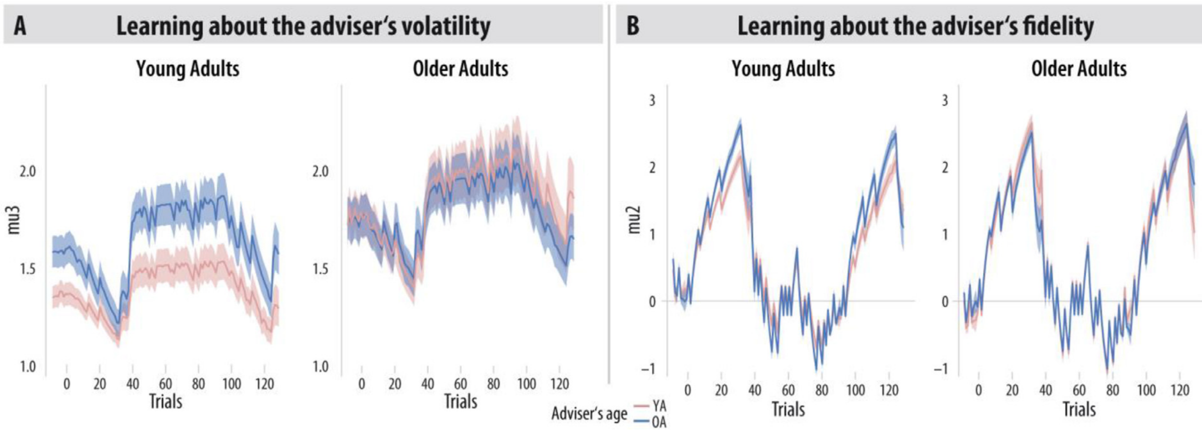

Fig. 5. Differences in belief evolution as a function of the adviser's age group. (A) Beliefs about the adviser's volatility as a function of participants' age, advisers' age, and phase. (B) YA assigned significantly more fidelity to OA advisers in the first stable phase than to YA advisers. No significant adviser's age effect on the belief about fidelity was found in OA.

ing predictions about the adviser's fidelity (i.e., their probability of offering accurate recommendations). We found that, before learning about the adviser, older adults had an *a priori* expectation that the adviser was more volatile, and thus changed their intentions more during 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 values (ϑ). 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/a mishap) 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. Rather, because they expect the adviser to change, they 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 fidelity 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 that 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 higher-level belief belief updating and the maintenance of a "more robust model of the world" in healthy aging (Moran et al., 2014). Similar to nonsocial contexts, we find that higher-level beliefs about intentions (i.e., on the third level of the HGF) 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 with 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 our study, which demonstrated that impressions about morally bad agents are more volatile and more rapidly updated than impressions about good agents (Siegel et al., 2018). Older adults in our study had a less positive (albeit more realistic) prior view of the adviser than younger adults and were more inclined to update this view than younger adults. This differs from previous findings, where older adults were shown to trust *more* in untrustworthy social information than younger adults. Meta-analytical evidence suggests that adult age difference with respect to trust are generally subject to variability as a function of the experimental set-up that is used in a specific study. This also might explain the discrepant results of the present study, as other than in previous research, in our experiment, participants were required to constantly update the trustworthiness of the adviser (Bailey and Leon, 2019). The current findings thus 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 et al., 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 as it evolved over the course of the experiment. 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 et al., 1985; Kanske et al., 2015; Reiter et al., 2017). Our findings thus extend the existing adult developmental literature toward 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. Thus, it appears important that future age comparative studies more explicitly take this dynamic nature into account. 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).

Parallel to age differences in nonsocial 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 (nonsocial) 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 develop differently (Reiter et al., 2017).

The present study was built on a nondeceptive previous study (Diaconescu et al., 2014). That previous study used true interactions between the player and a human adviser (Diaconescu et al., 2014), whereas the current study used pre-recorded videos of the advisers. The standardized advice structure we used in this study was pre-determined in a way that mimicked the behavior of the real human advisers in Diaconescu et al. Nevertheless, the videos used in the present study might have reduced the social aspect of the task and thus the ecological validity. Future studies could include a direct non-social comparison task (such as e.g., Diaconescu et al., 2014; Garvert et al., 2015) to investigate the specificity of the social effects.

Another recent study (Diaconescu et al., 2017) has used the same Bayesian modeling approach as presently to unravel neural correlates of social learning. The authors showed that learning at different levels of the hierarchy mapped onto distinct neuro-modulatory systems. Learning about the volatility of the others' intentions (as reduced in older adults in the current study) was associated with activation in the cholinergic basal forebrain. Learning about the adviser's fidelity (pronounced in older adults here) activated the dopaminergic midbrain. These differences between younger and older adults may be linked to aging-related changes in neuromodulation. Of note, both the cholinergic and the dopaminergic systems undergo marked change over the course of the lifespan (Bäckman et al., 2006; Dreher et al., 2008; Grothe et al., 2012; Schliebs, 2006 #5084). Evidence from earlier computational studies of aging neuromodulation (Li et al., 2001; Li et al., 2006) suggest that older adult's deficient neuromodulation (particularly dopaminergic modulation) would result in noisier neural information processing. More specifically, it was suggested that the computational consequences would be reduced precision of neurocognitive representations of learning experiences and increased 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 pharmacoimaging studies are needed to investigate the relationship between aging-related changes in neuromodulation and changes in dynamic social inference learning.

Besides age differences on the side of the observer, 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 before ("ageism," e.g., Cuddy and Fiske, 2002). Previous qualitative and cluster analyses showed that "doddering but dear", that is, inconsistent but warm, types of stereotypes exist about older adults in the general population. It is interesting to speculate that our quantitative-computational findings (maybe most tentatively interpreted as "fickle but dear" older adults) might mirror these qualitative findings in part. While more research is needed to clearly map the different approaches onto each other, our study might serve as an example of how to use computational approaches for stereotype research (Siegel et al., 2018; Tamir and 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 and 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.

Acknowledgments

The authors thank Abigail Calder and Anne 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 and to Joshua Grant for proof-reading. 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). AMFR is supported by a 2020 NARSAD Young Investigator Grant from the Brain & Behavior Research Foundation. AOD acknowledges support by the Swiss National Foundation (SNF PZ00P3_167952) and the Krembil Foundation.

Disclosure statement

None.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.neurobiolaging.2021.01.034.

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. Neurosci Biobehav Rev 30 (6), 791–807.
Bailey, P.E., Leon, T., 2019. A systematic review and meta-analysis of age-related differences in trust. Psychol Aging 34 (5), 674.
Bailey, P.E., Ruffman, T., Rendell, P.G., 2012. Age-related differences in social economic decision making: the ultimatum game. J Gerontol Ser B 68 (3), 356–363.
Bailey, P.E., Slessor, G., Rieger, M., Rendell, P.G., Moustafa, A.A., Ruffman, T., 2015. Trust and trustworthiness in young and older adults. Psychol Aging 30 (4), 977.
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. J Mem Lang 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. Front Psychol 8, 2048.
Bond Jr, C. F., & DePaulo, B. M. (2008). Individual differences in judging deception: accuracy and bias. Psychol Bull, 134(4), 477.
Cassidy, B.S., Shih, J.Y., Gutchess, A.H., 2012. Age-related changes to the neural correlates of social evaluation. Social Neurosci 7 (6), 552–564.
Cuddy, A.J., Fiske, S.T., 2002. Doddering but dear: process, content, and function in stereotyping of older persons. Ageism 3, 26.

Daunizeau, J., Adam, V., Rigoux, L., 2014. VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioural data. PLoS Comput Biol 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. PLoS Comput Biol 10 (9), e1003810. doi:10.1371/journal.pcbi.1003810.
Diaconescu, A.O., Mathys, C., Weber, L.A.E., Kasper, L., Mauer, J., Stephan, K.E, 2017. Hierarchical prediction errors in midbrain and septum during social learning. Soc Cogn Affect Neurosci 12 (4), 618–634. doi:10.1093/scan/nsw171.
Dreher, J.-C., Meyer-Lindenberg, A., Kohn, P., Berman, K.F., 2008. Age-related changes in midbrain dopaminergic regulation of the human reward system. Proc Natl Acad Sci 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: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.
Goh, J.O., Suzuki, A., Park, D.C., 2010. Reduced neural selectivity increases fMRI adaptation with age during face discrimination. Neuroimage 51 (1), 336–344.
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. Biol 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. Neurobiol Aging.
Henco, L., Diaconescu, A.O., Lahnakoski, J.M., Brandi, M.-L., Hörmann, S., Hennings, J., ..., Bolis, D., 2020. Aberrant computational mechanisms of social learning and decision-making in schizophrenia and borderline personality disorder. PLoS Comput Biol 16 (9), e1008162.
Henry, J.D., Phillips, L.H., Ruffman, T., Bailey, P.E., 2013. A meta-analytic review of age differences in theory of mind. Psychol 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: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., 2017. lmerTest package: tests in linear mixed effects models. J Stat Softw 82 (13), 1–26.
Lenth, R.V., 2016. Least-squares means: the R package lsmeans. J Stat Softw 69 (1), 1–33.
MacDonald, S.W., Karlsson, S., Rieckmann, A., Nyberg, L., Bäckman, L., 2012. Aging-related increases in behavioral variability: relations to losses of dopamine D1 receptors. Journal of Neuroscience 32 (24), 8186–8191.
Mathys, C.D., 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. J Am Geriatr Soc 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. Nat Commun 7, 11609.
Pak, R., Rovira, E., McLaughlin, A.C., Baldwin, N., 2017. Does the domain of technology impact user trust? Investigating trust in automation across different consumer-oriented domains in young adults, military, and older adults. Theor Issues Ergon Sci 18 (3), 199–220.
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[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. Sci Rep 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.
Li, S.C., Lindenberger, U., Sikström, S., 2001. Aging cognition: from neuromodulation to representation. Trends in Cognitive Sciences 5 (11), 479–486.
Li, S.C., von Oertzen, T., Lindenberger, U., 2006. A neurocomputational model of stochastic resonance and aging. Neurocomputing 69 (13–15), 1553–1560.
Schliebs, R., Arendt, T., 2006. The significance of the cholinergic system in the brain during aging and in Alzheimer's disease. Journal of neural transmission 113 (11), 1625–1644.
Sevgi, M., Diaconescu, A.O., Henco, L., Tittgemeyer, M., Schilbach, L., 2020. Social Bayes: using Bayesian modeling to study autistic trait–related differences in social cognition. Biol Psychiatry 87 (2), 185–193.
Siegel, J.Z., Mathys, C., Rutledge, R.B., Crockett, M.J., 2018. Beliefs about bad people are volatile. Nature Hum Behav 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 https://CRAN.R-project.org/package=afex.
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:10.1016/j.neuroimage.2009.03.025.
Sutter, M., Kocher, M.G., 2007. Trust and trustworthiness across different age groups. Games Econ Behav 59 (2), 364–382.
Suzuki, A., Ueno, M., Ishikawa, K., Kobayashi, A., Okubo, M., Nakai, T., 2019. Age-related differences in the activation of the mentalizing-and reward-related brain regions during the learning of others' true trustworthiness. Neurobiol Aging 73, 1–8.
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 Cogn Sci 22 (3), 201–212.