How do humans learn about other people?

Incorporating social knowledge structures into computational models

Koen M.M. Frolichs^{1,2}, Gabriela Rosenblau³, and Christoph W. Korn^{1,2}

¹ Institute for Systems Neuroscience, University Medical Center Hamburg-Eppendorf, Hamburg, Germany.

² Section Social Neuroscience, Department of General Psychiatry, University of Heidelberg, Heidelberg, Germany.

³Department of Psychological and Brain Sciences, George Washington University, Washington DC, USA.

Correspondence

Koen M.M. Frolichs and Christoph W. Korn

Section Social Neuroscience, Department of General Psychiatry, University of Heidelberg
Voßstraße 4

69115 Heidelberg, Germany

Koen.Frolichs@med.uni-heidelberg.de

<https://orcid.org/0000-0003-3298-1192>

Christoph.Korn@med.uni-heidelberg.de

<http://orcid.org/0000-0002-8228-2624>

Abstract

To navigate social interactions successfully, humans need to continuously learn about the personality traits of other people (e.g., how helpful or aggressive is the other person?). However, formal models that capture the complexities of social learning processes are currently lacking. In this study, we specified and tested potential strategies that humans could employ for learning about others. Standard Rescorla-Wagner (RW) learning models only capture part of the learning process because they neglect inherent knowledge structures and omit previously acquired knowledge. We therefore formalized two social knowledge structures and implemented them in novel hybrid RW models to test their usefulness across multiple social learning tasks. We named these concepts granularity (knowledge structures about personality traits that can be utilized at different levels of detail during learning) and reference points (previous knowledge formalized into representations of average people within a social group). In five behavioural experiments, results from model comparison and statistical analyses indicated that participants efficiently combined the concepts of granularity and reference points—with the specific combinations in models depending on the people and traits that participants learned about. Overall, our experiments demonstrate that variants of RW algorithms, which incorporate social knowledge structures, describe crucial aspects of the dynamics at play when people interact with each other.

Keywords:

Personality traits, social learning, Rescorla-Wagner, Big-Five, schemata

Introduction

Humans constantly meet new people and interact with them. Successfully navigating social interactions crucially depends on quick and accurate learning about others' personality from brief encounters. Humans are skilled in amassing social information by abstracting from concrete, situation-specific observations to personality traits. However, so far these learning processes have not been described in terms of underlying computational mechanisms. Here, we set out to specify how social knowledge structures shape learning about others.

Rescorla-Wagner (RW) models, that in the cognitive sciences fall under the wide umbrella term of reinforcement learning (RL) (Collins & Shenhav, 2022; Lockwood & Klein-Flügge, 2021; Chierchia et al., 2021), entail simple and robust algorithms that characterize dynamic learning processes across a wide range of (non-)social tasks (Sutton & Barto, 1998; Dayan & Niv, 2008; Dunne & O'Doherty, 2013; Ruff & Fehr, 2014; Niv & Langdon, 2016; Lockwood & Klein-Flügge, 2019; Zhang et al., 2019; Joiner et al., 2017; Palminteri et al., 2015). In general, RL models have been fundamental for understanding learning about the value of objects or situations regarding how rewarding these are (in the form of food, money, pain, etc.). RW models in specific describe learning as stepwise reductions of prediction errors (PEs) through outcomes or feedback, where the PE is the difference between a predicted value and the actually received outcome or feedback in a given time step. Learning results from using these PEs to update the estimates of future values.

Akin to non-social learning, social learning can be described by such RL and RW models (Ruff & Fehr, 2014; Olsson et al., 2020). Variants of these models capture how humans learn *from* others (i.e., observational learning: e.g., Rodriguez Buritica et al., 2015; Diaconescu et al., 2017; Hill et al., 2016; Najar et al., 2020) and *for* others (Garvert et al., 2015). Moreover, studies have shown that they can also account for how humans learn *about* others on specific characteristics such as trustworthiness, generosity or emotional states (Jones et al., 2011; Ruff and Fehr, 2014; Fareri et al., 2015; Zaki et al., 2016; Hackel et al., 2015; Stanley, 2016; Bellucci et al., 2019).

However, these models likely fail to explain the complexities of social learning about others' multi-faceted personality because they only rely on environmental feedback and are typically restricted to learning about the current situation or dimension at hand. Human personality cannot be accurately represented on a one- or low-dimensional continuum (such as nutritional or monetary value).

Research in personality psychology presupposes a multi-dimensional structure to human personality traits (Ostendorf, 1990; McCrae & Costa, 2008). A substantial literature in this field has revealed that the vast number of commonly used personality traits can be reduced to a few independent dimension or factors (Goldberg, 1990). The Big-Five personality factors (neuroticism, extraversion, openness to experience, agreeableness, and conscientiousness) are the most commonly discussed factors to describe various samples of people across the world (McCrae & Costa, 1999; Wiggins & Trapnell, 1997). According to influential theories, such factors can be represented in human cognition via schemata that bundle frequently co-occurring features (Cantor & Mischel, 1977; Mayer & Bower, 1986). As a consequence learning can be influenced by schemata with preconceived personality prototypes (i.e., stereotypes) (Mayer & Bower, 1986) and the group the other belongs to (in- versus out-group) (Kang et al., 2021). In some cases, people

might rely more on concrete prototypes and trait-exemplifying behaviours when judging others' personality. In situations people are experienced with, they can abstract from concrete exemplars to more general personality traits and to more fine-grained relations between them (Klein et al., 1992).

In order to capture the complexities of social learning about others' personality, computational models should implement such prototypes and abstractions to allow for multi-dimensional learning (Jolly & Chang, 2019; Lockwood et al., 2020). So far, studies that explore learning about others have not investigated personality traits specifically or have looked at single traits (Hartley & Sommerville, 2012). Focusing on learning about a single aspect of personality such as trustworthiness (Delgado et al., 2005; King-Casas et al., 2005) or on learning about specific groups of people (Lau et al., 2020; Amodio, 2008) makes learning processes amenable to standard RW models. But models on single aspects do not take the multi-dimensionality of human personality into consideration and thus cannot appropriately describe learning about personality structures in real life.

In this study, we aimed at testing how humans learn about the multi-dimensional personality of others. In order to account for the complexities of learning about human personality, we employed hybrid learning models that weigh prior experience, contextual knowledge, and RW (Camerer & Ho, 1999). We therefore constructed and tested a number of computational models with varying complexities, from a simple linear regression that functions as a baseline to hybrid models that combine standard RW learning models (Rescorla & Wagner, 1972) with two social knowledge structures that we refer to as *Reference Points* (RPs) and *Granularity* (G).

Reference points capture prior experience and expectations that humans have distilled about many other people (from in- or out-groups). Using these RPs, humans can compare newly encountered people with "average people." In this study, RPs thus function as points of comparison.

Granularity refers to the level of detail or abstraction with which humans represent the structure of personality traits. As soon as humans learn some new characteristic of a person, their whole view of that person might change (for better or worse), which implies that they update their expectations across many traits based on new information about a single characteristic. That is, receiving information on a single trait may be generalized to similar traits. In order to do so, humans may use representations (or schema) of the underlying similarities between personality traits (e.g., if you learn that someone is very polite you will also expect her to be friendly because these two traits are related). Alternatively, when information is sparse or no detail is needed, one can generalize individual traits into coarse representations that distil the most important parts of personality (i.e., akin to the factors of the Big-5).

We expect that the use of both reference points and granularity underlies social learning about others. Different learning situations might require information to be represented on a coarse or fine-grained level of detail. Using variations of social learning experiments and computational models, we tested how well both reference points and granularity accounted for learning behaviour. Models were evaluated against participants' behavioural data to illuminate what strategies participants might use. Additionally, we used model simulations to quantify the different learning

strategies based on their performance (when solving the same learning task as participants). This allowed us to evaluate if participants behaved in accordance with these best performing models given the set of our models. Crucially, several statistical analyses (e.g., regressions & correlations) were conducted to support the model-based analyses.

In recent work, we showed how adults and adolescents learn about others' preferences (Rosenblau et al., 2020; Rosenblau, Korn & Pelphrey, 2018) by trading off reference points (in the form of prior knowledge about the average population's preferences) and different levels of granularity (in the form of updating information based on past feedback about preferences for similar items). This type of updating requires social knowledge of the similarities between preferences.

The current study builds on our previous studies by applying a considerably wider range of models to learning about a much more generic type of social information, i.e., others' personality traits in various social contexts.

We tested how humans use social knowledge structures for social learning by comparing a set of models that differently formalize reference points and granularity. Specifically, we modelled granularity on two levels: a fine-grained level specifying the similarities between all relevant traits and a coarse-grained level, which can be seen as dimensionality reduction of the fine-grained level where only the factors of the Big-Five are represented. We also investigated which type of granularity affords best fitting learning outcomes and whether participants used optimal representations based on the group of people they were learning about. Additionally, we tested the potential of our models to capture stereotypes. Taken together, we expected that the winning models combine reference points and granularity to refine computational models of social learning.

Results

Experimental Outline

In all five experiments included here, participants rated a number of unknown people on personality traits. Participants never met these people. Instead, they were presented—trial-by-trial—with the self-ratings that these unknown people had previously given (**Figure 1A**). That is, participants could learn from the explicit self-views of other persons (and not via direct interactions).

In each trial, participants rated a person on a given trait (**Figure 1A**). Directly after each rating participants received feedback about the other person's self-rating on the given trait. Through this feedback, participants could learn about the other person's character traits by adapting their ratings to those of the person in question. In each experiment, participants learned consecutively about four or five other persons, i.e., four or five “profiles”. Every experiment was changed slightly to test consecutive hypotheses (see below, **Table 1**, and the section “Differences between experiments” in the methods; **Supplementary Tables 1-3** list the items used in the experiments). The instructions mentioned the possibility to learn about the other persons but participants were not monetarily incentivized to give an accurate impression of the others' traits.

Before we present the results, we briefly outline the different experiments and models. The first experiment functioned as our baseline from which the other experiments were varied.

Experiment 1 [*Real Profiles & Wide Traits*] included “real feedback” from people who had given self-ratings in a previous study on positive and negative traits from all Big-Five factors (Korn et al., 2012). A total of 60 traits were used. A minimum of 8 traits and a maximum of 19 traits were presented per factor (see **Supplementary Table 1**).

Experiment 2 [*Constructed Profiles & Narrow Traits*] consisted of artificially constructed feedback on 30 positive traits from each of the two factors agreeableness and conscientiousness. This experiment allowed a tighter experimental control on a narrow set of traits.

Experiment 3 [*Real Profiles & Narrow Traits*] combined aspects of Experiments 1 & 2 to clarify the differences in their results. Real feedback on 30 traits from each of the two factors agreeableness and conscientiousness was given.

Experiment 4 [*Fashion Models*] used real feedback from self-ratings of an out-group. We chose fashion models because they are members of an out-group that everybody knows but that few people personally interact with. The same traits as in Experiment 1 were used.

Experiment 5 [*IPIP Items*] consisted of feedback from real self-ratings on 50 items from the German translation of the International Personality Item Pool (IPIP). For each of the Big-Five factors, 10 sentences were presented. This experiment thus aimed at generalizing our findings across different types of items describing personality traits.

Figure 1. Overview of the experiment and models used

In this study, we tested computational models of how humans learn about others. In five distinct experiments participants performed a social learning task on several profiles of other persons (with every profile being presented in a separate run).

A) General overview of the learning task: Two trials for Experiment 1 (left) and Experiment 5 (right) are shown. During the learning task, participants estimated which self-ratings a person (here called profile) had given for specific traits (on a Likert scale from “does not apply at all” to “does apply very much”). After each estimate, participants received direct feedback in the form of the actual self-rating of that person. This process continued for all traits (in random order). The tasks shown here are translated to English, original German instructions can be found in **Supplementary Tables 1-3**.

B) To explore participant’s behaviour over the 5 experiments we constructed 5 main computational models (listed in the bottom right). Two main concepts varied in our computational models: First, **Reference Points (RPs)** represent what participants can use as a basis for estimating an average person (from an in- or out-group). Participants can use the RPs on each trait to compare this average person’s rating with their current estimate for the other person. Models 1, 3 and 5 used RPs and models 2 and 4 did not. In all experiments, we use a population mean of the average population (shown here for all items used in Experiment 5). Traits are ordered based on the Big-Five Factors that are

shown in different shades of grey. In Experiment 4, we additionally used the stereotypical ratings of fashion models and we constructed models that use Self-Ratings as RP (**Supplementary Figure 10**). Second, **Granularity (G)** refers to the level of detail in the represented structure of others' personality traits. The granularity matrix generalizes the PEs across similar items in two distinct ways (i.e., per Big-Five factors for "coarse granularity" and per individual traits based on how correlated they are to the current trait for "fine granularity"). Using these two concepts, models can be divided into three sets, which are depicted in three different colours:

1). **No Learning** (blue), consists of a single regression model, Model 1 [No Learning] that functions as a baseline model.

2). **Coarse Granularity** (pink), comprises two models that update their future estimates for the whole (Big-Five) factor to which the current adjective belongs, the first model in this set, Model 2 [Coarse Granularity] uses the standard Rescorla-Wagner (RW) function to update the factor estimates for every trial, the second model in this set, Model 3 [Coarse Granularity & Population RP] uses the aforementioned RW-model to update the factor estimates but also integrates information distilled from the population.

3). **Fine Granularity** (orange), consists of two models that update all adjectives. Crucially, the magnitude of updating on any trait depends on the correlation with the current trait (i.e., during learning one updates not only the estimate of the current trait but also those of all other traits based on how correlated they are to the current item). Like the Coarse models the first model in this set, Model 4 [Fine Granularity] updates all items according to the Fine Granularity, whereas the second model, Model 5 [Fine Granularity & Population RP] updates items not only according to the Fine Granularity but also the RP. More detailed descriptions of the models and their separate functionalities can be found in **Supplementary Figure 1**.

In total there are three free parameters of which the models can make use: α = learning rate (how fast one learns and forgets the past), γ = weighting (weighs between different components of complex models), and starting value (the initial value/ estimate i.e., P_0).

All models that use Rescorla-Wagner learning make use of PEs, an example of the PEs of a randomly selected participant is shown in the top right corner.

P = prediction, Int = intercept, RP = reference point, α = learning rate, PE = prediction error, γ = weighting parameter, F = (generalizes over Factor) coarse granularity, All = (generalizes over All items) fine granularity, SIM = similarity matrix.

Computational Models

In each of the five experiments, we tested (a minimum of) five computational models to explore which model best described participants' learning about others. In our models, we tested two concepts: *Reference Points* (RPs) and *Granularity* (G) (**Figure 1B**). In brief, reference points capture an “average person” (of a specific group). That is, participants may have an average person in mind when rating an unknown person and use this average person to guide their learning. Granularity refers to the level of generalization during learning, i.e., the level of detail in the similarity between represented traits: Coarse representations of trait relationships imply that participants only consider differences between the Big-Five factors—but do not distinguish between traits within a factor. Fine-grained representations take the similarities of all traits into account. In short, granularity generalizes updating based on PEs across similar items.

In the following, we verbally summarize the five models (please refer to the methods section, **Figure 1B** and, **Supplementary Figure 1** for mathematical and conceptual details about the models). Each model is described as if it was learning about one profile, models never combine information or learning across profiles.

Model 1 [*No Learning*] functions as our baseline model. This regression model entails a simple linear transformation of population means (on trait ratings) as reference points to predict others' traits.

Model 2 [*Coarse Granularity*] combines Rescorla-Wagner (RW) learning and coarse granularity. That is, it applies coarse granularity to the personality domain (e.g., learns a single value for each Big-Five factor).

Model 3 [*Coarse Granularity & Population RP*] combines Models 1 & 2 and thus allows to balance participants' use of population means as reference points and learning according to coarse granularity.

Model 4 [*Fine Granularity*] employs RW learning by updating the estimated values of all the traits on every trial based on how similar the other traits are to the trait that is currently presented.

Model 5 [*Fine Granularity & Population RP*] combines Models 1 & 4, i.e., population reference points and learning according to fine granularity.

In Experiment 4, three models are added that are functionally the same as Models 1, 3 & 5 but use stereotypical RPs calculated separately for an out-group (i.e., a group of fashion models).

Experiment 1: Real Profiles & Wide traits

In our first experiment, participants rated four people on 60 self-ratings of trait words sampled from a list of positive and negative German trait adjectives such as polite and aggressive. These traits were selected to be representative of the Big-Five factors, i.e., adjectives for each Big-Five factor were included albeit not in equal numbers (see **Supplementary Table 1**). The real self-ratings of the four other persons were combined with fabricated personal information (e.g., name and age) into four distinct veridical profiles.

Statistical analyses

Analyses that were not based on computational models indicated that participants were learning i.e., the absolute PEs decreased over trials (**Figure 2C, top**). A pairwise Pearson correlation between trial number and the mean of the absolute PE over all participants showed a large negative correlation, $r(58) = -.523$, $p < .001$. Furthermore, we conducted a general linear model (GLM) analysis that consisted of three separate regressors to predict participants' answer accuracy (i.e., higher accuracy means lower PE). Each of the regressors represented a substantial part of the models. Regressor 1) captures learning in the standard Rescorla-Wagner model by tracking the total number of previous trials for each item (i.e., if participants are learning one should see a decrease in PEs over trials). Regressor 2) captures the coarse granularity by tracking the total number of previous trials within a factor for each item (i.e., if participants learn based on each factor, one expects to see a decrease in PEs over trials within this factor). Regressor 3) assesses the fine granularity by computing the summed absolute correlations of the previous items with the current item (i.e., assumes the correlation is the information density of an item to the current item, the sum of all previous items thus predicts the decrease in PE). Results from this GLM indicated the use of the fine-grained correlation structures (i.e., regressor three) (**Figure 2D**).

Computational Models

Bayesian model comparisons revealed Model 5 [*Fine granularity & Population RP*] as the best fitting model among our set of five models according to both fixed- and random-effects analyses (**Figure 2A & Supplementary Figure 2A**). That is, participants relied on previous knowledge about an "average person" when judging the traits of others. Importantly, they not only used population averages as reference points to make estimates but they also scaled the extent to which they updated these estimates by the fine-grained similarity between the currently presented traits and potential other traits in the item set (in line with results from the GLM). These results indicate that participants used a representation of the similarities between traits for their reference group. They represented how single traits relate to one another and used this fine-grained structure in combination with average population ratings to make estimates about others.

Simulations

To assess the winning model's ability to capture learning we simulated its learning (**Figure 2C, bottom**). A pairwise Pearson correlation between trial number and the mean of the absolute PE over all simulations showed a negative correlation, $r(58) = -.563$, $p < .001$. These results were in-line with participants' results. Moreover, we tested which of the five models was the best performing strategy. This analysis disregards participants' behaviour and solely focusses on the models performing the task. These simulations give us an unbiased estimate which model (out of the set of available models) is best suited for the current task. This was achieved by fitting the models on the actual experimental task (i.e., profile responses) rather than participants' responses (see **Supplementary Figure 7** for a detailed explanation of the rationale). Analyses that quantified this model performance corroborated that Model 5 [*Fine granularity & Population RP*] was also the best performing model for this task (**Figure 2B**). In summary, these simulations indicated that participants used a good strategy for learning in this task.

Model Checks

To test the robustness and distinguishability of our models, we performed parameter recovery and calculated a confusion matrix. For parameter recovery, we independently simulated data 200 times using uniformly distributed and randomly sampled parameters with noise added in the last step. Parameter recovery compares known input parameters versus recovered output parameters and uses correlations as a metric of fit, where high correlations indicate better recovery. The correlations for the three parameters of the winning model were deemed satisfactory (see **Supplementary Figures 4-6**): learning rate ($r \geq .87$), weighting parameter ($r \geq .97$) and starting value ($r \geq .92$). Data were also simulated 200 times for construction of the confusion matrix,—now for every model separately. All models were fitted to this simulated data using BIC as the metric of model fit. Ideally, data simulated with a model will also be best fitted by that model, indicating that models can be distinguished. All five models were recovered correctly in $\geq 96\%$ of the simulations, which we deemed excellent (see **Supplementary Figure 3**). Having found evidence that participant employed social knowledge structures for learning about veridical profiles over a wide range of traits, we conducted a second experiment, in which participants performed the same task but were presented with artificially constructed profiles over a narrow range of traits. That is, Experiment 2 entailed more controlled feedback.

Figure 2. Overview of the main analyses for experiment 1. Results indicate that participants used fine-grained correlation structures.

a) Model comparison results from participants' data using fixed-effects analysis. The fixed-effects analysis uses the difference between the summed BIC values for each participant and each model with those from the losing model (here Model 2 [Coarse Granularity (CG)]). This is the reason that the losing model is always zero.

For this experiment, Model 5 [Fine Granularity (FG) & Population Reference Point (RP)] is the best fitting model. This model uses the average population as a reference point and fine granularity for generalization.

b) Simulated data for the 'best performing model'. The best performing model indicates which of the models in the set can simulate the most optimal task performance. That is, the simulated data is fitted on the actual task instead of participants' responses (see Supplementary Figure 7 for a detailed explanation). These simulations use the same fixed-effects analysis as for panel A. However, since these are model simulations, the exact BIC values depend on the total number of simulations (here $n = 36$ simulation). This means that only relative differences between BICs are interpretable.

As for participants' data, Model 5 [Fine Granularity (FG) & Population Reference Point (RP)] also performed best when testing for the best performing model. This indicates that participants used the best strategy (of the strategies in our set).

c) A decrease of the prediction errors (PEs) over time (i.e., trials in the experimental runs) can be interpreted as learning because participants' estimates got closer to the feedback they received. Both plots display the average absolute PEs over time (black line, grey patch: SEM). We calculate a pairwise Pearson correlation between trial numbers and the mean absolute PEs to determine if the PEs decrease over time (least squares line, LSLine, red line).

Top) Participants' data shows a decrease in the PEs over time ($p: .523$). **Bottom)** Simulated data from the best fitting model (Model 5) shows a similar decrease in PEs over time, indicating that the models 'learned' in a similar way to participants.

d) General linear model on core features of our models. Three separate regressions predicted the accuracy (i.e., the PE) per trial per participant. These regressors roughly corresponded to 1) Rescorla-Wagner RL, 2) the coarse models, and 3) the fine models. Detailed descriptions of the regressors are given in the results of experiment one and the methods section (*standard behavioural analysis*).

Each regressor is indicated by a different colour, individual data points are the parameter estimates per participant (summarized by the boxplots). Conclusions should be taken with caution since all three regressors were highly correlated (ρ between .76 - .92). [One-sided t-test; * indicates $p < 0.05$, ** indicates $p < 0.001$, no correction for multiple comparisons].

CG = coarse granularity, FG = fine granularity, BF = Bayes factors, RP = reference point, # = "number of", PEs = prediction errors, SEM = standard error of the mean, RL = reinforcement learning.

Experiment 2: Constructed Profiles & Narrow Traits

In this experiment, we removed the veridical structure of trait self-ratings by artificially constructing feedback for 60 positive trait adjectives belonging to two Big-Five factors, agreeableness and conscientiousness (30 adjectives belonged to each factor) for four profiles. An average value was determined for each factor based on independent self-ratings from a previous study (Korn et al., 2012). To create feedback for each adjective separately, normally distributed noise was added to the average factor to create 30 distinct trait ratings. It is important to note that through this procedure we kept the reference point component consistent with average trait ratings in the population. We only changed the relationship between traits, i.e. the trait similarity structure. We hypothesized that for this experiment using a fine-grained similarity structure (i.e., single trait similarities) should be less advantageous than employing coarse granularity (e.g., average values per Big-Five factor). Otherwise, the experimental set-up was analogous to the first experiment (i.e., participants were not aware that the profiles were artificial nor that they would only learn about two factors). Participants were not given any clues that would encourage them to use a coarse instead of fine granularity structure during learning.

Statistical Analyses

To assess if participants were learning over time (i.e., showed a decrease of PEs over trials) we calculated the Pearson correlation coefficient to compare the mean absolute PEs with the trial number and found a negative correlation, $r(58) = -.546$, $p < .001$ (**Figure 3C, top**). Calculations of a GLM with three regressors, each representing an integral part of our computational models, indicated that participants learned over trials but also seemed to use both coarse and fine granularities during learning (**Figure 3D**).

Computational Models

As hypothesized before we indeed found that the winning model for this experiment was Model 3 [*Coarse Granularity & Population RP*]. Model 3 assumes that participants use a coarse representation of traits for generalizing. It only keeps track of a factor average combined with using an "average person" as RP. Both fixed- and random-effects analyses showed that Model 3 was the best fitting model (**Figure 3A & Supplementary Figure 2B**). This indicates that—in absence of innate similarity structure within traits—participants use the coarser granularity combined with the

population mean as a reference point. Some comment that the modelling results are more trustworthy than the statistical analyses.

Simulations

Model simulations with the winning model showed a similar decrease in the absolute PE over time as participants' data $r(58) = -.489$, $p < .001$ (**Figure 3C, bottom**). Simulations for the best performing model had a better fit for Model 2 [*Coarse Granularity*] than for Model 3 (**Figure 3B**), which indicates that in this experiment participants properly used coarse granularity but could improve their performance by abandoning the use of the reference points altogether.

Model Checks

Tests for model robustness and distinguishability were implemented by performing parameter recovery and calculating a confusion matrix. Parameters for the winning model were recovered satisfactory (learning rate, $r \geq .85$, weighting parameter, $r \geq .98$, and starting value, $r \geq .75$) (**Supplementary Figures 4-6**). The confusion matrix indicated that models were distinguishable with all models being recovered in $\geq 99\%$ from the simulated data (**Supplementary Figure 3**). In Experiment 2, we not only manipulated the feedback but also included traits that belonged to two instead of the five Big-Five personality factors originally used in Experiment 1. To rule out the possibility that this narrower item set was responsible for the observed changes in the winning models, we conducted a third experiment in which participants received veridical feedback (like Experiment 1) but with items that belonged to only two factors (like Experiment 2).

Figure 3. Overview of the main analyses for experiment 2.

Please refer to the caption of figure 2 for a more detailed description of each panel.

a) For experiment 2, Model 3 [Coarse Granularity & Population RP] is the best fitting. This model uses the average population as a reference point and coarse granularity for generalization.

b) Simulated data for the ‘best performing model’. Unlike participants’ data, Model 2 [Coarse Granularity] was the best performing model. This indicates that participants could have used a more optimal strategy by disregarding the RP.

c) Top) Participants’ data shows a decrease in the PEs over time ($p: -.546$), this indicates participants were learning over time. **Bottom)** Simulated data from the best fitting model (Model 3) shows a similar decrease in PEs over time, indicating that the models ‘learned’ in a similar way to participants.

d) All three regressors were significant indicating participants learned over time but also made use of both coarse and fine granularity. However, these regressors were correlated conclusions should thus be drawn with caution. [One-sided t-test; * indicates $p < 0.05$, ** indicates $p < 0.001$, no correction for multiple comparisons].

Experiment 3: Real Profiles & Narrow Traits

Experiment 3 differed from the previous experiments in two distinct ways. First, as in Experiment 2, participants only rated others on traits belonging to two Big-Five factors (agreeableness & conscientiousness). This was a narrower item set compared to the one used in Experiment 1, in which traits from all five factors were included. Second, in contrast to Experiment 2, self-ratings

were sampled from real self-ratings (given by participants from Experiment 2). Similar to the previous experiments, participants rated 60 trait adjectives in total (30 for each of the two factors). As the innate personality structure between traits remained intact, we hypothesized that, similar to Experiment 1, participants would use a fine-grained representation of the traits combined with a reference point of the population mean when rating others traits, i.e., we expected Model 5 to win.

Statistical Analyses

Using Pearson correlation between trial number and absolute PEs revealed a negative correlation ($r(58) = -.42, p < .001$) indicating that participants learned over the course of a run (**Figure 4C, top**). In a further analysis we calculated a GLM with three regressors. These three regressors captured learning over time as well as coarse- & fine-granular learning. We found all three regressors to be significant, indicating that participants learned over time but also made use of both the granularity structures (**Figure 4D**). For an in-depth and more stringent look at the knowledge structures used by the participants we used computational models.

Computational Models

Using a reduced number of factors and veridical feedback in this experiment, we replicated our results from Experiment 1. Model comparison showed that the winning model according to our measures of fixed- and random-effects was Model 5 [*Fine granularity & Population RP*] (**Figure 4A & Supplementary Figure 2C**). This indicates that, in accordance with our hypothesis, even when learning on a reduced number of factors, participants still used fine-grained similarities between traits combined with the population means as reference points to aid learning.

Simulations

We performed model simulations to assess whether the models learned over time. These simulations, using the winning model, showed a decrease in the absolute PE akin to participants ($r(58) = -.764, p < .001$) (**Figure 4C, bottom**). Model simulations for the best performing model confirmed that Model 5, which assumes fine-grained representations of traits combined with RPs in the form of population means, was also the best performing model of the five models in the set (**Figure 4B**).

Model Checks

We tested for model distinguishability by calculating a confusion matrix that showed that all models could be retrieved with $\geq 98\%$ accuracy (**Supplementary Figure 3**). Parameter recovery was also robust with high correlations between employed and recovered parameters (learning rate, $r \geq .83$, weighting parameter, $r \geq .97$, and starting value, $r \geq .78$) (**Supplementary Figures 4-6**).

Conclusion

In the first three experiments, we found evidence that participants consistently relied on average population ratings as reference points. Next we wanted to test if participants would deviate from these standard population averages as a reference point when learning about a subpopulation that could be perceived as dissimilar to participants. To achieve this, we varied the group of the profiles from students, which are the same as the participants (i.e., an in-group), to an out-group (i.e., fashion models), who are likely to be perceived as having a different personality from the students.

Experiment: 3

Figure 4. Overview of the main analyses for experiment 3.

Please refer to the caption of figure 2 for a more detailed description of each panel.

a) Similar to experiment 1, Model 5 [Fine Granularity & Population RP] is the best fitting model for experiment 3. This indicates This model uses the average population as a reference point and fine granularity for generalization.

b) Simulated data for the ‘best performing model’. Unlike participants’ data, Model 2 [Coarse Granularity] was the best performing model. This indicates that participants could have used a more optimal strategy by disregarding the RP.

c) Top) Participants’ data shows a decrease in the PEs over time ($p: -.546$), this indicates participants were learning over time. **Bottom)** Simulated data from the best fitting model (Model 3) shows a similar decrease in PEs over time, indicating that the models ‘learned’ in a similar way to participants.

d) All three regressors were significant indicating participants learned over time but also made use of both coarse and fine granularity. However, these regressors were correlated conclusions should thus be drawn with caution. [One-sided t-test; * indicates $p < 0.05$, ** indicates $p < 0.001$, no correction for multiple comparisons].

Experiment 4: Fashion Models

Based on our experience with Experiments 1-3, Experiment 4 was preregistered on the Open Science Framework (<https://osf.io/8r6gv>) and conducted accordingly. In this experiment we tested participants on the self-ratings of an out-group, i.e., fashion models. We chose fashion models as the out-group because we assumed participants to have prior knowledge (or opinions) about them coupled with a high probability that they would not know a fashion model personally. This meant that participants might change their reference point to a stereotypical reference point. Like the other experiments, self-ratings were combined with fabricated personal information into profiles. Participants were tasked with learning four profiles on all five of the Big-Five factors (i.e., the same stimuli as in Experiment 1). In order to test for stereotypical views, we added three computational models to our regular five models. In these additional models, we changed the previous reference points to an average of fashion model ratings that were assessed before the learning task. Our hypothesis was that participants would still use the fine granularity structure but combine it with this new stereotypical reference point instead of using the student reference point.

Statistical Analyses

First, we investigated whether participants held stereotypical views of the fashion models. Before starting the learning task participants were asked to rate their impression of an average fashion model on all 60 traits that were used in the experiment. From these ratings, we calculated a new reference point that was based on these perceived fashion model self-ratings. To compare the previously used student reference points (M: 6.01, SD: .69) with the perceived fashion model reference points (M: 5.13, SD: .32) we conducted a paired sample t-test, this indicated that there was a significant difference between the two reference points, $t(59) = -9.7137$ $p < .001$.

As in the other experiment we conducted a Pearson correlation analysis to test for a relation between the trial numbers and absolute PEs. This produced a negative correlation $r(58) = -.307$, $p < .001$, indicating participants learned over time (**Figure 5C, top**). A further analysis that used a GLM to test for the use of RL and both coarse and fine granularities indicated, in-line with previous results, that participants used fine-granular representations (**Figure 5D**).

Computational Models

Model comparison confirmed our preregistered hypothesis, both fixed- and random-effects analyses indicated Model 5-STE [*Fine granularity & Stereotype RP*] as the winning model (**Figure 5A & Supplementary Figure 2D**), suggesting that participants used a fine-grained representation of the personality structure together with stereotypical reference points.

Simulations

Simulations of the winning model found a decrease in the abs PE $r(58) = -.788$, $p < .001$, indicating that these models “learned” in a similar way to participants (**Figure 5C, bottom**). Surprisingly, simulations for the best performing model indicated that Model 1 [*No Learning*] would have been the best strategy for this task (**Figure 5B**). After analysing the specific profiles used in experiments 1 and 4 we concluded that there was a higher correlation between the fashion model profiles and the standard population RP (i.e., the average student ratings) compared to the profiles for experiment 1 (**Supplementary Figure 8**). This means that participants did not need all the information that the more complex models provided but rather that the simple *No Learning* model

sufficed in capturing the main complexities of each profile. Interestingly, the second best performing model was Model 5 [*Fine granularity & Population RP*], indicating that the very next best strategy would have also used the standard population RP in favor of the stereotypic RP. The fact that participants still used Model 5-STE [*Fine granularity & Stereotype RP*] even though better strategies were available show both the pervasiveness of the stereotypic RPs and that similarity learning functions as a default strategy.

Model Checks

Both the parameter recovery and the confusion matrix were robust. Correlations between simulated and recovered parameters for the winning model were satisfactory (learning rate $r \geq .81$, starting value $r \geq .74$ and weighting $r \geq .94$). The confusion matrix indicated that all models were distinguishable with all models being retrieved with $\geq 97\%$ accuracy (**Supplementary Figures 3-6**).

Conclusion

In the first four experiments, we found evidence that participants used both the granularity and the reference point in multiple meaningful ways for experiments with trait words as stimuli. In order to test the scope of these models we next tested them on a different set of stimuli. Specifically, not on trait words but statements reflecting traits (e.g. “I’m the life of the party” and “I insult people”).

Figure 5. Overview of the main analyses for experiment 4.

Please refer to the caption of figure 2 for a more detailed description of each panel. In this experiment participants learned about an out-group (i.e., fashion models) instead of the regular in-group (i.e., students). Furthermore, we added 3 models to our original set of 5. These additional models allowed us to capture stereotypic inclinations that participants might have. In brief, every model in the original set that used the population RP (i.e., Models 1, 3, and 5) was duplicated and given a stereotypical RP. This stereotypical RP was calculated by averaging expected fashion-model ratings participants gave before the learning task. Every model that uses the stereotypical reference point is indicated with its original model number with -STE added (e.g., Model 1-STE means the first model but now with the stereotypical reference point instead of that of the student population).

a) As hypothesized, the best fitting model was Model 5-STE [Fine Granularity & Stereotypic RP]. This model uses the expected self-ratings from models as a reference point and fine granularity for generalization.

b) Simulated data for the ‘best performing model’. Contrary to participants’ data the best performing model was model 1 [No Learning]. This indicates that participants use too complex a strategy for learning about the fashion models. The reason this simpler strategy is better suited is explored in **Supplementary Figure 8**.

c) Top) Participants’ data shows a decrease in the PEs over time ($\rho: -0.307$), an indication of learning over time. **Bottom)** Simulated data from the best fitting model (Model 5-STE) shows a large decrease in PEs over time. This indicated that the models ‘learned’ in a similar way to participants.

d) Only the third regressor (summed absolute correlations) was significant, indicating participants used the fine granularity during learning. Due to the high correlations between regressors one should be careful to draw conclusions. [One-sided t-test; * indicates $p < 0.05$, ** indicates $p < 0.001$, no correction for multiple comparisons].

Experiment 5: IPIP items

In our final experiment, preregistered in conjunction with Experiment 4 (<https://osf.io/8r6gv>), we assessed the robustness and generalizability of our models by testing them on a different set of stimuli. For this experiment, we tested the same sample of participants as in Experiment 4 immediately after they finished Experiment 4. We changed the learning task from the trait adjectives to the German translation of the 50-item International Personality Item Pool (IPIP). We selected five profiles from a large online sample (Open Source Psychometrics Project; <https://openpsychometrics.org/>) and used this same sample to calculate the population mean, i.e., the reference point for each trait. We expected that our five models would perform just as well on this experiment since both reference points and granularity can be applied to the personality structures of the IPIP in the same manner. This also led us to expect that participants would use the same models as they did in the previous experiments (i.e., Model 5). We thus hypothesized that participants would use a fine-grained representation of the item similarity structure and the population mean as a reference point.

Standard Analyses & Simulations

A correlation analysis of the absolute PEs found a small negative correlation ($r(48) = -.311$, $p < .001$) indicating that participants learned over time (**Figure 6C, top**). Furthermore, a GLM with separate regressors that each represent a core part of our models (i.e., Rescorla-Wagner learning, and the use of both coarse and fine granularity) did not find evidence for their use by participants, potentially as a result of the smaller answer scale (1-5) and a lesser number of items per factor (**Figure 6D**).

Computational Models

Fixed- and random-effects analyses showed that Model 4 [*Fine Granularity*] performed best (**Figure 6A & Supplementary Figure 2E**). This indicated that participants used a fine-grained representation of the item similarities but did not seem to make use of the reference point.

Simulations

Simulating the winning model on the same dataset revealed a large negative correlation $r(48) = -.868$, $p < .001$, in line with participants' data (**Figure 6C, bottom**). Simulations for the best performing model showed that model 4 (**Figure 6B**) was indeed the best strategy to use, indicating that the reference point did not add enough information to be useful.

Model Checks

Parameter recovery was robust with satisfactory correlations between simulated and recovered parameters (learning rate $r \geq .74$, weighting $r \geq .89$ and starting value $r \geq .51$). Retrieval of models on their own simulated data for the confusion matrix was achieved with an accuracy $\geq 92\%$ for all models (**Supplementary Figure 3-6**).

Conclusion

Results from this experiment are a first indication that the models can be applied to a broader scope of social learning situations.

Experiment: 5

Figure 6. Overview of the main analyses for experiment 5.

Please refer to the caption of figure 2 for a more detailed description of each panel. Participants learned about the German translation of the 50-item International Personality Item Pool (IPIP).

a) Contrary to our hypothesis, Model 4 [Fine Granularity] was the best fitting model.

b) Model simulations confirmed that Model 4 was indeed the best performing model. Indicating that the Population RP did not add enough information to be useful.

c) **Top)** Participants' data shows a decrease in the PEs over time ($\rho: -.311$), indicating participants learned over time. **Bottom)** Simulated data from the best fitting model (Model 4) shows a large decrease in PEs over time. Replicating participants' learning during the task.

d) None of the regressors was significant. This might be due to the smaller answer scale [1-5] and less items in the factors. [One-sided t-test; * indicates $p < 0.05$, ** indicates $p < 0.001$, no correction for multiple comparisons].

Additional Analyses

Positive / Negative Models

Additionally, we fitted all models described above with separate learning rate parameters for positive and negative PEs (**Supplementary Figure 9**). This allowed us to investigate whether participants learned differently from positive and negative feedback. For all experiments, except for experiment 4, these models were worse than the standard models. In Experiment 4, Model 12 [Fine Granularity & Stereotypical RP], with separate free parameters for positive and negative PEs, performed best. Fitted parameters for this model indicated participants learned more from

positive PEs. Interestingly, the total number of positive PEs over all participants was a lot higher than the negative ones (2953 vs 1873), indicating participants might have had too negative a perception of the fashion models.

Self-Ratings as Reference Points

When lacking experience with other groups, a consistent and reliable reference point would be to use one's own self-ratings on traits. For all models that use the population RP we added models that use the participants' self-ratings as RP (**Supplementary Figure 10**) and fitted these models on all experiments. For all experiments, the models using self-ratings were worse than those using the population RP, indicating that participants relied on a more accurate population RP than on their own self-ratings.

Discussion

In the present study, we investigated which potential strategies participants use in social learning about others. We formalized two concepts, which we refer to as reference points and granularity, and constructed a number of models to test their relevance for several social learning tasks over five different experiments. We found evidence for the flexible use of both reference points and granularity across all five experiments using both model-based and model-free analyses.

General Summary

When participants learned about veridical profiles from real people on a wide range of traits, they made use of reference points that summarised the average person of the other group. On top of that, participants in Experiment 1 used a learning algorithm that combined Rescorla-Wagner (RW) learning and a fine granularity social knowledge structure i.e., they generalized updating via PEs according to the similarities between traits.

To delineate this concept of granularity further, we constructed profiles for Experiment 2. These artificial profiles lacked the fine-grained similarities, with only the coarse-grained structure being intact. Remarkably, this manipulation led to different behaviour. Participants made use of coarser personality structures while maintaining the use of reference points. Not only the profiles changed in Experiment 2 but also the number of factors used. To make sure this did not influence the results, we kept the same number of factors (i.e., two factors) but used real profiles once again for Experiment 3. In this experiment, model fit indicated that participants displayed the same behaviour as they did in Experiment 1. That is, participants used a fine-grained similarity structure and reference points based on the average population during learning. As these past three experiments consistently indicated participants' use of reference points, we wanted to explore next whether participants would change their reference points according to the group the person they are learning about belongs to. In order to achieve this, we presented participants with profiles from an out-group (i.e., fashion models) in Experiment 4. Model comparison revealed the use of fine granularity and a distinct reference point specific to this out-group. This indicates that participants held differing views about the out-group under consideration. These prior assumptions about the out-group influenced the employed reference points—while keeping intact participants' understanding of the underlying similarities between traits. These findings illustrate—in combination with the other experiments—that human trait learning seems to rely on both reference points and granularity. In the final experiment, Experiment 5, we explored the generalizability of these models to new stimuli by changing the task from learning about trait words to learning about personality statements. Model comparison suggested participants still used the fine granularity but no longer the reference point for these stimuli. This could indicate that the reference point might not add enough information for such statements; potentially because participants' representation had become too abstracted.

All models displayed high robustness and replicability over all five experiments. Simulations done separately for each experiment indicated that both parameter recovery and confusion matrices were always within a satisfactory range (**Supplementary Figures 3-6**). Results were robust across experiments and for all implementations of the computational models (Wilson

& Collins, 2019) and statistical analyses (e.g., regressions & correlations) largely corroborated these results. Finally, best fitting parameters for simulations and participants' data largely overlapped (**Supplementary Figures 11-16**). Participants seemed to use both granularity and reference points throughout the experiments, often together in their most complex form (i.e., Model 5), exemplifying the complexity of human social learning. These results, coupled with the subtle differences in learning that are picked-up by the models for the varying experiments, highlight the usefulness of both granularity and reference points in computational models and indicate their importance for guiding social learning. In the following, we discuss these model components.

Prediction Errors

Across five experiments, we found evidence for the use of PEs in learning, which is in line with a host of previous work on (non-)social learning (Rodriguez Buritica et al., 2015; Diaconescu et al., 2017; Garvert et al., 2015). From a statistics perspective we found that participants showed, on average, a decrease of the absolute PEs over time for all experiments, importantly these results were replicated by simulations of the winning models (**Panel C, Figures 2-6**). These findings reinforce the validity of using the hybrid Rescorla-Wagner learning models, because the PEs are integral to these models.

In the current study we focused on learning about others, in this context we did not expect any difference in positive and negative feedback (i.e., PEs). Nonetheless, we performed an additional analysis that constructed models based on the existing five such that they took into account positive and negative feedback separately. Results from all five experiments largely corroborated our expectations: the regular models were better in every experiment, with a better fit for the positive/ negative models only being achieved in Experiment 4. In future research, we expect that positive and negative PEs are more influential when learning about oneself, as previous work has shown (Sharot et al., 2011; Korn et al., 2014). For example, stereotypical or plain false beliefs about others (as often seen in patients with borderline personality disorder Barnow et al., 2009) may result in excessive positive or negative PEs.

Reference Points

We defined reference points as “an average person (of a group) one has in mind” and formalized them as the average of a group of independent self-ratings (i.e., individual trait ratings calculated from an independent sample). RPs were used as points of comparison during learning by participants in all experiments except for Experiment 5. We surmise the absence of RPs in Experiment 5 to be related to the smaller scale of the answer options (i.e., 1-5 compared to 1-8 for the other experiments) which decreases their usability as well as the relatively small amount of items within a factor (10).

The most striking use of the RP was exhibited in Experiment 4, where participants were tasked with learning about people from an out-group (i.e., fashion models). During this learning task they used a different (stereotypical) RP than the one they had been using for the other experiments where they learned about people from an in-group (i.e., students). Interestingly, the best performing models indicated participants would be better off using Model 1 [No Learning].

Additional analysis that compared the fashion model profiles and the student profiles from experiment one showed higher correlations among the fashion model profiles and the standard population RP. This indicates that one could potentially estimate the average fashion model accurately with only the population RP available (**Supplementary Figure 8**). Notably, participants were persistent in their use of a learning model scaled by fine granularity knowledge and the stereotypical reference point. We interpret this finding as a case of “stereotypical reference points” in line with Jussim et al., (1995) who define stereotypes as cognitive categories that people use when thinking about groups and about individuals from those groups. It is widely accepted that stereotypic expectancies can guide learning (Hamilton, Sherman, and Ruvolo, 1990), as exemplified in Experiment 4.

Finally, we introduced a separate set of analogous models with participants’ self-ratings as RPs for all experiments (**Supplementary Figure 10**). These models introduce a ground truth i.e., one always has access to their own traits. Results from these experiments indicate that the average (perceived student) RPs to always be slightly better, indicating the higher accuracy from such perceived group averages compared to self-ratings.

Granularity

Granularity was introduced to create computational models that are more akin to real-life learning, which necessitates generalization across similar stimuli and situations. Previous learning models only allowed updating about isolated traits. However, in a real-life situation humans rarely learn about traits in an isolated manner but rather update their whole view of a person’s personality (e.g., if you see someone being impolite your view of that person as a whole changes in a negative manner).

Granularity was defined on two levels of detail with which participants could represent the personality space. Participants seemed to use both levels of granularity in our models depending on the specifics of each experiment. We argue that coarse granularity roughly corresponds to the use of schemata (Mayer & Bower, 1986). Schemata are representations of features that often co-occur and influence learning and perception (Mayer & Bower, 1986). These representations or cognitive structures contain units of information and their links (Fiske & Linville, 1980; Taylor & Crocker, 1981). Likewise, our implementation of coarse granularity entailed the factors of the Big-Five (i.e., items within a factor occur together most often and are summarized by a single average value). Conversely, updating with fine granularity linked all traits together based on their correlations. Furthermore, splitting granularity in two levels of detail was inspired by work of Klein et al., (1992) who showed that trait-representations are stored as two distinct memories (episodic and semantic). Episodic memories correspond to single exemplars of social information and semantic representations are more abstract and generalized.

Both coarse and fine-grained structures described participants’ behaviour depending on experimental manipulations. Interestingly, this indicates that participants represented and accessed both types of structures when beneficial. Our findings might indicate that humans flexibly change their representations of social knowledge to match implicit task demands. Importantly, participants could only derive clues about how to adapt their strategy from trial-level item and feedback

information, as the general task frameworks and instructions were kept very similar across experiments. We believe that this flexible updating of abstract representations also allows for the ability to generalize across widely differing stimuli, as evidenced in Experiment 5.

The granularity of knowledge structures employed in this study were chosen to represent two extreme levels (coarse and fine). In many cases, humans probably do not use either of the extremes but a level of generalization that lies on a continuum between these extremes. Because coarse- and fine-granularity are two levels of granularity, they are inherently related. We attempted to tease their unique contribution apart with a GLM analysis, this was problematic due to the high correlation between the two regressors. Results from the confusion matrices of all experiments indicated that models could differentiate between the different levels of granularity. This made us confident that the computational modelling approach was better suited in distinguishing between these two levels of granularity. In future work we aim to investigate the representations at various levels of granularity across this continuum.

Scope of models

Social learning is a very broad domain. Even within a certain social learning domain, there is large situational variability that determines which learning strategy is most suitable. In five consecutive experiments that vary with respect to participant populations, items and the to-be learned profiles, we introduced a defined set of computational models that generalize across experiments and pick up the systematic and subtle differences between experiments. The goal of this manuscript was to introduce the modelling framework in a clear and concise manner and not be exhaustive in introducing its potential applications. Nonetheless, we would like to argue for the potential broad scope of their applicability to a majority of social learning settings. As a primer for this broader scope we included experiment 5, in which we used a different set of stimuli for which, as expected, the models still functioned like the former experiments.

We conjecture that these models are especially useful for learning about personality traits, as shown over the five experiments. However, a very similar model space has also successfully been applied to learning about others' preferences (Rosenblau et al., 2020). The models integrated similarity structures that can be derived from preference items (coarser grained structures such as candy and fruit and fine-grained item-level correlations). A fine-grained similarity learning model captured how participants learned about others' preferences best.

Moreover, recent studies have found neural evidence that humans encode specific social knowledge structures employed in our models (i.e., fine grained correlation matrices) (Stolier et al., 2020; Thornton & Tamir, 2021) suggesting a central role of item-level similarity structures in social cognition. The relative simplicity of our models makes them adaptable to various multidimensional representations across different learning domains. These results are corroborated outside of the social domain by Roweis & Saul (2000) who highlight human expertise in representing complex abstract structures akin to the granularity structure introduced here.

In a general modelling framework we argue that our models relate to differences on a continuum between model-free (MF) and model-based (MB) RL (Dolan & Dayan, 2013). The details captured by granularity reveal complexities that MF-RL does not capture: MF learns a single

value for any given stimulus in a given learning situation. Likewise, our models with coarse granularity learn a summary value for each of the Big-Five factors. In their simpler form, we thus think of our models using coarse granularity to be more similar to MF-RL rather than to employ full MB-RL (i.e., coarse granularity functions as a look-up chart of traits). The more complex models with fine granularity can be viewed as more like MB-RL as they incorporate a full representation of the traits and the similarities between them. This split—akin to differences between MB-RL and MF-RL—suggests that our models can distinguish between more costly optimal models and more efficient heuristic models, similar to models on optimal versus heuristic decision-making. (Korn & Bach, 2019; Korn & Bach, 2018).

Limitations

Studying social interactions in a well-controlled experimental setting forces limitations on a study (i.e., in the current experiment some limitations were necessary to establish the current modelling paradigm). First, even though we used RW learning models to capture participants' behaviour. The task, in which participants received direct and exact feedback, bears some resemblance to a supervised learning problem. Supervised models have thus far not been used to explain our task and could pose as interesting models for future research.

A related limitation is that the task and thus feedback was received in exact numbers. In a social situation these number would normally be substituted for a verbal description (e.g., “I am very generous”) or an action (e.g., sharing resources shows ones generosity). When learning about others in most social settings people seldom receive exact or direct feedback. Precise number feedback as used in our task is more common in the work or school setting as part of performance evaluations. Future studies can generalize from precise numbers to less precise, verbal quantifications to resemble the most commonly used form of social evaluation.

Furthermore, learning about others' personality is sometimes more directly related to actions e.g., approaching someone who seems friendly. Expanding the task to include actions based on feedback about someone's personality (e.g., a cooperative foraging task) could therefore further heighten the external validity.

The total amount of trait items used is on the small side i.e., for experiment 5 we used 50 items (10 per factor) but there exist questionnaires with 120 or even 300 items. Expanding the total number of items could offer a more precise look at learning about personality traits. Specifically it would offer a more rigorous look at the continuum between the coarse and fine-grained extremes.

Future Research

Future research should explore how people learn about the reference points and granularity in the first place. Based on the suggestions by Klein et al. (1992) that memory moves from a reliance on exemplars to more abstract representations, it would be interesting to explore if the concept of granularity also evolves over time during learning—concurrently exploring if humans use even more levels of granularity between the current two extremes and how these levels get adjusted. Furthermore, if and how humans switch within reference points and granularity within and between tasks are interesting avenues for future research. Reference points could be explored with a host of

other groups (e.g., other out-groups, (racist) stereotypes, etc.), especially since learning from others seems to be influenced by whether they belong to an in- or out-group (Kang et al., 2021; Hein et al., 2010).

The literature in social neurosciences points towards parts of the (medial) prefrontal cortex as one of the most important regions for social cognition (Tamir & Thornton, 2018; Olsson et al., 2020; Anzellotti & Young, 2020). It would be interesting to explore if and how the concepts that we introduce in these experiments are represented in the brain and if the changes mentioned above can be tracked as dynamic processes in the human prefrontal cortex (Mende-Siedlecki, 2018).

As mentioned previously, we expect that the learning models presented here will also be applicable to different domains of learning such as other domains of social learning (e.g., learning for others; Ruff & Fehr, 2014)—and potentially to different types of non-social, abstract learning.

Conclusion

In summary, we tested how humans learn about the multi-dimensional personality of others. We specifically focused on how social knowledge structures shape this learning and introduced two social knowledge structures that can be mathematically specified in computational models. We found evidence that participants use prior convictions akin to schemata to learn about others who belong to different groups. Crucially, participants displayed a representation of the multidimensional similarity structure of personality traits and were able to apply this on several levels of abstraction (within constraints). The computational models introduced in this study are robust, simple, and widely applicable to multidimensional (social) learning situations.

Methods

General

We present five experiments. All experiments are similar to each other and vary on a couple factors per experiment in order to test specific components within our models. In particular, these variations aim at distinguishing: Reference Points, Granularity and overall robustness of the models.

Participants

Participants for all experiments were recruited through online advertisement. Experiments 1-3 were administered to independent participant samples. Experiments 4 & 5 were tested on the same sample of participants. Participants' details per experiment are summarised in **Table 2**. Participants had to meet the following inclusion criteria: 1) age between 18 and 40 years, 2) German native speakers, 3) normal or corrected to normal vision, and 4) no history of neurological or psychiatric disorders. In Experiments 1-3, participants could only partake in one of the experiments; participation in one experiment thus directly excluded them from partaking in any of the four other experiments. The same group of people took part in Experiments 4 & 5 and were tested in the same testing session. The study was conducted in accord with the Declaration of Helsinki and approved by the local research ethics committee (Ethik-Kommission der Ärztekammer Hamburg, Number: PV5746). All participants gave written informed consent using a form approved by the ethics committee and were compensated the regular hourly fee for behavioural studies. On average, one experiment took around 30 minutes.

Preregistration

Experiments 4 & 5 were preregistered on the Open Science Framework prior to data collection (<https://osf.io/8r6gv>). In this preregistration, we specified our expected sample size, exclusion criteria, measured variables and computational models. Most of these considerations also applied to the first three experiments, which we did not preregister.

Task

All five experiments shared the same structure, but differed in their content (i.e., words and profiles). The first experiment is described below. Any differences between Experiment 1 and the following experiments are detailed in the section “differences between experiments.” The main task in the experiments was a social learning task followed by (self-) rating tasks.

Participants performed a social learning task about different trait profiles, where a profile consisted of sparse information about a person (name, age, time since starting university studies and time spent living by themselves) and self-ratings on a selection of traits selected to resemble items found in the Big-Five (**Supplementary Tables 1-3**). The information and self-ratings from these profiles were selected from an independent previous study (Korn et al., 2012), with fictitious and randomly assigned common names. Instructions for the whole experiment were given orally and presented on the screen at the start of the experiment. Participants were given the opportunity to ask questions after the instructions and could start the experiment themselves through a button

press. The learning task consisted of four runs, i.e., one run per profile. Participants learned about a different profile in each run (**Figure 1A**). As described above, the person of the profile in question was briefly introduced at the start of the run. Two of the profiles were given female names and two of the profiles were given male names. Each experimental trial started with the presentation of a trait word (e.g., generous, arrogant), a blank space underneath, and the rating scale (1=does not apply at all; 8=does apply very much). Participants had four seconds to estimate how they thought the persons in question rated themselves on this trait by choosing the appropriate number [1-8] on the keyboard. Directly after this estimate participants received feedback (i.e., self-ratings from the person of the profile) for a duration of two seconds, for a total of 60 trials. This means that participants could learn about each person's self-ratings over time.

After the social learning task, participants were asked to give their self-ratings on 80 traits, i.e., the full list of 40 positive and 40 negative words used in previous studies (Korn et al., 2012; 2014). A subsection of 30 positive and 30 negative traits was used in the learning part of the experiment.

Consistent with the previous studies, the anchor points for all the rating scales were “1=does not apply at all” and “8=does apply very much.” All experiments were presented using the MATLAB toolbox Cogent (www.vislab.ucl.ac.uk).

Differences between experiments

Tables 1 & 2 summarise the differences between all five experiments. The first experiment was used as a reference to subsequent experiments. In Experiments 2-5, we manipulated key features of the first experiment and tested whether these changes resulted in different fits of the computational models. **Supplementary Tables 1-3** list the stimuli used.

Experiment 2: Constructed Profiles & Narrow Traits

Experiment 2 consisted of artificial profiles of a narrower trait space. That is, 60 trait words were selected pertaining to the two Big-Five factors agreeableness and conscientiousness (with 30 positive traits for each of the two factors). Specifically, 15 trait words were selected from the original list of 80 traits from the previous studies. A set of 45 trait words were newly introduced, which resulted in a total of 125 trait words used overall. In Experiments 2-4, participants made self-ratings on all of these 125 words.

The “self-ratings” of the profiles about which participants in Experiment 2 learned were simulated. Specifically, we specified a mean value for each of the two factors, which could be 5 or 7 on the 8-point scale. We specified four profiles with the four possible combinations of a mean of 5 or 7 for the two factors (i.e., 5 & 5, 5 & 7, 7 & 5, and 7 & 7). For each trait word, we individually added random variability ($SD = 1$) using the standard random normal distribution (`randn`) in MATLAB. After adding this random variation, ratings were rounded to their nearest integer and capped to the anchor points [1-8]. All four profiles were given names that matched the gender of the participant.

Experiment 3: Real Profiles & Narrow Traits

Profiles in Experiment 3 were veridical self-ratings belonging to two of the Big-Five factors (agreeableness & conscientiousness). We selected these profiles from self-ratings given by participants of Experiment 2.

Experiment 4: Fashion Models

Profiles in Experiment 4 were selected from a pre-defined reference group. That is, from a sample of 30 female fashion-models, who provided online self-ratings in a separate study conducted for a Bachelor's thesis (unpublished), from these 30 self-ratings we selected four profiles for the learning task. Information given to participants in the learning task was: female first name (fictionalized), age, years of work as a model, and time spent living by themselves.

At the start of the experiment, we asked participants for pre-ratings to capture their view of the average fashion model's personality on the used stimuli. After giving these pre-ratings, participants completed the standard learning task. They rated four fashion model profiles and received trial-by-trial feedback about the fashion models' own ratings. Both the pre-rating and learning task were performed on the same 60 traits and five factors as in Experiment 1.

Experiment 5: IPIP items

A different set of items was used for Experiment 5 (**Supplementary Table 3**). Instead of the trait words used in Experiments 1-4, we used 50 items from the German translation of the International Personality Item Pool (IPIP), also known as "Lexical Big-Five factor markers" (Goldberg, 1992). These items were rated on a 5-point scale as those in the original questionnaire (1= does not apply at all; 5=does apply very much).

Information for selection of the profiles was taken from a large online dataset with 50 self-report items on the IPIP from over 1 million participants (Open Source Psychometrics Project; <https://openpsychometrics.org/>). For our learning task, we selected five profiles with average ratings on 4 out of the 5 factors (mean score within 1 SD) and a divergent score on the remaining factor (mean score above 1 SD). That is, each profile is divergent on one out of five factors. After the learning task, participants judged themselves on the same 50 items.

Data Analysis

Missing Data

The same exclusion criteria for missing data were applied for all experiments. Within an experiment, a single run would be removed if more than 20% of the answers were missing (due to participants not answering within four seconds). Data from a participant was completely removed if more than 10% of all answers were missing. This resulted in the exclusion of three complete data sets (2 data sets in Experiment 5 and 1 data set from Experiment 4) and 9 separate profiles over all experiments combined (see **Table 2** for details). These exclusion criteria are in line with our preregistration for Experiments 4 & 5.

Statistical Behavioural Analysis

To test whether the PEs have a downward trend over time (an indication of learning), we calculated the average of the absolute PEs per profile per trial for all participants. Profile data were then averaged into one vector (with length being the number of trials) of the average absolute PEs over time per experiment. To determine whether there was a significant decrease of PEs over time, we calculated the Pearson correlation coefficient on these absolute average PEs and the corresponding trial number. A negative Pearson correlation indicates a decrease in the absolute PEs over the trials. The same procedure was applied to data that was simulated based on the winning (i.e., best fitting) model for that experiment.

Moreover, for each experiment separately, we calculated a GLM that explained the accuracy per trial (i.e., the prediction error) with three regressors. Each regressor was initially tested in a separate regression and captured an integral part of our computational models. In brief, the regressors were the following:

Regressor 1: total number of previous trials seen for each item, (this assesses the relationship between the decrease of the PEs and the number of items seen previously and thus learning in the standard Rescorla-Wagner model)

Regressor 2: total number of previous trials that are from the same factor as the current item (this assesses the relationship between the decrease of PEs and the number of items encountered from a specific factor and thus investigates the behaviour captured by the coarse granularity models)

Regressor 3: the summed absolute correlations of the previous items with the current item (this assesses the fine granularity models where the information content of all the previous items is weighted by their correlation to the current item).

In a second-level analysis, participants' individual parameter estimates were subjected to a one-sided one-sample t-test to test if the slope was significantly different from zero in the negative direction (indicating a decrease of the absolute PE over time).

Computational Models

We created computational models that use standard Rescorla-Wagner learning and expand these models with two sources of information: Granularity and Reference Point. Through our computational models, we explore to what degree participants use these concepts (Figure 1B, Supplementary Figure 1). We test participant learning on individual profiles but do assume that participants use the same strategy (i.e., model) throughout the experiment.

Calculation of the PEs

Models 2-5 make use of the standard Rescorla-Wagner learning model, this model makes use of the prediction error (PE) to update the prediction (P) for the following trial. The prediction error for all models is the feedback (F) on a certain trial (t) minus the prediction (P) on that trial.

$$PE_t = F_t - P_t$$

Granularity

Granularity refers to the level of detail with which participants represent the others' traits. This means that one can either have a summary value per (Big-5) factor, or learn a separate value for each trait per person (**Figure 1B**, and **Supplementary Figure 1D, E**). The former, Coarse Granularity, is mathematically defined by having a single summary value for each factor which will be updated only when one learns about a trait that belongs to that factor. The latter, Fine Granularity, assumes one learns a separate value per trait but also updates all other traits based on how similar (i.e., correlated) they are to this current trait. In experiments 1-4, these similarity matrices were calculated using Pearson's correlation on self-ratings from two published independent laboratory studies (Korn et al., 2012, 2014) as well as data from previous online studies (see Oganian, Heekeren, Korn, 2018). Specifically, we selected four participants from the 27 participants in Korn et al. (2012) for the four profiles in Experiment 1. The self-ratings of the remaining 23 participants as well as the self-ratings from the 78 healthy control participants in Korn et al. (2014) were used for the similarity matrix. These 101 participants who completed laboratory studies gave ratings on the 80 traits of the original list. Additionally, for the calculation of the similarity matrix we included online ratings of 734 participants, who each gave self-ratings on 50 pseudorandom traits from the overall list of 125 traits. For Experiment 5, we calculated the correlation matrix from the (> 1 million) sample from the Open Source Psychometrics Project (<https://openpsychometrics.org/>).

Reference Points

Reference points refer to a-priori expectation of a person's personality. This means that a person has an idea of the trait ratings from an average person who belongs to the same group as the person who is being judged and uses these ratings next to regular learning (**Supplementary Figure 1B, C**) (e.g., most people from this specific group are generous I therefore should rate them higher on generous, but I perceive them, on average, to have less diligence, so I should score them lower on this trait). The reference point data was calculated by taking the mean value per trait of the same datasets that were used to calculate the similarity matrices. In Experiment 4, we used the stereotypical ratings given by participants before the learning task. To test whether there was a difference between the self-ratings of the independent study and the stereotypical ratings we calculated an independent samples t-test on the trait averages.

Model 1: No Learning

Model 1 assumes that participants perform a linear transformation of the reference point (RP) population mean to predict (P) others' personalities. This model performs like a standard linear regression where b_0 represents the intercept and b_1 the slope.

$$P = RP * b_1 + b_0$$

Model 2: Coarse Granularity

Updates an average factor value (F, factor) for the next trial estimate based on the prediction error (PE) and the current estimate. That is, on a trial-by-trial basis the model updates a value that represents the average value for this factor. The speed of learning is determined by the free parameter α (bounded [0 1]), with a second free parameter; starting value (bounded [1 8]), determining the value at which each factor will be initialized. That is, the starting value determines the first ‘guess’ the model makes when a new factor or profile is presented (i.e., the value at $P_{(0,F)}$).

$$P_{(t+1,F)} = P_{(t,F)} + \alpha * PE$$

Model 3: Coarse Granularity & Population RP

Model 3 expands Model 2 by adding the reference point (RP) in form of a population mean per trait. During each trial model 3 learns an average factor value like model 2 but it adds information based on the Reference Points (i.e., the average population rating on this trait). Information is integrated using the weighting parameter γ (bounded [0 1]). This parameter determines how much participants rely on just the RW from model 2 or the RP (0.5 indicating that both are used equally). Like Model 2, Model 3 also uses the parameters learning rate (α) and starting value for its initial estimate for a new factor or profile. Parameters α , γ , and the starting value are free parameters.

$$P_{(t+1,F)} = \gamma * RP + (1 - \gamma) * (P_{(t,F)} + \alpha * PE)$$

Model 4: Fine Granularity

Model 4 employs fine-grained granularity and updates all items (All) at once based on how correlated they are to the current item (**Supplementary Figure 1E**). This similarity matrix (SIM) was calculated before the experiment based on separate samples (see section Granularity). Thus on a trial-by-trial basis Model 4 updates the current trait based on the PE and the learning rate (i.e., it is perfectly correlated with itself) all the other traits will get updated based on the PE, the learning rate and their similarity to the current item (i.e., their correlation), which means that traits that are more similar get updated more. Model 4 also makes use of the free parameter starting value. Because this model updates all items for every step of learning the starting value is not just initialized for the first value but rather all items. This ensures that the model can updates all items from the onset of learning.

$$P_{(t+1,All)} = P_{(t,All)} + \alpha * PE * SIM$$

Model 5: Fine Granularity & Population RP

Like the coarse granularity models, Model 5 expands the previous model (i.e., Model 4) by adding the reference point (RP) in the form of the population mean. That is, for every trial learning happens (as in model 4) but additional information in the form of an average trait rating is added. The degree to which this information (i.e., RP and Granularity) is integrated is determined by γ . Similar to Model 4, Model 5 uses the starting value to initialize all item estimates at the onset of learning. Free parameters are α , γ , and starting value.

$$P_{(t+1,All)} = \gamma * RP + (1 - \gamma) * (P_{(t,All)} + \alpha * PE * SIM)$$

Model additions: Stereotypical Reference Points

In the preregistration for Experiment 4, we added a variation of the reference point to the models specified above. Next to the standard reference points, we also tested models using the stereotypic view held by the participant prior to learning. This stereotypic view in the reference point (STE) is added for Models 1, 3 & 5.

$$\begin{aligned} \text{Model 1-STE} \quad & P = STE * b_1 + b_0 \\ \text{Model 3-STE} \quad & P_{(t+1,F)} = \gamma * STE + (1 - \gamma) * (P_{(t,F)} + \alpha * PE) \\ \text{Model 5-STE} \quad & P_{(t+1,All)} = \gamma * STE + (1 - \gamma) * (P_{(t,All)} + \alpha * PE * SIM) \end{aligned}$$

Model additions: Positive / Negative Learning Rate

To explore if participants learned differently from positive and negative feedback, we constructed additional models that use two parameters for the learning rate. These two parameters thus have a separate learning rate for positive and negative PEs. All models except for model 1 [No Learning] were adapted in this manner (i.e., Models 2-5 all make use of the following two α 's that depend on the sign of the PE in the current trial).

$$\begin{aligned} a &= \alpha^- \text{ if } PE < 0 \\ a &= \alpha^+ \text{ if } PE \geq 0 \end{aligned}$$

Model additions: Self-Ratings as Reference Points

Similar to the stereotypical models in Experiment 4 we included additional models that used participants' self-ratings as RPs. To achieve this the models that originally used the population RP (i.e., Model 1, Model 3, and Model 5) were duplicated and adapted to use each participants' self-ratings, resulting in 11 total models for Experiment 4 and 8 total models for the other Experiments. Model comparisons were conducted as previously specified.

$$\begin{aligned} \text{Model 1-Self} \quad & P = Self * b_1 + b_0 \\ \text{Model 3-Self} \quad & P_{(t+1,F)} = \gamma * Self + (1 - \gamma) * (P_{(t,F)} + \alpha * PE) \\ \text{Model 5-Self} \quad & P_{(t+1,All)} = \gamma * Self + (1 - \gamma) * (P_{(t,All)} + \alpha * PE * SIM) \end{aligned}$$

Model Fitting and Comparison

For model fitting, the free parameters are all initialized at the average between the maximum and minimum bounds (i.e., 0.5 for both α and γ). The first prediction (P_0) was fitted as the free parameter: starting value.

Optimal model parameters were determined using linear least squares estimation. In detail, optimization used a non-linear Nelder-Mead simplex search algorithm (implemented in the MATLAB function *fminsearch*) to minimize the sum of squared errors (SSE) of prediction over all trials for each participant. Model evidence was then assessed using the Bayesian information criterion (BIC) using the formula:

$$BIC = n * \ln\left(\frac{SSE}{n}\right) + k * \ln(n)$$

Here, n is the number of trials and k is the number of free parameters in the model (BIC thus penalizes for an increase in model complexity). Models were compared using both fixed- and

random-effects analyses. For fixed-effects analyses we calculated the log-group Bayes factor (BF): all model BIC scores are summed across participants and then subtracted from the value of the reference model (worst scoring model) where the best performing model has the lowest score. Furthermore, for random-effects analyses, we calculated the posterior exceedance probability. This measures the likelihood that any model is more frequent in the comparison set than all other models, corrected for the possibility that observed differences are due to chance. This random-effects family wise comparison was performed using the Bayesian Model Selection (BMS) procedure implemented in the MATLAB toolbox SPM12 (<http://www.fil.ion.ucl.ac.uk/spm/>; *spm_BMS*, Rigoux et al., 2014).

Best Performing Models (in the set)

Normally, models are fit on the answers participants give during the task, this helps us estimate which of the strategies (i.e., models) participants used during learning. But in order to understand the usefulness of the strategies we estimate the best performing model (in the set). This is achieved by fitting the models on the profile answers (i.e., the same way the participants experienced the experiment) instead of the participant answers (**Supplementary Figure 7**). Comparing these best performing models with the ones that best fit the participants' strategies we can determine whether participants' strategies were well fit to the task demands.

Confusion Matrix

To test whether data simulated with a model would also be best fit by that model, we created confusion matrices for every experiment separately. For every confusion matrix we drew random parameter values between [0.2 0.8] 200 times. These values were used to simulate data for the profiles using every model. After the data was simulated, noise was added from a standard normal distribution. Then all models were fitted to the simulated data to find out the extent to which simulated data from a particular model is best fitted by this model opposed to all other models (**Supplementary Figure 3**). Ideal recovery is achieved when the confusion matrix results in the identity matrix.

Parameter Recovery

To test whether parameters could be recovered consistently, we simulated data for all five experiments. For every experiment separately we simulated data for the profiles by drawing two random parameter values (between [0 1]) from a random distribution 200 times. With these parameters data was simulated using every model. Then noise (sampled from a normal distribution) was added. After this, all models were fitted to this data and parameters were recovered. We measured fit through correlations; where a higher correlation means better fit. Winning models were marked using green squares in **Supplementary Figures 4-6**. For Experiment 4, we omitted the stereotype models as these are functionally the same as the standard models when it comes to simulating data.

Tables

Table 1. Overview of the stimuli for the five experiments

Experiment	Stimuli	n Factors	Profiles	n Profiles	Scale	n Trials
1 (Original)	60 traits	5	Real (Online)	4	1-8	240
2 (Constructed Profiles)	60 traits	2	Constructed	4	1-8	240
3 (Two Factors)	60 traits	2	Real (Online)	4	1-8	240
4 (Fashion Models)	60 traits	5	Real (Fashion)	4	1-8	240
5 (IPIP items)	50 IPIP items	5	Real (IPIP)	5	1-5	250

In Experiments 1-4, we used personality adjectives (i.e., trait words such as generous, diligent; see *Supplementary Table 1* for Experiments 1 & 4 and *Supplementary Table 2* for Experiments 2 & 3). Sixty adjectives were presented per profile (i.e., for each of the four people about which participants learned). We used traits from all five factors of the Big-Five or traits from only two factors (i.e., agreeableness and conscientiousness). In Experiment 5, for each of the five profiles, participants saw 50 items from the German translation of the IPIP (International Personality Item Pool, which consists of lexical Big-Five factor markers; see *Supplementary Table 3*). Profiles for Experiments 1 were selected from self-ratings of people from an unrelated sample of a previous lab study (Korn et al., 2012). Profiles for Experiment 2 were constructed by specifying the mean for the two factors and randomly adding noise according to a specified SD. Profiles for Experiment 3 were selected from the self-ratings of participants in Experiment 2. Profiles for Experiment 4 were selected from self-ratings given by a group of female fashion models for a related online study. All four selected persons have worked internationally as fashion models for several years. Profiles for Experiment 5 were selected from self-ratings of a large online dataset on the IPIP with over 1 million participants (Open Source Psychometrics Project; <https://openpsychometrics.org/>). We selected five profiles with average ratings on 4 out of the 5 factors (mean within 1 SD) but divergent scores on the remaining factor (mean above 1 SD). That is, each profile was divergent on another factor.

1202 **Table 2. Overview of the participants for the five experiments**

Experiment	n Tested	Age (mean)	Age (SD)	n Female	n Final	n Profiles removed	n Sessions removed
1 (Original)	35	24.37	3.66	23	35	0	1
2 (Constructed Profiles)	42	23.67	3.19	28	42	0	0
3 (Two Factors)	59	25.37	4.95	30	59	0	2
4 (Fashion Models)	30	25.38	4.91	15	29	1	0
5 (IPIP items)	30	25.38	4.91	15	28	2	8

1203 The same group of people participated in Experiments 4 & 5.

1204

1205 **Data Availability**

1206 The behavioural data that support the findings of this study are publicly available on Github
1207 (<https://github.com/dnhi-lab/PerLe>) and on Zenodo (<http://doi.org/10.5281/zenodo.4697286>)

1208

1209 **Code Availability**

1210 The code used for analysis is publicly available on Github (<https://github.com/dnhi-lab/PerLe>) and
1211 on Zenodo (<http://doi.org/10.5281/zenodo.4697286>)

1212

References

- Amodio, D. M. (2008). The social neuroscience of intergroup relations. *European review of social psychology*, 19(1), 1-54.
- Anzellotti, S., & Young, L. L. (2020). The Acquisition of Person Knowledge. *Annual review of psychology*, 71, 613-634.
- Barnow, S., Stopsack, M., Grabe, H. J., Meinke, C., Spitzer, C., Kronmüller, K., & Sieswerda, S. (2009). Interpersonal evaluation bias in borderline personality disorder. *Behaviour research and therapy*, 47(5), 359-365.
- Bellucci, G., Molter, F., & Park, S. Q. (2019). Neural representations of honesty predict future trust behavior. *Nature communications*, 10(1), 1-12.
- Brewer, W. F., & Treyens, J. C. (1981). Role of schemata in memory for places. *Cognitive psychology*, 13(2), 207-230.
- Camerer, C., & Hua Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4), 827-874.
- Cantor, N., & Mischel, W. (1977). Traits as prototypes: Effects on recognition memory. *Journal of personality and social psychology*, 35(1), 38.
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: the good, the bad and the ugly. *Current opinion in neurobiology*, 18(2), 185-196.
- Delgado, M. R., Frank, R. H., & Phelps, E. A. (2005). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nature neuroscience*, 8(11), 1611-1618.
- 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.
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2), 312-325.
- Dunne, S., & O'Doherty, J. P. (2013). Insights from the application of computational neuroimaging to social neuroscience. *Current opinion in neurobiology*, 23(3), 387-392.
- Frolichs, K.M.M., Korn, C.W. (2021, April 16). PerLe v1.0 [Computer software]. Zenodo. <http://doi.org/10.5281/zenodo.4697286>
- Hayes-Roth, B. (1977). Evolution of cognitive structures and processes. *Psychological review*, 84(3), 260.
- Fareri, D. S., Chang, L. J., & Delgado, M. R. (2015). Computational substrates of social value in interpersonal collaboration. *Journal of Neuroscience*, 35(21), 8170-8180.

- Fiske, S. T., & Dyer, L. M. (1985). Structure and development of social schemata: Evidence from positive and negative transfer effects. *Journal of Personality and Social Psychology*, 48(4), 839.
- 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.
- Goldberg, L. R. (1990). An alternative" description of personality": the big-five factor structure. *Journal of personality and social psychology*, 59(6), 1216.
- Hackel, L. M., Doll, B. B., & Amodio, D. M. (2015). Instrumental learning of traits versus rewards: dissociable neural correlates and effects on choice. *Nature Neuroscience*, 18(9), 1233.
- Hartley, C. A., & Somerville, L. H. (2015). The neuroscience of adolescent decision-making. *Current opinion in behavioral sciences*, 5, 108-115.
- Hein, G., Silani, G., Preuschoff, K., Batson, C. D., & Singer, T. (2010). Neural responses to ingroup and outgroup members' suffering predict individual differences in costly helping. *Neuron*, 68(1), 149-160.
- Hill, M. R., Boorman, E. D., & Fried, I. (2016). Observational learning computations in neurons of the human anterior cingulate cortex. *Nature communications*, 7(1), 1-12.
- Joiner, J., Piva, M., Turrin, C., & Chang, S. W. (2017). Social learning through prediction error in the brain. *NPJ science of learning*, 2(1), 1-9.
- Jolly, E., & Chang, L. J. (2019). The flatland fallacy: Moving beyond low-dimensional thinking. *Topics in cognitive science*, 11(2), 433-454.
- Jones, R. M., Somerville, L. H., Li, J., Ruberry, E. J., Libby, V., Glover, G., ... & Casey, B. J. (2011). Behavioral and neural properties of social reinforcement learning. *Journal of Neuroscience*, 31(37), 13039-13045.
- Jussim, L., Nelson, T. E., Manis, M., & Soffin, S. (1995). Prejudice, stereotypes, and labeling effects: Sources of bias in person perception. *Journal of Personality and Social Psychology*, 68(2), 228.
- Kang, P., Burke, C. J., Tobler, P. N., & Hein, G. (2021). Why We Learn Less from Observing Outgroups. *Journal of Neuroscience*, 41(1), 144-152.
- King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., & Montague, P. R. (2005). Getting to know you: reputation and trust in a two-person economic exchange. *Science*, 308(5718), 78-83.
- Klein, S. B., Loftus, J., Trafton, J. G., & Fuhrman, R. W. (1992). Use of exemplars and abstractions in trait judgments: A model of trait knowledge about the self and others. *Journal of Personality and Social Psychology*, 63(5), 739.

- Korn, C. W., & Bach, D. R. (2019). Minimizing threat via heuristic and optimal policies recruits hippocampus and medial prefrontal cortex. *Nature human behaviour*, 3(7), 733-745.
- Korn, C. W., Prehn, K., Park, S. Q., Walter, H., & Heekeren, H. R. (2012). Positively biased processing of self-relevant social feedback. *Journal of Neuroscience*, 32(47), 16832-16844.
- Korn, C. W., Sharot, T., Walter, H., Heekeren, H. R., & Dolan, R. J. (2014). Depression is related to an absence of optimistically biased belief updating about future life events. *Psychological medicine*, 44(3), 579-592.
- Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2, 4.
- Lau, T., Gershman, S. J., & Cikara, M. (2020). Social structure learning in human anterior insula. *Elife*, 9, e53162.
- Lockwood, P. L., Apps, M. A., & Chang, S. W. (2020). Is there a ‘social’ brain? Implementations and algorithms. *Trends in Cognitive Sciences*.
- Lockwood, P. L., & Klein-Flügge, M. (2019). Computational modelling of social cognition and behaviour—a reinforcement learning primer. *Social Cognitive and Affective Neuroscience*.
- Mayer, J. D., & Bower, G. H. (1986). Learning and memory for personality prototypes. *Journal of Personality and Social Psychology*, 51(3), 473.
- Mende-Siedlecki, P. (2018). Changing our minds: the neural bases of dynamic impression updating. *Current opinion in psychology*, 24, 72-76.
- McCrae, R. R., & Costa Jr, P. T. (2008). The five-factor theory of personality.
- McCrae, R. R., Costa, P. T., de Lima, M. P., Simões, A., Ostendorf, F., Angleitner, A. & Chae, J. H. (1999). Age differences in personality across the adult life span: parallels in five cultures. *Developmental psychology*, 35(2), 466.
- Najar, A., Bonnet, E., Bahrami, B., & Palminteri, S. (2020). The actions of others act as a pseudo-reward to drive imitation in the context of social reinforcement learning. *PLoS biology*, 18(12), e3001028.
- Niv, Y., & Langdon, A. (2016). Reinforcement learning with Marr. *Current opinion in behavioral sciences*, 11, 67-73.
- Olsson, A., Knapska, E., & Lindström, B. (2020). The neural and computational systems of social learning. *Nature Reviews Neuroscience*, 21(4), 197-212.
- Ostendorf, F. (1990). *Sprache und Persönlichkeitsstruktur. Zur Validität des Fünf-Faktoren-Modells der Persönlichkeit*.

- Palminteri, S., Khamassi, M., Joffily, M., & Coricelli, G. (2015). Contextual modulation of value signals in reward and punishment learning. *Nature communications*, 6(1), 1-14.
- Rescorla, R. A., & Wagner, A. R., (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. *Classical conditioning II: Current research and theory*.
- Rice, G. E. (1980). On cultural schemata. *American ethnologist*, 7(1), 152-171.
- Rodriguez Buritica, J. M., Eppinger, B., Schuck, N. W., Heekeren, H. R., & Li, S. C. (2016). Electrophysiological correlates of observational learning in children. *Developmental science*, 19(5), 699-709.
- Rosenblau, G., Korn, C. W., & Pelphrey, K. A. (2018). A computational account of optimizing social predictions reveals that adolescents are conservative learners in social contexts. *Journal of Neuroscience*, 38(4), 974-988.
- Rosenblau, G., Korn, C. W., Dutton, A., Lee, D., & Pelphrey, K. A. (2020). Neurocognitive mechanisms of social inferences in typical and autistic adolescents. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*.
- Ruff, C. C., & Fehr, E. (2014). The neurobiology of rewards and values in social decision making. *Nature Reviews Neuroscience*, 15(8), 549-562.
- Sharot, T., Korn, C. W., & Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality. *Nature neuroscience*, 14(11), 1475-1479.
- Stanley, D. A. (2016). Getting to know you: general and specific neural computations for learning about people. *Social Cognitive and Affective Neuroscience*, 11(4), 525-536.
- Stolier, R. M., Hehman, E., & Freeman, J. B. (2020). Trait knowledge forms a common structure across social cognition. *Nature human behaviour*, 4(4), 361-371.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning* (Vol. 135). Cambridge: MIT press.
- Tamir, D. I., & Thornton, M. A. (2018). Modeling the predictive social mind. *Trends in cognitive sciences*, 22(3), 201-212.
- Thorndyke, P. W., & Hayes-Roth, B. (1979). The use of schemata in the acquisition and transfer of knowledge. *Cognitive Psychology*, 11(1), 82-106.
- Thornton, M. A., Weaverdyck, M. E., Mildner, J. N., & Tamir, D. I. (2019). People represent their own mental states more distinctly than those of others. *Nature communications*, 10(1), 1-9.

- Wiggins, J. S., & Trapnell, P. D. (1997). Personality structure: The return of the Big Five. In *Handbook of personality psychology* (pp. 737-765). Academic Press.
- Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *Elife*, 8, e49547.
- Zaki, J., Kallman, S., Wimmer, G. E., Ochsner, K., & Shohamy, D. (2016). Social cognition as reinforcement learning: feedback modulates emotion inference. *Journal of Cognitive Neuroscience*, 28(9), 1270-1282.
- Zhang, L., Lengersdorff, L., Mikus, N., Gläscher, J., & Lamm, C. (2020). Using reinforcement learning models in social neuroscience: frameworks, pitfalls and suggestions of best practices. *Social cognitive and affective neuroscience*, 15(6), 695-707.

Acknowledgments

We thank Antonia Wesseloh, Kevin Rozario, Cristian Ioan, and Alina Dinu for help with data acquisition and Jan Gläscher, Tessa Rusch, Lisa Doppelhofer, and Benjamin Kuper-Smith for helpful discussions. KMMF and CWK are supported by the German Research Foundation (DFG); specifically by an Emmy Noether Research Group [392443797]. GR is supported by the National Institute of Mental Health (NIMH, R01MH116252) and the Simons Foundation for Autism Research Bridge to Independence Award. The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions

All authors designed the experiments, developed the analysis procedures, and wrote the paper. KMMF and CWK collected and analysed the data.

1426 How do humans learn about other people?
1427 Incorporating social knowledge structures into
1428 computational models

1429 **Supplementary Materials**

Supplementary Figure 1. Highlighting the main parts of our modeling approach. All data is simulated and purely intended for visualization purposes.

A. Behavioural Data. An idealized learning response pattern over time. The black line represents to-be-learned values on a given task (e.g., self-ratings on personality traits by a certain person). The gray line represents a participant's answers trying to estimate the self-ratings. This participant is learning because over time her answers get closer to the outcome (self-ratings). Similarly, this is indicated by the decrease in the absolute prediction error (PE, transparent gray) over time. Using computational models we can attempt to capture such behavior. In the following sections we explain several models of differing complexity that all focus on different pieces of information regarding the self-ratings.

Building a model

In general, the black line represents self-ratings on personality traits. A real-life experiment would have randomly ordered items but we have ordered them based on five arbitrary factors for easier visualization, where each factor consists of multiple trials. Shaded (grey and white) backgrounds indicate the factor and coloured shaded regions in the lower parts of the plots indicate the PEs of the models with the same colour.

B. Factor Reference Points. A simple model could make use of only a Reference Point (RP) per factor (i.e., a person is represented as having a single mean value per factor). Added are two such RPs, one that fits the data reasonably well (green) and one that does not fit the data well (red). The shaded area at the bottom of the plot represents the absolute PE for both RPs. The worse fit of the red RP is indicated by the higher values of the absolute PE.

C. Individual Reference Points. A more complicated model could have RPs for each trait separately since this can capture idiosyncrasies within factors e.g., in Factors 3 and 4 these RPs capture that there is a group of higher and lower rated items. Something that is not present in the previous Factor RP model.

D. Reinforcement Learning. Standard Rescorla-Wagner models use the discrepancy between the expected and actual outcome values (PE) to incrementally update the estimate for the next time step, the magnitude of the updating is determined by the learning rate (usually denoted as: α). In this figure, an agent learns an average value for two factors i.e., individual items get presented that belong to one of two factors, the participant automatically categorizes these into their respective factor and learns an average value per factor. The horizontal coloured lines are the expected averages for the first (blue) and second (green) factor where shades indicate the time steps in learning. The first estimate for the first factor (blue) results in a PE of 3.4, after updating this estimate again the resulting PE is: 1.8, a third estimation results in a PE of: 1. Estimates in this figure are spaced out for visualization purposes, for the same reason the factors have been presented in a ordered fashion instead of interspersed.

E. Similarity Reinforcement Learning. Rescorla-Wagner models learn only one value at a time, this can be a trait or an average value (e.g., per factor). However, real-life learning happens broadly from few examples. A mechanism that can explain such generalized learning is to update all traits based on how correlated they are to the current trait.

In the top section, the black line represents a subsection of self-ratings on two factors. For visualization sake it is assumed that the participant expects each trait rating to be zero (grey dots on the x-axis). In this example the participant gets feedback on item 5 (black dot), which results in a PE of 5.11 (i.e., the difference between the expectation [0] and the outcome [5.11], the vertical black line between the two black dots). Instead of multiplying this by the learning rate and updating solely this item, the participant can update all traits based on their correlation to this current item. The similarity matrix on the right shows the correlations between items. Since the model is learning about the fifth item, the fifth row (could also be the fifth column) is selected to update all other items. On the bottom, this row has been enlarged and correlation coefficients have been added. Each item is then updated by multiplying the PE, the learning rate, and the correlation coefficient resulting in the new estimates (red dots). Item 5 is updated according to the standard RL. While only receiving feedback about item 5, updates occur for all other correlated items in the set. Items for which the correlation is (close to) zero are not simultaneously updated because learning about item 5 gives no or very little information about uncorrelated items.

F. Combining Factor RP & Similarity RL. Using a weighting parameter (γ) we can combine the previously discussed models into combination models. The shaded blue lines are several weighted combinations of the Factor RP and the Similarity RL. Notably, if the weighting parameter is either 0 or 1 it is simply one of the subcomponents (Factor RP or Similarity RL) of this combination model. In this example, equally weighting both subcomponents (γ : 0.5) most accurately follows the self-ratings and thus would be the best fitting model. Akin to the learning rate, the γ is usually a free parameter that is estimated during model fitting.

Supplementary Figure 2. Protected exceedance probabilities for all experiments.

The losing performing model is always chosen as the baseline. Bayesian model selection is performed using random-effects (RFX) analysis. Granularity (G) and Reference Points (RP).

- A) Consistent with FFX analysis, Model 5 [Fine Granularity & Population RP] was the best fitting model.
- B) In line with participant behaviour, RFX-analysis indicated Model 3 [Coarse Granularity & Population RP] to be the best model.
- C) Model comparison indicated that model 5 performed best according to FFX and RFX analyses.
- D) Models using a stereotypical RP have a darker colour than their “twins” with the standard RP. Model comparison evidence is low for Experiment 4, with RFX analysis being slightly ahead for model 8 [Fine Granularity and Stereotypical RP].
- E) Model comparison indicated Model 4 as the best model according to both FFX and RFX analyses.

Fit Model

A	1	2	3	4	5
1	96.5	1	1.5	0.5	0.5
2	0	99.5	0.5	0	0
3	0	0.5	99.5	0	0
4	0	0	0	99.5	0.5
5	0	0	0	1	99

B	1	2	3	4	5
1	99	0	0	0.5	0.5
2	0	100	0	0	0
3	0	0.5	99.5	0	0
4	0	0	0	99.5	0.5
5	0	0	0	1	99

Simulated Model

C	1	2	3	4	5
1	98	0	1	1	0
2	0	99.5	0.5	0	0
3	0	1	99	0	0
4	0	0	0	99	1
5	0	0	0	2	98

D	1	2	3	4	5
1	97.5	0.5	0.5	0.5	1
2	0	99	1	0	0
3	0	0.5	99.5	0	0
4	0	0	0	99.5	0.5
5	0	0	0	0	100

E	1	2	3	4	5
1	96.5	0	1.5	1.5	0.5
2	0	97.5	2.5	0	0
3	1	6.5	92	0	0.5
4	0	0	0	99	1
5	0.5	0	0	7	92.5

Models

1. No Learning
2. **C**oarse **G**ranularity
3. **CG** & Population RP
4. **F**ine **G**ranularity
5. **FG** & Population RP

Supplementary Figure 3. Confusion matrices for every experiment.

To test whether data created with a model would also be best fitted by that model, a confusion matrix was created for every experiment separately by randomly sampling parameter values 200 times. With these random parameter values, data was simulated and noise, sampled from a standard normal distribution, was added. After this, the models were fitted to the simulated data. Ideal recovery is reflected by the identity matrix (i.e., all values on the diagonal being 100%). Model recovery for all experiments was good, indicating that our models were recovered correctly. A-E represent experiments 1-5.

Slope/ Alpha

Simulated

Supplementary Figure 4. Parameter recovery for the first free parameter.

To check whether model fitting returns meaningful parameter values, we performed parameter recovery on all models and all experiments separately. Akin to the process of model recovery (Supplementary Figure 2), we simulated data using random parameter values, to which noise was added. After this, all models were fitted to this simulated data and parameters recovered. We measured recovery by calculating the Pearson's Correlation coefficient (ρ) where a value closer to 1 indicated better recovery. Models are ordered along the columns, every row is a separate experiment. Parameter recovery for all experiments and all models was well within acceptable bounds. Green boxes indicate the winning model for that specific experiment. For Experiment 4, we only added the standard 5 models, choosing to omit the stereotype models. This was done because the stereotype models are functionally the same when it comes to simulated data and thus would not add any information.

1518

1519

1520

1521

1522

1523

1524

1525

1526

Supplementary Figure 5. Parameter Recovery for the second free parameter.
See **Supplementary Figure 3**, for an explanation of how we conducted parameter recovery. The Gamma parameter was only used in models 3 [Coarse Granularity and Population Reference Point] and 5 [Fine Granularity and Population Reference Point]. Models are on the rows, experiments on the columns. Recovery for all models and all experiments is very good across the whole range of values. Green boxes indicate when a model was the best fit for that specific experiment.

Intercept/ Starting Value

Supplementary Figure 6. Parameter recovery for the third free parameter.

See **Supplementary Figure 3**, for an explanation of how we conducted parameter recovery. Parameter recovery for the starting value (Models 2-5) and intercept (Model 1) is acceptable for all models.

Every row is an experiment, green boxes indicate this was the winning model for that experiment.

Supplementary Figure 7. The rationale of the best performing models (in the set).

Data in this figure are selected from two participants (top and bottom row) from Experiment 2. For the axis, data have been sorted according to time in experiment and according to factor i.e., items have been sorted based on when they were encountered within the factors agreeableness and conscientiousness. The best performing models are encountered in **Figure 2C** and **Figure 5E-H**

The left panels (A & D) show the profile answers in grey (this is the feedback participants received) and participant answers are indicated in green. During regular model fitting, models are fitted on these participant answers (rightmost panels C & F), here we see that the top participant's answers are best explained by model 5 (orange) and the answers by the bottom participant by model 1 (blue).

The middle column (B & E) shows the 'best performing models', these are models fitted on profile answers (grey line) i.e., models are fitted as if they have full information about the profiles, like participants had. For both profiles model 5 is the best performing model in our set of models.

Comparing the best fitting models with the models that best explained participant data helps us in understanding whether participants used an effective strategy or not.

Supplementary Figure 8. Pairwise correlations of profiles and population averages for experiment 1 and 4. The correlation coefficients between profiles and the RP seem to be higher for experiment 4 than experiment 1. This additional analysis was performed to understand why the best performing model for experiment 4 (**Figure 5G**) indicated that the No Learning model was best. One potential explanation is that the fashion model profiles are more correlated than the profiles from other studies. Thus needing less information to do the task efficiently. Interestingly enough, participants still seemed to use a complex strategy [Model 5] during this task.

Supplementary Figure 9. Positive/ Negative models (P/N) compared with the standard models.

Positive/ Negative models (P/N) have two parameters for the learning rate (α), splitting it up for positive and negative PEs. This allowed us to check whether participants learned differently from positive and negative PEs.

A. The standard models tested for Experiments 1-5 corresponding to B-F. Experiment 4 (see E) has the standard models with the addition of the models in the grey boxes that make use of the stereotypical reference point. B-F. In general, the P/N models perform worse compared to the regular models (also see Figure 2 and Figure 5). The Model [Fine Granularity & stereotypical RP] with the split in P/N learning rates only outperforms the other models in Experiment 4 (see E).

1569

1570
1571
1572
1573

Supplementary Figure 10. Model comparison of all models (see **Figure 2** and **Figure 5**) with additional models that use the participants' self-ratings as reference point (RP). Models that use self-ratings as their RP are always a little worse than those that use the population average RP.

Supplementary Methods

Optimal Parameter Space

In combination with the best performing model (**Figure 2C** and **Figure 5E-H** and see **Supplementary Figure 6** for a detailed explanation), we performed an extra analysis in an attempt to visualize the results of the model fitting procedure in more detail.

First we simulated all models and their respective fits over the whole parameter space on simulated profiles. Since our models had a relatively limited parameter space (i.e., maximum of 3 free parameters), we could explore their parameter space with relatively little computations. Each model's parameter space was mapped through a grid-search on all parameters, this was executed by creating a grid using 100 steps between the parameters bounds (i.e., [0 1] or [1 8]), this meant that the coordinates on the grid corresponded to the parameter settings. These parameter settings were run for every dataset separately using the best performing models (where the fit was determined using the sum of squared errors). Models 1, 2, and 4 have 2 free parameters and thus result in 2-dimensional plots, whereas models 3 and 5 have 3 free parameters. To display models 3 and 5 in a similar manner to the other models, we took the best fitting "slice" in the z-axis dimension (i.e., for the weighting parameter γ).

On this parameter space (where each axis represents distinct values on a parameter), we superimposed each participants' fitted parameters (grey diamonds for models that best fit participants behavior and white dots for the worse fitting models) to check if fitting procedures went correctly and how fitted parameters behaved for every model.

Finally, we hoped that visualizing the data as such we could easier catch a-priori unexpected ways that the models could fit the data (e.g., by ignoring one parameter completely).

These visualizations indicated that fitting procedures went well for all models (e.g., most fitted parameters fell within the optimal range) and all experiments (**Supplementary Figure 11-16**).

Supplementary Figure 11. Optimal parameter space with fitted parameters from participants in Experiment 1. There is a large overlap between optimal simulated SSE (red patch) and participants' best fitting parameters (grey dots).

Models 1, 2 & 3 (A, B & C) have two free parameters (*intercept* and *slope* for Model 1 and α and *start value* for Models 2 & 3). With only two parameters the whole parameter space can be displayed on a Cartesian plane. Models 3 & 5 (D & E) have three parameters which means their parameter space fills a "cube" (i.e., the three parameters result in three dimensional plots). In order to display these models in a similar manner to the other models, we chose a "slice" out of this cube where the best fit was achieved for the γ parameter (γ -value specified top left). In all plots, fitted parameter values per participant were plotted for the depicted model. White dots indicate participants for which the depicted model did not achieve the best fit. Grey diamonds indicate participants for which the depicted model did achieve the best fit out of all models tested.

Supplementary figure 12. Optimal Parameter Space for Experiment 2.

Same logic as Supplementary Figure 11, all parameters estimates fell within the optimal range for each model.

Supplementary figure 14. Optimal Parameter Space for Experiment 4.
Same logic as Supplementary Figure 11.

Supplementary figure 15. [Extension of Supplementary Figure 14] Optimal Parameter Space for Stereotypical Models of Experiment 4.

Experiment 4 had three additional (stereotypical) models that are displayed here in the spots of the models they were modelled after. That is, the models **A**, **B**, & **C** in this figure were functionally the same as models **A**, **C**, & **E** in **Supplementary Figure 14** with the exception of their reference points. Evidently from the plots, this change in RP did not have a major influence on the fitting procedure.

Supplementary figure 16. Optimal Parameter Space for Experiment 5.

Same logic as **Supplementary Figure 11**. Model fitting for experiment 5 resulted in lower SSE than the other experiments. This might have been caused by the smaller answer scale ([0 5] compared to [0 8]).

Item number	German trait word used	Approximate English translation
Factor 1: neuroticism – positive traits		
1	gelassen	composed
2	locker	easy-going
3	selbstständig	self-reliant
4	souverän	confident
Factor 1: neuroticism – negative traits		
5	ängstlich	anxious
6	feige	cowardly
7	launisch	moody
8	unentschlossen	indecisive
9	unsicher	insecure
10	wehleidig	whiny
Factor 2: extraversion – positive traits		
11	enthusiastisch	enthusiastic
12	lebenslustig	fun-loving
13	schlagfertig	articulate
14	tatkräftig	dynamic
Factor 2: extraversion – negative traits		
15	humorlos	humorless
16	kalt	cold-hearted
17	scheu	unassertive
18	unnahbar	inapproachable
Factor 3: openness to experience – positive traits		
19	kreativ	creative
20	offen	open-minded

21	spontan	spontaneous
22	tolerant	tolerant
23	vielseitig	versatile
24	wissbegierig	inquisitive
Factor 3: openness to experience – negative traits		
25	bieder	overly conservative
26	engstirnig	narrow-minded
27	träge	lazy
28	voreingenommen	biased
Factor 4: agreeableness – positive traits		
29	einfühlsam	empathetic
30	freundlich	friendly
31	gesellig	sociable
32	großzügig	generous
33	hilfsbereit	helpful
34	höflich	polite
35	respektvoll	Respectful
36	vertrauenswürdig	trustworthy
37	zuverlässig	reliable
Factor 4: agreeableness – negative traits		
38	aggressiv	aggressive
39	arrogant	arrogant
40	egoistisch	selfish
41	eitel	conceited
42	gehässig	spiteful
43	großmäulig	loud-mouthed
44	hinterhältig	conniving
45	rücksichtslos	inconsiderate
46	stur	stubborn

47	unsympathisch	unpleasant
Factor 5: conscientiousness – positive traits		
48	aufrichtig	honest
49	bescheiden	modest
50	diszipliniert	organized
51	effizient	efficient
52	fleißig	hard-working
53	kompetent	competent
54	ordentlich	tidy
Factor 5: conscientiousness – negative traits		
55	chaotisch	chaotic
56	inkonsequent	inconsistent
57	leichtsinnig	foolhardy
58	pedantisch	pedantic
59	unpünktlich	tardy
60	zwanghaft	obsessive

The 60 trait adjectives used here were a sub-selection of the 80 trait words used in previous studies (Korn et al., 2012, 2014). Half of the selected words were positive and half were negative (according to ratings in the earlier study; Korn et al., 2012). The trait adjectives were hand-annotated according to the Big-Five categories. These hand-annotations were compared to a larger list of German trait adjectives, which relied a slightly different factorization than the Big-Five (Ostendorf, 1990). We tried to use English translations consisting of one word. As is often the case for translations, a mixture of words could give a better idea of the word meaning (especially in the case of words with rather similar meaning).

1647 **Table S2. Stimuli used in Experiments 2 & 3**

Item number	German trait word used	Approximate English translation
Factor 4: agreeableness – positive traits		
1	altruistisch	altruistic
2	anständig	decent
3	aufrichtig	honest
4	begnügung	unpretentious
5	diplomatisch	diplomatic
6	ehrlich	candid
7	fair	fair
8	freigiebig	bountiful
9	generös	open-handed
10	gerecht	just
11	großmütig	magnanimous
12	großzügig	generous
13	hilfsbereit	helpful
14	kollegial	like a good colleague
15	konstruktiv	constructive
16	kooperativ	cooperative
17	loyal	loyal
18	rücksichtsvoll	considerate
19	selbstlos	selfless
20	solidarisch	solidary
21	spendabel	generous
22	treu	faithful
23	unbestechlich	incorruptible
24	uneigennützig	disinterested
25	unterstützend	supportive

26	verlässlich	dependable
27	verständnisvoll	understanding
28	vertrauenswürdig	trustworthy
29	verzeihend	forgiving
30	zuverlässig	reliable
Factor 5: conscientiousness – positive traits		
31	akkurat	accurate
32	anpackend	energetic
33	arbeitsam	industrious
34	beharrlich	persistent
35	diszipliniert	organized
36	effizient	efficient
37	eifrig	keen
38	engagiert	engaged
39	entschlusskräftig	decisive
40	fleißig	hard-working
41	flink	swift
42	fokussiert	focused
43	geduldig	patient
44	genau	meticulous
45	gewissenhaft	conscientious
46	gründlich	thorough
47	konsequent	consistent
48	leistungsorientiert	achievement-oriented
49	ordentlich	tidy
50	ordnungsliebend	orderly
51	pflichtbewusst	dutiful
52	planvoll	tactical
53	pünktlich	punctual

54	sorgfältig	diligent
55	systematisch	systematic
56	tatkräftig	dynamic
57	tüchtig	strenuous
58	verantwortungsbewusst	responsible
59	verantwortungsvoll	responsible
60	zielstrebig	goal-oriented

1648 Fifteen of the 60 trait adjectives in this list were selected from the list of 80 trait words used in previous studies (Korn et al., 2012,
1649 2014). All selected words were positive. The trait adjectives were hand-annotated according to the Big-Five categories. We tried to
1650 use English translations consisting of one word. As is often the case for translations, a mixture of words could give a better idea of
1651 the word meaning (especially in the case of words with rather similar meaning).
1652

1653 **Table S3. Stimuli used in Experiment 5**

Item number	German sentence used	English translation
Factor 1: neuroticism / emotional stability		
1	Ich gerate leicht in Stress.	I get stressed out easily.
2	Ich bin die meiste Zeit entspannt.	I am relaxed most of the time.
3	Ich grüble über Dinge.	I worry about things.
4	Ich fühle mich selten deprimiert (bedrückt).	I seldom feel blue.
5	Ich bin leicht zu beunruhigen.	I am easily disturbed.
6	Ich bin leicht aus der Fassung zu bringen.	I get upset easily.
7	Ich wechsele oft meine Stimmung.	I change my mood a lot.
8	Ich habe häufig Stimmungsschwankungen.	I have frequent mood swings.
9	Ich bin leicht zu reizen.	I get irritated easily.
10	Ich fühle mich oft deprimiert (bedrückt).	I often feel blue.
Factor 2: extraversion		
11	Ich bringe eine Party in Schwung.	I am the life of the party.
12	Ich spreche nicht viel.	I don't talk a lot.
13	Ich fühle mich wohl unter Menschen.	I feel comfortable around people.
14	Ich halte mich im Hintergrund.	I keep in the background.
15	Ich beginne Gespräche.	I start conversations.
16	Ich habe wenig zu sagen.	I have little to say.
17	Ich spreche mit vielen verschiedenen Leuten auf Partys.	I talk to a lot of different people at parties.
18	Ich mag es nicht, Aufmerksamkeit auf mich zu ziehen.	I don't like to draw attention to myself.
19	Ich habe kein Problem im Zentrum der Aufmerksamkeit zu stehen.	I don't mind being the center of attention.
20	Ich bin still unter Fremden.	I am quiet around strangers.
Factor 3: openness to experience		
21	Ich besitze einen großen Wortschatz.	I have a rich vocabulary.

22	Ich habe Schwierigkeiten abstrakte Ideen zu verstehen.	I have difficulty understanding abstract ideas.
23	Ich habe eine lebhafte Vorstellungskraft.	I have a vivid imagination.
24	Ich bin an abstrakten Ideen nicht interessiert.	I am not interested in abstract ideas.
25	Ich habe ausgezeichnete Ideen.	I have excellent ideas.
26	Ich habe kein gutes Vorstellungsvermögen.	I do not have a good imagination.
27	Ich verstehe Dinge schnell.	I am quick to understand things.
28	Ich gebrauche schwierige Wörter.	I use difficult words.
29	Ich verbringe Zeit damit über Dinge nachzudenken.	I spend time reflecting on things.
30	Ich bin voller Ideen.	I am full of ideas.

Factor 4: agreeableness

31	Ich empfinde wenig für andere.	I feel little concern for others.
32	Ich bin interessiert an anderen Menschen.	I am interested in people.
33	Ich beleidige andere.	I insult people.
34	Ich fühle mit anderen Menschen mit.	I sympathize with others' feelings.
35	Ich bin nicht interessiert an den Problemen anderer Leute.	I am not interested in other people's problems.
36	Ich habe ein weiches Herz.	I have a soft heart.
37	Ich bin nicht wirklich interessiert an anderen.	I am not really interested in others.
38	Ich nehme mir Zeit für andere.	I take time out for others.
39	Ich fühle die Gefühle anderer.	I feel others' emotions.
40	Ich bringe Leute dazu, sich wohl zu fühlen.	I make people feel at ease.

Factor 5: conscientiousness

41	Ich bin immer vorbereitet.	I am always prepared.
42	Ich lasse meine Sachen herumliegen.	I leave my belongings around.
43	Ich achte auf Details.	I pay attention to details.
44	Ich verursache großes Durcheinander.	I make a mess of things.

45	Ich erledige lästige Routinearbeiten unmittelbar.	I get chores done right away.
46	Ich vergesse oft, Dinge an ihren Platz zurückzulegen.	I often forget to put things back in their proper place.
47	Ich mag Ordnung.	I like order.
48	Ich drücke mich vor meinen Pflichten.	I shirk my duties.
49	Ich folge einem Plan.	I follow a schedule.
50	Ich bin anspruchsvoll in meiner Arbeit.	I am exacting in my work.

1654 We used 50 items from the German translation of the IPIP (International Personality Item Pool, which consists of lexical Big-Five
1655 factor markers). In the task, sentences were changed to the third person singular (e.g., “she gets stressed out easily”).
1656