Learning About the Self: Motives for coherence and positivity constrain learning from self-relevant social feedback

Jacob Elder^a, Tyler Davis^b, & Brent L. Hughes^a

^aDepartment of Psychology, UC Riverside, Riverside, CA

^bDepartment of Psychological Sciences, Texas Tech University, Lubbock, TX

Correspondence to:

Brent L. Hughes

Department of Psychology

University of California Riverside

900 University Avenue

Riverside, CA 92521, U.S.A.

Email: bhughes@ucr.edu

MOTIVATED SELF-LEARNING

Abstract

People learn about themselves from social feedback, but desires for coherence and positivity constrain how feedback is incorporated into the self-concept. We develop a network-based model of the self-concept and embed it in a reinforcement learning framework to provide a computational account of how motivations shape self-learning from feedback. Participants ($N = 46$) received feedback while self-evaluating on traits drawn from a causal network of trait semantics. Network-defined communities were assigned different likelihoods of positive feedback. Participants learned from positive feedback but dismissed negative feedback, as reflected by asymmetries in computational parameters that represent the incorporation of positive versus negative outcomes. Furthermore, participants were constrained in how they incorporated feedback: self-evaluations changed less for traits more important to coherence of the network. We provide a computational explanation of how motives for coherence and positivity jointly constrain learning about the self from feedback that makes testable predictions for future clinical research.

Keywords: self-concept, reinforcement learning, network analysis, motivated cognition, positivity, coherence

MOTIVATED SELF-LEARNING

Statement of Relevance

Our self-concepts are not static structures but rather fluctuate from social feedback throughout our lifetimes. However, we are not mere reflections of feedback we receive; there are a number of constraints that affect how people update self-beliefs. We offer a new computational account of how people update self-concepts to jointly maintain positivity and coherence while learning from social feedback. Our approach formalizes the self-concept as a network of interconnected traits that depend upon one another, which constrains how people update individual trait knowledge to maintain global self-concept coherence. We embed this network in a reinforcement learning framework and show how positive information is assimilated into self-concepts more readily than negative information, and how learning depends upon network structure. Our mechanistic network model offers the first formal theory connecting positivity and coherence constraints in self-learning and provides testable predictions for future research on self-concept malleability across mental health fields.

MOTIVATED SELF-LEARNING

Our self-concepts are not static structures but rather fluctuate from social feedback throughout our lifetimes (Bem, 1972; Mead, 1934; Vygotsky, 1978). However, we are not mere reflections of feedback we receive—we desire to view ourselves positively (Alicke & Sedikides, 2009; Taylor & Brown, 1988), and to maintain coherent and consistent views of ourselves (Markus & Wurf, 1987; Swann et al., 2003). How people acquire and maintain self-evaluations from social feedback may parallel how people more generally learn from feedback. Specifically, research demonstrates that positivity biases (Lefebvre et al., 2017), self-relevant biases (Koban et al., 2017; Müller-Pinzler et al., 2019), and associations between feedback (Vaidya & Badre, 2020) can shape how people modify their behavior from feedback. We unite these findings by considering how feedback-based learning (i.e. Reinforcement Learning) algorithms may be similarly employed to learn about the multifaceted self-concept (see Lockwood et al., 2020). We develop a computational framework to describe how the dual drives of positivity and coherence constrain learning about the self-concept from social feedback.

Although there is extensive evidence that people strive for self-concept positivity and coherence, research often pits these drives as antagonistic frameworks as opposed to dual constraints on self-representation. On the one hand, people frequently engage in self-enhancement processes in response to social feedback (Dunning, 1999; Hepper & Sedikides, 2012; Vandellen et al., 2010). People preferentially seek positive feedback over negative feedback (Pyszczynski et al., 1985), overestimate the positive feedback they receive (Hepper et al., 2011; Rodman et al., 2017; Somerville et al., 2010), and are more likely to update self-views from positive than negative feedback (Korn et al., 2012). Likewise, people selectively forget or discard negative feedback (M. Anderson & Hanslmayr, 2014; Frey et al., 1986; Schröder-Abé et al., 2007), and overcompensate in response to negative feedback (Beer et al., 2013; Greenberg &

MOTIVATED SELF-LEARNING

Pyszczynski, 1985; Hughes & Beer, 2013). On the other hand, people are also motivated to maintain coherent self-concepts (Markus & Wurf, 1987), even at the expense of positivity (Swann et al., 2003). For example, individuals seek interpersonal feedback that is consistent with their self-views, even when it is unfavorable (Swann et al., 1992), and they resist inconsistent feedback (Swann & Hill, 1982). However, this area of research lacks a formal explanation of how motivations to maintain positive and coherent self-views dually constrain self-concept malleability.

We develop a computational framework for how people dynamically update self-views to maintain positivity and coherence. We formalize the self-concept as a causal network of traits embedded in an RL agent that learns asymmetrically from positive and negative outcomes. Using this framework, we are able to make predictions about how people update their beliefs about individual traits. From a causal network perspective, traits that are critical for maintaining network coherence are those that are central in the network’s structure (Sloman et al., 1998). Our causal trait network was developed by having an independent sample of participants provide semantic evaluations of the causal relationships between traits (Elder et al., 2020). For example, in this model, the trait “Outgoing” causes “Sociable” and “Fun”, and “Fun” in turn causes “Witty”. By extension, an individual’s self-views on “Outgoing” should thereby bear implications on their self-views on “Sociable” and “Fun”. This network model describes how people represent traits and their relationships, and how this structured representation shapes self-concept maintenance and malleability. The network model instantiates normative beliefs about the compatibility and causal dependence amongst traits, in contrast to how traits vary and cause each other in actual personalities. If people believe they are outgoing they ought to endorse traits that depend on being outgoing, and if they believe they are not outgoing they ought not endorse

MOTIVATED SELF-LEARNING

such traits. Coherence is achieved by maintaining consistency between self-perceptions for traits in the network, which suggests that highly central traits, that many traits depend on, are particularly critical for maintaining coherence.

Our causal network theory draws an important distinction between outdegree and indegree centrality. Outdegree centrality describes the number of traits that depend upon (or are outputs/children of) a given trait. A trait is more outdegree central if it has many traits that semantically depend on it (e.g. “Outgoing” causes “Sociable”). Indegree centrality describes the number of traits that cause (or are inputs/parents of) a given trait (e.g. “Witty” caused by “Fun”) or the number of other traits a trait depends on for its meaning. A trait that depends on many other traits will be more indegree central. Both types of centrality can be important for self-concept coherence, but likely relate to self-concept malleability in different ways. Higher outdegree traits are more important to self-concept coherence, as changes to higher outdegree traits would result in a cascade of changes to their dependent traits. Thus, higher outdegree traits should be more resistant to change. Higher indegree traits, in contrast, receive more inputs from other traits during feedback, and may thus be more malleable due to their sensitivity to many influences.

In order to test predictions from this computational framework, we employ a task that assigns social feedback according to network-defined communities of traits. Participants evaluated themselves on all network traits, and after each trait rating, received social feedback on each trait (ostensibly based on an interview). Each trait community received different likelihoods of positive feedback, which determined the feedback assigned to each trait. One simple test of network structure is whether participants learn how they are perceived on particular trait communities by discriminating the different levels of feedback they received.

MOTIVATED SELF-LEARNING

Our computational framework generates several predictions about how participants will maintain self-concept positivity and coherence. First, we expect that the motivation for positivity will manifest as a tendency to learn more from positive than negative feedback, such that network communities associated with higher probabilities of positive feedback will exhibit greater change in self-evaluations over the course of learning and from pre- to post-feedback. Computationally, this observation will be reflected in asymmetrical learning rates (RL parameters that describe how much participants incorporate feedback). We predict that the learning rate for positive feedback will be higher than for negative feedback. The tendency to asymmetrically process social feedback when learning about the self-concept may help people to maintain positive self-views. Thus, we tested whether personality measures generally reflecting positive self-views and psychological adjustment are associated with learning rates. Second, we expect that a node's outdegree and indegree centrality will impact how much people update self-evaluations based on feedback. Specifically, we predict that self-evaluations for higher outdegree traits will change less from feedback, and self-evaluations for higher indegree traits will change more from feedback.

Method

Participants

48 undergraduate student participants were recruited via the UC Riverside credit pool in compliance with approved UC Riverside Institutional Review Board protocols and provided informed consent for their participation. The convenience sample was 67.3% female, from ages 18 to 24 ($M_{\text{Age}} = 19.45$), and was 32.7% Asian, 20.4% Hispanic, 16.3% Mixed, 12.2% African-American, 10.2% Caucasian, and 8.2% Other. Two subjects were omitted due to not believing the experimental manipulation, and two participants' questionnaire measures were missing due

MOTIVATED SELF-LEARNING

to protocol administration errors. The target minimum sample size of 44 was determined by a prospective power analysis for our primary within-subjects manipulation, testing the average difference in self-evaluations between the five trait communities that received feedback. We assumed a moderately small effect of $\eta_p^2 = 0.03$ (effect size shown in related literature, e.g. Beer et al., 2013; Korn et al., 2012), with a correlation among repeated measures of 0.45 and power above 80%.¹ We collected data in weekly posted slots, until the minimum was reached. Once the minimum sample size was reached, data collection continued until the conclusion of the academic quarter in which it started.

We manually inspected the data for quality. We conducted further quality control by flagging any subjects for which their quantity of missing responses was greater than 15% or used the same response for over 70% of responses, at both at both initial self-evaluations and re-evaluations. No participants met these criteria, and thus no participants were flagged and excluded.

Trait Network Development

148 traits were selected from a list of 292 positive traits (N. Anderson, 1968; Hampson et al., 1987; Kirby & Gardner, 1972). The list was filtered to 148 by randomly drawing 150 traits across 10,000 simulations and identifying which traits achieved the greater consistency across five dimensions of normative data: desirability, prevalence, category breadth, observability, and interpersonal (see Table S1 for a list of all traits). Two traits were removed from the list of 150 for having a mean desirability below 4.0 (for more detail, see Elder et al., 2020).

On each trial, participants were presented with one trait word as a target trait. Participants were then presented with a list of the remaining 147 positive trait words and asked, “Which traits

¹ In the final dataset, initial self-evaluations demonstrated repeated measures correlations of 0.493 and re-evaluations demonstrated repeated measures correlations of 0.411.

MOTIVATED SELF-LEARNING

does [TARGET TRAIT] depend upon?”. Participants were able to nominate as many trait words as they believed were applicable. Participants completed a total of 10 trials with 10 randomly selected traits as the target trait. This procedure was executed for all 148 positive traits and 148 negative traits, and at the conclusion, 144 positive traits were presented 12 times as the target word while 4 positive traits were presented 13 times as the target trait across participants. An adjacency matrix of 148 rows by 148 columns for trait words was computed based on the number of causal relationships nominated by participants. For each nomination of dependency by a participant, the corresponding cell (i.e., edge) within the adjacency matrix increased by 1 (e.g., if a participant nominates “Outgoing” as causing “Nice”, then 1 is added to the row for “Nice” and the column for “Outgoing”). This adjacency matrix was then thresholded using an a priori threshold of 25% to remove edges with lower numbers of endorsements. Moreover, thresholding to remove edges with a lower number of endorsements should filter idiosyncrasies in dependency nominations that do not reflect general consensus of trait relationships (for more detail, see Elder et al., 2020).

Outdegree centrality was defined as the number of directed edges from a target word to other words (i.e. how many traits depend on a target trait). Indegree centrality was defined as the number of directed edges from other words to the target word (i.e. how many traits a target trait depends on). Additionally, we identified communities of traits in the network by using a walktrap community detection algorithm (Pons & Latapy, 2005). This algorithm uses the length of random walks from a node in a network to other nodes in order to identify densely interconnected groups of traits. We defined the communities used to assign feedback from an initial community detection procedure that detected five communities. However, later analyses uncovered a coding error that excluded 1 trait word, that when corrected, led the same algorithm

MOTIVATED SELF-LEARNING

to detect four total communities. Given the role of the original community construction in the experimental design, we retained the original solution for all analyses and modeling. All 148 traits from the positive trait network were incorporated into the subsequent experimental protocol (see Figure 1 for network visualization).

Experimental Procedure

Participants first completed an initial lab visit in which they were administered a consent form and received the cover story for the experiment. They were informed that they would be undergoing an interview that would be recorded and shared with three to five members of UC Riverside undergraduate admissions committee. During the interview, participants were asked a range of questions about their personal characteristics, goals, and interests (see Table S2 for interview questions), and interviews lasted approximately 10-20 minutes. Following the interview, participants completed several questionnaires and were scheduled to return to the lab for the second appointment to complete the social evaluative task (approximately 7 to 10 days later).

Participants returned to the lab for the second visit to complete a social evaluative task on a computer. Participants were led to believe that in the time between the first and second visit, 3-5 members of the UC Riverside admissions committee had reviewed the participant's video interview and evaluated the participant on all 148 trait words based on their interview. Participants were told that during the task, they would be asked to evaluate themselves on array of traits. However, while evaluating themselves, they will also observe how the committee members had evaluated them. It was not made explicit whether they should take these evaluations into consideration during their own self-judgments.

MOTIVATED SELF-LEARNING

At the conclusion of the second visit, participants were debriefed and informed that the feedback was administered via a pseudo-random algorithm and there were no committee members that had reviewed their interview. They also underwent a series of offboarding questions asked to assess their confusions, engagement, and belief in the authenticity of the manipulation. Following this exit interview, they were awarded a credit for completion and exited the lab.

Social Evaluative Feedback Task

The experimental task was programmed in MATLAB's Psychtoolbox and was similar to other social evaluative feedback paradigms (Eisenberger et al., 2011; Hughes & Beer, 2013; Korn et al., 2012; Somerville et al., 2010; Will et al., 2017). During the task, participants evaluated themselves on all 148 positive traits from the trait network on a 1 (not at all) to 7 (very much) scale in response to the prompt, "To what extent does the following trait describe you?". The number that participants selected as self-descriptive was framed in an orange square once the response was made. Participants were permitted to respond in as much time as they needed, and could opt out of self-evaluating if they were unfamiliar with a given trait word. After each self-evaluation, participants observed a brief inter-trial interval of a constant 0.20 seconds and were then presented with the ostensible social evaluative feedback. Feedback appeared with the prompt "The reviewers see you as..." with the trait at the center of the screen, and a red square framed around the assigned feedback. The orange score denoting the participant self-evaluation remained on the screen for the feedback phase (e.g. "+3" if reviewer feedback was 7 and participant self-evaluation was 4).

Feedback was administered via a pseudo-random algorithm. Five different probabilities of positive feedback— 90%, 70%, 50%, 30%, or 10%— were randomly assigned to each of the

MOTIVATED SELF-LEARNING

five trait network communities for each participant. For instance, a trait belonging to the 70% community would have a 30% probability of receiving negative feedback (i.e. on average, feedback is lower than the self-evaluation) and a 70% probability of receiving positive feedback (i.e. on average, feedback is higher than the self-evaluation). Given that all traits used in the study were positive traits, higher numerical feedback on the 1 to 7 scale was considered more positive. Once feedback was determined according to a given community's probability, the feedback number on the ordinal scale of 1 to 7 was assigned according to criteria related to the participant's response and the determined feedback valence (see Table S3 for feedback assignment rules). For example, if the presented trait belonging to the 70% community was assigned negative feedback due to a 30% probability, and the self-evaluation was a 4, the feedback would be randomly assigned from 1 to 3. If the trait was assigned positive feedback due to the 70% probability, feedback would be randomly assigned a number from 4 to 7. Therefore, feedback was contingent upon participant self-evaluations (see Table S4 for means and SD of feedback and other variables; see Figure 2 for task schemata).

Self-Report Measures

Rosenberg Self-Esteem Scale. A 10-item questionnaire that assesses individual differences in global personal self-esteem was administered (Rosenberg, 1965). The scale demonstrated excellent reliability in current sample ($\omega = .93$).

Self-Concept Clarity Scale. A 12-item questionnaire was administered that assesses individual differences in how clearly defined, internally consistent, and stable beliefs about the self are (Campbell et al., 1996). The scale demonstrated good reliability in the current sample ($\omega = .87$).

MOTIVATED SELF-LEARNING

Multidimensional Assessment of Interoceptive Awareness Scale Version 2. A 37-item questionnaire was administered that assess individual differences in conscious levels of interoception, “the process by which the nervous system senses, interprets, and integrates signals originating from within the body,” (Mehling et al., 2018). The scale demonstrated excellent reliability in the current sample ($\omega = .91$).

Dialectical Self-Views Scale. A 32-item questionnaire was administered that assesses individual differences in tolerance for contradictory views about oneself, maintenance of both positive and negative self-views, the expectation for self-views to change, and cognitive holism (Spencer-Rodgers et al., 2015). The scale demonstrated good reliability in the current sample ($\omega = .82$).

Singelis Self-Construct Scale. A 30-item scale was used to assess individual differences in independent (15-item subscale) and interdependent (15-item subscale) self-construct styles (Singelis, 1994). The independent and interdependent subscales exhibited good reliabilities respectively ($\omega = .83$ for both).

Beck Depression Inventory. A 21-item scale was used to assess individual differences in characteristic attitudes and symptoms of depression (Beck et al., 1996). The scale exhibited excellent reliability ($\omega = .93$).

Liebowitz Social Anxiety Scale. A 24-item scale was administered to assess individual differences in social anxiety related symptomatology (Liebowitz, 1987). The scale exhibited excellent reliability ($\omega = .96$).

Self-Monitoring Scale. A 18-item scale was administered to assess individual differences in concerns and ability to modify behaviors to meet the perceptions of others (Snyder, 1974). The scale exhibited acceptable reliability ($\alpha = .78$).

MOTIVATED SELF-LEARNING

Rejection Sensitivity Scale. A 18-item scale was administered to assess individual differences in sensitivity to perceived or actual social rejection (Downey & Feldman, 1996). The scale exhibited good reliability ($\omega = .83$).

Big Five Inventory–2 Extra-Short Form (BFI-2-XS). A 15-item short scale of the Big Five Inventory (Soto & John, 2017) was administered to assess individual differences in the traits factors negative emotionality/neuroticism, openness, agreeableness, extraversion, and conscientiousness. Openness demonstrated poor reliability ($\omega = .58$), neuroticism demonstrated relatively poor reliability ($\omega = .61$), conscientiousness demonstrated poor reliability ($\omega = .54$), extraversion demonstrated acceptable reliability ($\omega = .75$), and agreeableness demonstrated very poor reliability ($\omega = .32$). As noted by the creators, the low reliability is to be expected for brief scales that prioritize content validity over internal consistency.

Computational Models

In order to capture how individuals dynamically learn from social feedback, and how the motivation for positivity shapes this learning, we employed a reinforcement learning approach. We used a standard reinforcement learning model (Rescorla & Wagner, 1972) in which the trait communities were treated as Pavlovian cues, C , with *social expectations*, SE describing the model's estimate of expected rating for the current community. Weights (W_t) were defined as the learned associations for the five trait communities. Social expectations were operationalized as estimates of others' perceptions of them on a given community of related traits based on prior social feedback. The weight for the current trial's trait community is the current social expectation. Thus, the model describes how the expected rating (SE_{t+1}) of a trait community (C_t), as determined by the trait present on a given trial, is updated from its existing social expectation (SE_t), according to the error term (δ_t) and how much of the error term is incorporated by the

MOTIVATED SELF-LEARNING

learning rate (α). Weights (W_t) were initialized at the midpoint 4 to represent participants' neutral expectations regarding others' perceptions of them, consistent with the RL convention of initializing values between possible outcomes (Zhang et al., 2020). We tested a variety of different models and arrived at the best fitting model via model comparison (see Figure S1 for comparisons and Supplementary Text for model descriptions). Our final, best performing model included four free parameters: a negative prediction error learning rate (α_n ; bounded [0, 1]), a positive prediction error learning rate (α_p ; bounded [0, 1]), a (inverse) feedback sensitivity parameter (β ; bounded [.5, 9]), and a decay parameter (ϕ ; bounded [0, 1]).

At each trial, an individual's expectation (SE_t) corresponded to the weight (W_t) for the current trait's community (C_t):

$$SE_t = W_t(C)$$

Feedback (in whole numbers ranging from 1 to 7) is presented after the self-evaluation. Feedback is scaled by a feedback sensitivity parameter (inverse β), that accounts for range attenuation of responses at the natural boundaries of the Likert scale:

$$F_t = \frac{1}{1 + e^{\beta^* - (F_t - 4)}}$$

Feedback at the high (e.g. 6 or 7) or low (e.g. 1 or 2) ends of the scale have less range to modify participants SEs in positive or negative directions, respectively. A lower feedback sensitivity reflects that feedback outside the midpoint is enhanced and processed as more extreme than it is (e.g. 2 is processed as approximately 1), whereas a higher feedback sensitivity reflects that feedback outside the midpoint is blunted and processed as less extreme than it is (e.g. 1 is processed as approximately 2). In other words, individuals with a higher (inverse) feedback sensitivity parameter are more insensitive to differences in feedback such that feedback shrinks towards the midpoint whereas individuals with a lower feedback sensitivity parameter

MOTIVATED SELF-LEARNING

are more sensitive to differences in feedback such that feedback expands towards extremes (see Figure S2 for visualization of feedback sensitivity scaling feedback).

Feedback is centered at 0 for the sigmoid function (i.e. feedback sensitivity), and then rescaled back to the Likert scale after transformation. It is not possible for transformed feedback to be below 1 or above 7, and thus feedback and expectations will never exceed the boundaries of the Likert scale:

$$F_t = (F_t * 6) + 1$$

The error term consists of the difference between trial-level transformed feedback (F_t) on a given trait community (C_t) from the current expected rating (SE_t) of the community. The error term, also known as *prediction error (PE)*, is computed as:

$$\delta_t = F_t - SE_t$$

Due to separate learning rates for positive and negative prediction errors, reflecting asymmetrical learning, participants incorporate prediction errors differentially depending on whether it is positive or negative.

if $\delta_t > 0$:

$$SE_{t+1} = SE_t + \alpha_p * \delta_t$$

else if $\delta_t \leq 0$:

$$SE_{t+1} = SE_t + \alpha_n * \delta_t$$

At the beginning of the next trial, the weights are allowed to decay towards the midpoint, representing the forgetting of associations across trials:

$$W_t = W_t * \varphi + 4 * (1 - \varphi)$$

On each trial for which a trait is presented, the SE for the community that it belongs to is updated. For example, if “Sociable” is presented on a trial, its broader community of 47 traits is

MOTIVATED SELF-LEARNING

applied as C_t . If “sociable” received scaled feedback of 6 and the community (C_t) had an expectation (SE_t) of 5, this would result in an error (δ_t) of 1, while only .25 of which would be incorporated in updating the future expectation (SE_{t+1}) for C_t due to the .25 positive learning rate (α_p). In summary, the model describes how people vary in sensitivity to feedback depending on its position on the scale, and learn differently from positive and negative PEs, while also forgetting learned outcomes over time.

Model Fitting and Comparison.

We fit the RL model using a two-stage procedure that incorporated both individual and group information. First, model free parameters were fit using ordinary least squares to generate parameter estimates for each subject. Next, we pooled these parameters across participants and used their means as a prior to regularize participants’ individual parameters. Ridge penalized OLS estimation was then used to adjust individual parameter estimates based on the group priors. This regularization procedure helps to stabilize individual estimates, produces more accurate estimates of true parameters, and reduces the influence of outliers (Daw, 2011; Lockwood et al., 2016):

$$\sum_{i=1}^n (Evaluation - SE)^2 + \lambda * \sum_{k=1}^m Estimate^2$$

Ridge penalized OLS regularized estimates towards the estimate priors, adjusted by a constant penalty parameter (λ). We arrived at an appropriate penalty ($\lambda = .15$) by iterating through penalties from .05 to 1 at increments of .05, and determining which penalty produced parameter estimates that were most recoverable and predictive of behavior. Given that the parameters are on different scales, we z-scored the priors based on their OLS estimated means and standard deviations, such that regularization penalizes towards the mean centered at 0. Parameters were fit to each subject’s self-evaluations using the “L-BFGS-B” optimization

MOTIVATED SELF-LEARNING

algorithm from the `optimx` package, available in R (Nash & Varadhan, 2011). In order to compare models, we employed a formulation of Akaike Information Criterion (AIC) for residual sums of squares:

$$AIC = 2k + n * \ln \left(\frac{RSS}{n} \right)$$

The above formula describes n as the number of trials for participant i , RSS as the residual sums of squares for participant i , while penalizing k for the number of free parameters estimated in the model. A lower AIC value reflects a better performing model. AIC values were summed across subjects to estimate model performance.

Parameter Recovery

In order to determine whether parameters are identifiable (Lockwood & Klein-Flügge, 2020; Wilson & Collins, 2019; Zhang et al., 2020), we tested whether parameters could be recovered from simulated data. We performed three different tests of parameter recovery: (1) randomly simulating parameters, generating behavior from simulated parameters, and testing whether the simulated parameters could be recovered during fitting to generated behavioral data; (2) generating behavioral data using the original fitted parameters from the 46 participants, and testing whether the original participants' parameters could be recovered during fitting to generated behavioral data; (3) testing whether behavior generated from the recovered parameters paralleled the observed experimental behavior.

For each parameter from the best fitting model, we identified five equally spaced intervals between the 25th percentile and 75th of the parameters' distributions. At each interval, we drew 50 values for a given parameter and added Gaussian noise equivalent to 1 SD of the parameter to increase the range of possible parameters simulated for positive learning rate, negative learning rate, and decay. For feedback sensitivity, due to its non-normal distribution, we

MOTIVATED SELF-LEARNING

generated noise from a uniform distribution centered at each interval ranging from its negative standard deviation to its positive standard deviation. If the noise caused a parameter to reach the boundary of the parameter space, we set the parameter to its boundary. To simulate 250 participants (e.g. Palminteri et al., 2017), we randomly sampled without replacement from each of the newly generated parameters to determine a parameter set for a given participant. Using these generated parameter sets, we simulated behavior by generating feedback under the same algorithm used for the original task and according to the trial-by-trial social expectations determined by the parameters. We rounded trial-by-trial simulated estimates to the nearest whole number to emulate the Likert behavioral responses of the participants. Then, as with the original behavioral data, we fit parameters to the simulated behavioral data. We then correlated the fitted parameters with the “true” parameters generated from the simulations to estimate whether parameters were recoverable.

The second parameter recovery simulation was aimed at identifying how recoverable parameters were while maintaining the observed covariance structure of the original fitted subject parameters. This is in contrast to the previous recovery test which used simulated parameters that were randomly generated at intervals with noise and were then randomly shuffled into parameter sets, resulting in simulated parameters that are independent and uncorrelated, and thus do not share the covariance structure of the original participants’ fitted parameters (Vaidya & Badre, 2020). In order to account for this, participant fitted parameters were used to generate new behavioral data, and parameters were then fit to behavioral data generated by the original participant parameters. Correlations were estimated between the fitted parameters and the original participant parameters in order to estimate whether the original fitted parameters are recoverable, while accounting for the original covariance structure of parameters.

MOTIVATED SELF-LEARNING

As a final test, we assessed whether the recovered parameters paralleled the observed behavior in the task. Behavioral data was generated from the recovered model parameters and compared to the original behavioral data to determine whether simulated behavior from the recovered parameters can reproduce key elements of the observed experimental behavioral data. Specifically, we conducted a statistical test on the simulated data that paralleled a test conducted on original data, and examined whether it produced similar outcomes.

Analysis Plan

Mixed models were implemented in R using lme4 (Bates et al., 2015) and Satterthwaite's approximation was used for determining p-values in lmerTest (Kuznetsova et al., 2017) packages. Randomization tests were implemented using multicon. Marginal and conditional R^2 , and semipartial R^2 for linear mixed models were estimated using r2glmm (Edwards et al., 2008; Jaeger et al., 2017). Likelihood ratio tests were performed to determine models best supported by the data. All mixed models included crossed random effects (Baayen et al., 2008), modeling both traits and subjects as random factors. Maximal random effects were tested, and were removed as needed if unsupported by the data (i.e. low variance estimates) or if the model failed to converge.

Feedback-Based Differences in Self-Evaluations. We first conducted preliminary analyses to examine aggregate changes in self-evaluations as a function of feedback. A paired samples t-test was performed on change scores (the difference in re-evaluations from initial self-evaluations for each trait) from negative and positive feedback. A 5 x 2 repeated measures ANOVA was performed on self-evaluations with the feedback probability as a within-subjects independent variable containing five levels (90%, 70%, 50%, 30%, 10%) and time as a within-subjects independent variable containing two levels (initial self-evaluation and re-evaluations).

MOTIVATED SELF-LEARNING

Effect sizes were reported for both t-tests and ANOVAs. As a post hoc test for a linear effect of feedback probability on self-evaluations, we conducted two separate mixed models, for separate response variables containing initial self-evaluations and re-evaluations, with trials nested within subjects, and subjects and traits as random factors. In the models, orthogonal polynomial contrasts were applied to the five-level categorical factor consisting of each feedback probability, and the effect was allowed to vary between subjects as random slopes while slopes were set as fixed for traits. 1st order comparisons were examined to consider the linear relationship between the five feedback probabilities, and whether self-evaluations increase incrementally as the probability of positive feedback increases.

Computational Model Descriptives. An advantage of reinforcement learning models is that parameter fits can offer insights into individual differences in how participants learn from outcomes (Yechiam et al., 2005). We computed basic summary statistics (e.g. mean, median) to examine how participants varied on their parameter fits. A non-parametric paired samples permutation t-test was conducted (100,000 permutations), due to the non-normal distributions of learning rate parameters, to examine within-subjects differences between positive and negative learning rates. The null hypothesis is that the distributions of learning rates among positive and negative prediction errors are identical, while the alternative hypothesis is that the distribution of learning rates among positive prediction errors are systematically lower or higher than among negative prediction errors. For the two-tailed test, the p-value is estimated as the smaller proportion of permuted mean differences that are either lower or greater than the observed mean difference.

Trial-by-Trial Learning. We sought to test whether the RL model can effectively predict participant self-evaluations, from the prior feedback presented. In order to avoid

MOTIVATED SELF-LEARNING

overfitting, leave-one-subject-out cross-validation was employed: for subject i from sample N , subject i 's free parameters were omitted and the mean of free parameters, from 1 to $N - i$, was determined. Mean parameters determined by the leave-one-out procedure were included in the RL model for subject i , such that any predictability produced from RL model would not be a result of participant i 's data and overfitting, but rather from robustness of the model itself. We tested a mixed model that contained trials nested within subjects, with initial self-evaluation as the response variable, and random slopes for SE for subjects and fixed slopes for traits.

Additionally, we tested whether self-evaluations can be predicted by what has been learned about neighboring, connected traits. Given that the traits are learned within a network of interconnected traits, we used neighborhood link-based features to predict self-evaluations. For each trial, we computed values for the average PE for all neighbors of the current trial's trait. We then used traits' average neighboring PEs, by the trial the trait was observed, to predict self-evaluations. Therefore, this provides a dynamic model of how social feedback is learned by neighbors and within a network position. SE was retained in the model to determine if neighboring traits' PEs explain self-evaluations above and beyond the SEs of trait communities. In this approach, only traits' neighbors that had been observed prior to the current trait were included in the average. Thus, the trials prior to when a trait had been observed were treated as missing data. The effect of neighbors' average PEs were estimated as random slopes for subjects and fixed slopes for traits.

Analysis of Self-Evaluation Change. A residualized change approach (predicting re-evaluations while controlling for initial self-evaluations) was employed to test for changes in self-views. In order to incorporate SEs into the change analysis, the SE for the last trial of each of the five communities during the feedback phase were applied as the final learned SEs to

MOTIVATED SELF-LEARNING

predict re-evaluations (i.e. five final learned SEs per subject). We tested a mixed model with initial self-evaluations, outdegree centrality, and indegree centrality were entered as fixed slopes. PEs and SEs were entered as random slopes for subjects and fixed slopes for traits. Both outdegree and indegree centrality were tested as interactions with both PEs and SEs. Model comparison revealed that outdegree centrality was supported as an interaction with PE while indegree centrality was supported as an interaction with SE (see Table S9 for model comparisons).

Simulating Replications with Cross-Validation. While our study was well-powered for the focal within-subjects analyses, the participant sample was limited and therefore places constraints on generalizability. We thus employed cross-validation to simulate replications and strengthen generalizability claims (Koul et al., 2018; Yarkoni & Westfall, 2017). Data was randomly split into 10 folds. For each of the 10 folds, a fold was held out and the linear mixed model was fit to the data. Using the model trained on the nine folds, out-of-sample prediction was conducted on the remaining fold. Subsequently, cross-validation was repeated with different folds across 40 iterations. We determined out-of-sample predictive accuracy by extracting Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), marginal R^2 , and conditional R^2 for each model. As a further test of predictive accuracy, we also fit OLS regression models predicting target from prediction. Box and whisker plots were qualitatively examined for dispersion of model fit metrics across folds, while means were compared between models. Additionally, the mean marginal and conditional R^2 's for all folds for the final model were compared to the in-sample marginal and conditional R^2 to determine if coefficients tested on out-of-sample data performed comparably to the original model.

MOTIVATED SELF-LEARNING

Individual Differences and Computational Parameters. To examine how personality variables may relate to how an individual incorporates social feedback, all individual differences measures were correlated with the positive learning rate and negative learning rate. Due to the large number of individual differences and parameters to test, we employed randomization tests and adjusted correlational p-values using the Benjamini & Hochberg (1995) FDR procedure to control inflation of the false positive rate due to multiple comparisons. Due to the limited sample size and demographic range, the generalizability of between-subjects inferences should be approached with caution.

Results

Positive Self-Evaluations in Response to Aggregate Feedback

We first examined whether participants differ in self-evaluations due to broad feedback categories such as the valence of feedback and the probability of positive feedback by trait community. Preliminary analyses revealed that participants' self-evaluations varied as a function of feedback. A paired samples t-test on valence of feedback (positive vs. negative) revealed significant differences in change scores between the two types of feedback, such that participants change their self-evaluations more after positive ($M = 0.31$, $SD = 0.39$) than negative ($M = -0.07$, $SD = 0.49$) feedback from initial self-evaluations to re-evaluations, $t(45) = 6.60$, $p < .001$, $d = 0.973$. This finding broadly supports the hypothesis that people are positively biased in how they incorporate social feedback.

The 5 x 2 repeated measures ANOVA was tested for feedback probability-based differences in self-evaluations as well as changes in self-evaluations across time. The analysis revealed a main effect of feedback probability, $F(4,180) = 2.69$, $p = .033$, $\eta_p^2 = 0.06$, such that participants differentiated their self-evaluations according to the feedback probabilities of the

MOTIVATED SELF-LEARNING

trait community. Additionally, there was a main effect of time, $F(1,45) = 4.86$, $p = .033$, $\eta_p^2 = 0.10$, reflecting that self-evaluations changed from initial self-evaluations to re-evaluations. Finally, we also identified an interaction between time and feedback probability, $F(4,180) = 8.22$, $p < .001$, $\eta_p^2 = 0.16$, suggesting that self-evaluations were differentiated across time more for particular probabilities than others. Consistent with this, the 90% probability community is the most differentiated of all probability groups, reflecting that participants may be most sensitive to feedback if it is largely positive. Meanwhile, only the 10% probability group decreased from initial self-evaluation to re-evaluation, reflecting that participants are resistant to incorporate negative feedback on traits unless it is nearly entirely negative (see Figure 3 for visualization of change between probabilities and across time, Figure S3 for large visualization with raw data, and Figure S13A for raw data panel).

We conducted further tests to examine how these feedback probabilities were differentiated, as the means between feedback probabilities revealed step-wise increases with probability order. Indeed, a post hoc mixed model employing 1st order orthogonal polynomial comparisons demonstrated a linear effect such that initial self-evaluations were higher for traits assigned to higher probability groups ($\beta = .12$, $SE = .06$, $CI = [0.007, 0.230]$, $t(45) = 2.14$, $p = .038$, $sr^2 = 0.003$). While participants experienced the social feedback, they appeared to be sensitive to the different feedback probabilities, and began to align their self-evaluations in an order consistent with the feedback probabilities. This linear effect was stronger at re-evaluations, after all feedback had been learned. The post hoc mixed model testing for 1st order orthogonal polynomial comparisons revealed a strong linear effect between probabilities at re-evaluations ($\beta = 0.29$, $SE = 0.06$, $CI = [.165, .407]$, $t(44) = 4.75$, $p < .001$, $sr^2 = 0.02$). Tests reveal that self-evaluations were becoming differentiated by probability as early as initial self-evaluations, and

MOTIVATED SELF-LEARNING

that the differentiation became greater by the re-evaluation phase (see Table S4 for summary statistics; see Table S5 for ANOVA table). The interaction between time and feedback probability demonstrates that participants express a positivity bias in both their receptivity to positive feedback and their nearly entire dismissal of negative feedback, but it does not allow inferences regarding the underlying processes.

Asymmetries in Self-Related Learning

In order to assess how people learn from feedback, we fit RL models to each participant's responses. The best performing model contained separate learning rates for trials in which the prediction error was either positive or negative. A permutation-based paired t-test shows that fitted positive ($Mean = 0.50$, $Median = 0.42$, $SD = 0.41$) and negative ($Mean = 0.12$, $Median = 0$, $SD = 0.26$) learning rates per participant were significantly different from one another ($Observed\ Difference = .38$, $Mean\ Permuted\ Difference = .0003$, $p < .001$). The parameter fits reflect that people learn incrementally about positive outcomes related to the self, while neglecting negative outcomes related to the self, serving as further evidence of the prevalent motivation for positivity (see Figure 4 for scatterplot of positive and negative learning rates, Figure S4A for correlation matrix of parameters, and Figure S5 for histograms of parameters).

Social Expectations and Neighboring Prediction Errors Predict Self-Evaluations

As a proof-of-concept of the validity of the model, we evaluated how the model accounts for trial-by-trial self-evaluations and tested whether trial-by-trial community-based expectations (i.e. SEs) predicted self-evaluations, reflecting a global approach (i.e. groups of traits) to characterizing self-evaluations. Secondly, we leveraged the network structure to test a more local approach (i.e. neighbors of traits), by interpolating the average PEs of network neighbors. In the test of the model, community-based SEs significantly predicted trial-by-trial self-

MOTIVATED SELF-LEARNING

evaluations ($\beta = .09$, $SE = .02$, $CI = [.046, .130]$, $t(48) = 4.21$, $p < .001$, $sr^2 = .009$). Given that there is substantial variability ($\sigma_{\text{Subject}} = .17$, $\sigma_{\text{Trait}} = .10$) in trait self-evaluations, and this effect can only be observed insofar as participants can infer the trait structural relationships and the associated feedback, this effect size is fairly substantial. Findings suggest that as participants learn how they are perceived on different trait communities from feedback, they begin to dynamically shift their self-evaluations to match their learned expectations. However, they incorporate prior feedback into expectations in an asymmetrical manner, resulting in generally favorable expectations and positively-skewed updating of self-evaluations across learning (see Table S6 for model parameters, Figure S6 for large visualization of predicted effects with raw data, and Figure S13B for raw data panel).

In addition to the more global analysis examining how community SEs predict participant self-evaluations, we explored how PEs among local network connections (i.e. neighbors) may uncover additional information about participants' self-evaluations. We found that PEs among neighbors ($\beta = .09$, $SE = .02$, $CI = [.057, .118]$, $t(44) = 5.78$, $p < .001$, $sr^2 = .009$) strongly predicted self-evaluations. Furthermore, once these neighboring features were incorporated, the trial-by-trial effect of community SE remains a strong predictor of self-evaluations ($\beta = .07$, $SE = .02$, $CI = [.029, .111]$, $t(48) = 3.48$, $p = .001$, $sr^2 = .006$), suggesting that both neighbor- and community-level learned information contribute to learning and shape self-evaluations. Both the global, conventional approach and the local, averaged PEs of neighbors' approach offer useful tests of the RL model reflecting how social feedback may be learned at different levels of information-processing. Results indicate that participants use not only community-level information about a trait's expected feedback (i.e. SE) in forming self-evaluations, but also local information about how feedback from neighboring traits in the network deviates from

MOTIVATED SELF-LEARNING

expectation (i.e. PE). Self-evaluations may be updated to maintain consistency between neighbors and coherence among the broader network of interconnected traits. Cross-validation performed on the trial-by-trial learning mixed models provides evidence of out-of-sample generalizability and indicate that overfitting was avoided (see Table S7, Figure S7, and Supplementary Text for cross-validation details).

We additionally tested whether the feedback previously experienced and expectations formed during learning predict later self-evaluations after learning. PEs ($\beta = .13$, $SE = .02$, $CI = [.088, .177]$, $t(28) = 6.01$, $p < .001$, $sr^2 = .03$) and learned SEs ($\beta = .21$, $SE = .03$, $CI = [.148, .271]$, $t(30) = 6.90$, $p < .001$, $sr^2 = .06$) predicted participant re-evaluations, while controlling for initial self-evaluations. Therefore, participants used both the previously observed trait-specific PEs and their learned community-based SEs to inform their updated self-evaluations. Consistent with the analysis of trial-by-trial learning, changes in self-evaluations after learning are also influenced by expectations at the local (i.e. trait-level PEs) and global (i.e. community-level SEs) levels of the network (see Table S8 for main effects of residualized change model parameters).

Self-Evaluation Updates Depend on Network Features

Next, we tested how network structure imposes constraints on how people learn about themselves from social feedback and incorporate PEs and SEs. To this end, we examined how outdegree and indegree impacted the malleability of traits from their initial to final evaluations. The effect of PEs on trait re-evaluations interacted with outdegree centrality ($\beta = -.03$, $SE = .01$, $CI = [-.051, -.016]$, $t(6393) = -3.73$, $p < .001$, $sr^2 = .002$), such that self-evaluations for traits with greater outdegree centrality change less from PEs after learning from social feedback (see Figure 5 for visualization of outdegree centrality interaction, Figure S8 for large visualization with raw data, and Figure S13C for raw data panel). To the extent that a trait has more dependencies,

MOTIVATED SELF-LEARNING

participants update self-evaluations less based on PEs, reflecting one way that people maintain coherence in the self-concept. The degree to which a trait is essential to self-concept coherence (based on its causal relationships) may make it less susceptible to change. Further, the effect of SE on trait re-evaluations interacted with indegree centrality ($\beta = .01$, $SE = .01$, $CI = [.001, .037]$, $t(6388) = 2.12$, $p = .03$, $sr^2 = .001$), such that the RL model is better at predicting self-evaluations for higher indegree traits than lower indegree traits. This suggests that individuals may maintain coherence by modifying self-views on traits that are most susceptible to influence and most contingent on other self-views (see Figure S9 for visualization of indegree centrality interaction, Figure S10 for large visualization with raw data, and Figure S13D for raw data panel).

Furthermore, it supports the causal, directed structure of the network, as indegree and outdegree centrality demonstrate different interaction patterns (see Table S9 for centrality interaction model parameters and model comparisons). We again simulated replications with cross-validation, which supported the out-of-sample performance of the model, and suggests inadvertent overfitting was avoided (see Table S10, Figure S11, and Supplementary Text for cross-validation details).

Parameter Recovery

“True” parameters randomly generated from simulations and parameters fit to the simulated behavioral data were correlated with each other to estimate the extent to which parameters were recoverable. High correlations should reflect that behavioral data generated by a particular set of parameters will also produce parameter fits consistent with the behavioral data. Parameters, particularly the learning rate parameters, were determined to be recoverable, as fitted parameters were strongly correlated with simulated parameters: $\alpha_p = .95$, $\alpha_n = .91$, $\phi = .61$, and $\beta = .74$ (see Supplementary Figure S4B for correlation matrix).

MOTIVATED SELF-LEARNING

However, as this data was randomly generated, the previous test of parameter recovery assumes independence between parameters that is not observed in the original parameters. In order to more effectively address the recoverability of parameters that share the original covariance structure, we used the original participant parameters to simulate new behavioral data to which we fit the model parameters. We correlated the original parameters with the fitted (simulated) parameters to test recoverability. Again, we found that parameters were determined to be recoverable: $\alpha_p = .95$, $\alpha_n = .97$, $\phi = .91$, and $\beta = .62$ (see Supplementary Figure S4C for correlation matrix). Decay appears to benefit the most in terms of recoverability when considering the covariance structure and non-independence of parameters. Feedback sensitivity exhibits the lowest recoverability and may be a noisier parameter to estimate in general.

Lastly, we generated behavioral data from the recovered participant parameters in order to test the extent to which these parameters reproduce key elements of the experimental behavioral data. We conducted a time by feedback probability within-subjects ANOVA, identical to the ANOVA described above, on the simulated data. We identify a main effect of feedback probability $F(4,180) = 28.59$, $p < .001$, $\eta_p^2 = 0.98$, supporting that simulated behavior discriminates between feedback probabilities as the observed behavior does. We also identify a main effect of time, $F(1,45) = 12.45$, $p < .001$, $\eta_p^2 = 0.21$, reflecting that simulated behavior produces differences from initial self-evaluations to re-evaluations, paralleling the observed data. Unlike the original behavioral data, we do not observe a feedback probability by time interaction, $F(4,180) = 1.83$, $p = .12$, $\eta_p^2 = 0.04$, reflecting that the recovered parameters may not fully capture how responses change more for particular feedback probabilities than others *between* learning and re-evaluation. However, this is unsurprising as the model stops learning when it

MOTIVATED SELF-LEARNING

stops receiving feedback, and thus cannot, in its present form, account for any changes in self-evaluations prior to re-evaluation that do not occur in response to immediate feedback.

Altogether, model parameters are recoverable, and behavior produced by recovered parameters parallels observed participant behavior. It is important to note that our fitted and simulated RL agents only represent trait values at the community level and assume that these are uniform across the community (see Discussion), and thus our simulated data does not contain sub-community information (immediate neighbors, trait centrality) from the independently constructed network portion of the model.

Individual Differences in Learning from Positive Feedback

In order to test the extent to which asymmetrical learning of social feedback reflects individual differences, we correlated learning rate parameters estimated during model fitting with all individual differences collected through self-report. Randomization tested correlations with positive learning rates revealed an average observed correlational effect size ($r = .28$) greater than the expected average absolute correlational effect size under the null hypothesis ($r = .13$) and a quantity of significant associations (8) exceeding the expected number significant under the null hypothesis (.80). Negative learning rates revealed an average observed correlational effect size ($r = .15$) slightly greater than the expected average absolute correlational effect size under the null hypothesis ($r = .13$) and a quantity of significant associations (1) slightly exceeding the expected number significant under the null hypothesis (.78). Decay and feedback sensitivity parameter correlations were likely to be spurious and were omitted from the primary analyses (see Table S11 for all correlations).

After FDR-correction, self-concept clarity ($r = .40$, $p_{FDR} = .022$, $p = .007$), conscientiousness ($r = .40$, $p_{FDR} = .022$, $p = .007$), social anxiety ($r = -.51$, $p_{FDR} < .001$, $p <$

MOTIVATED SELF-LEARNING

.001), depressive symptoms ($r = -.41$, $p_{FDR} = .022$, $p = .006$), and rejection sensitivity ($r = -.54$, $p_{FDR} < .001$, $p < .001$) revealed significant associations. Extraversion ($r = .34$, $p_{FDR} = .065$, $p = .024$), self-esteem ($r = .32$, $p_{FDR} = .034$, $p = .033$) and neuroticism ($r = -.32$, $p_{FDR} = .073$, $p = .036$) revealed uncorrected associations $p < .05$ (positive learning rate correlations are displayed in Figure 6). After FDR-correction, negative learning rate revealed no significant associations, but uncorrected, depressive symptoms were significantly associated with negative learning rate ($r = .36$, $p_{FDR} = .247$, $p = .015$). Thus, individual differences reflective of psychological maladjustment and mental health dysfunction are negatively correlated with the positive learning rate, or modifying self-views in response to positive PEs. The negative learning rate correlation is the converse of this, where depressive symptoms are positively correlated with negative learning rate, or modifying self-views in response to negative PEs.

Discussion

We test a network model of the self-concept embedded within an RL framework in order to examine how people update self-representations based on social feedback, and how this updating process is constrained by drives for self-concept positivity and coherence. Consistent with the drive to achieve positive self-views, participants learned more from positive feedback than negative feedback, which was largely neglected. Consistent with the drive to maintain coherence, prediction errors were less likely to cause updates for higher outdegree traits. These results reveal how people simultaneously maintain positivity and coherence of self-representations, and shed light on the specific computational mechanisms underlying these processes. Together, the integration of network analysis and RL approaches conferred combined insights into how motivations shape learning about the self.

MOTIVATED SELF-LEARNING

We provide a mechanistic account of how people accomplish self-concept positivity. Past work has shown that people are positively biased in their self-evaluations (Hughes & Zaki, 2015; Sharot & Garrett, 2016), particularly in response to social feedback (Alicke & Sedikides, 2009; Hughes & Beer, 2013; Korn et al., 2012; Somerville et al., 2010). Here, we show that people accomplish this by learning asymmetrically from positive over negative prediction errors. This mirrors asymmetrical learning observed across other contexts (Dorfman et al., 2019; Gershman, 2015; Niv et al., 2012), including positivity biases in learning about future outcomes (Lefebvre et al., 2017). Conversely, other distinct but related research on self-learning suggests that people are negatively biased when learning about their own task performance (Müller-Pinzler et al., 2019). It may be adaptive to be positively biased in learning about specific trait self-beliefs (i.e. *Who am I?*), while also adaptive to be negatively biased in learning about task performance that can improve (i.e. *How did I perform?*). Consistent with a potential adaptive function underlying positively biased learning about trait self-beliefs, we observe a multitude of individual differences implicated in positive self-views and adjustment that are associated with positive learning rate, such as self-esteem, rejection sensitivity, depressive symptomatology, and social anxiety. This aligns with ongoing work characterizing the computational underpinnings of aberrant self-relevant learning in psychiatric disorders (Will et al., 2017, 2020). For example, depressed individuals learned more from negative feedback relative to healthy controls (Garrett et al., 2014), and socially anxious individuals exhibit greater sensitivity to negative feedback (Hopkins et al., 2021).

Structural features of the self-concept and the motivation for coherence impose additional constraints on its malleability and how people update self-views. The PEs of neighboring traits in the network predict self-evaluations for a given trait, providing evidence that structural

MOTIVATED SELF-LEARNING

influences at different levels of organization shape self-concept learning and coherence maintenance. Moreover, self-views for more outdegree central traits were less likely to change based on PEs. If outdegree central traits were modified, they may cascade to alter self-evaluations for other related traits and temporarily destabilize self-views. As such, traits with higher outdegree centrality may require more or stronger evidence in order to destabilize traits that are important to coherence. It may not simply be undirected similarity to other traits that is integral to how people update self-views, but instead coherence-preserving relationships that depend on causal structure to constrain malleability.

Although prior research has used RL to understand social feedback learning (Hackel & Amodio, 2018; Lockwood & Klein-Flügge, 2020), this was the first application of RL that examines how individuals learn from latent self-related feedback, based on semantic associations between traits. Prior studies have used explicitly labeled cues to observe how people update behavior according to reward payout. For example, recent work examined how participants update self-esteem ratings as a function of approval from social groups, where each group's approval tendency was denoted by different colors (Will et al., 2017). In contrast, participants in our study learned others' perceptions of them on traits from the relationships between previously observed traits. For this to occur, people must be sensitive to the structure of the network. Thus, self-evaluations are not influenced strictly by the feedback received for a particular trait, but also by updates for nearby traits in the same community and updates for neighboring, adjacent traits. This network framework thus enables new insights into social learning by examining how latent mental constructs (e.g., others' perceptions on groups of interconnected traits) are learned from social feedback within a structured concept space. For example, this approach could be extended

MOTIVATED SELF-LEARNING

to understand how people learn about social groups that emphasize different communities of traits.

Our novel combination of RL and network approaches presents a computational theory of self-concept malleability and coherence during social experience. We show that learning occurs from local network structure (e.g., neighbors) above and beyond the global network structure (e.g., communities) used for our experimental manipulation. Future models should incorporate finer-grained structural relations when examining self-related learning about trait values, by considering how feedback travels across a structure via causal associations (Vaidya & Badre, 2020; Wu et al., 2020). Additionally, future research with larger and more representative samples may examine mental health and well-being outcomes related to self-relevant learning to develop strategies that increase adjustment. Finally, we acknowledge that the self-concept is composed of more than traits words used in our model, and future work can expand the network's scope beyond traits to include other facets of the self-concept (e.g., social roles, group membership, specific behaviors). We believe this approach may allow a deeper understanding of how motivations interact in shaping self-perceptions, and the myriad conditions by which the self-concept may remain stable or change across time.

Open Sciences Statement

The study was not formally preregistered. However, de-identified data analysis, scripts, and materials are available at the Open Science Framework (OSF) at https://osf.io/j6bkc/?view_only=5766c02a668749e9b8a4d79355c274b6.

Acknowledgements

We thank Samuel Gershman (Harvard University) for his input on the computational modeling. We thank the research assistants, Leanne Esconde, Bryant Ma, James Sobrino, and

MOTIVATED SELF-LEARNING

Julia Hopkins, for their assistance with data collection. We thank Bernice Cheung (University of Oregon) for her contributions in piloting and developing the network.

References

- Alicke, M. D., & Sedikides, C. (2009). Self-enhancement and self-protection: What they are and what they do. *European Review of Social Psychology*, 20, 1–48.
<https://doi.org/10.1080/10463280802613866>
- Anderson, M. C., & Hanslmayr, S. (2014). Neural mechanisms of motivated forgetting. *Trends in Cognitive Sciences*, 18(6), 279–292. <https://doi.org/10.1016/j.tics.2014.03.002>
- Anderson, N. (1968). Likableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology*, 9(3), 272–279. <https://doi.org/10.1037/h0025907>
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48.
<https://doi.org/10.18637/jss.v067.i01>
- Beck, A. T., Steer, R. A., Ball, R., & Ranieri, W. (1996). Comparison of Beck Depression Inventories -IA and -II in psychiatric outpatients. *Journal of Personality Assessment*, 67(3), 588–597. https://doi.org/10.1207/s15327752jpa6703_13
- Beer, J. S., Chester, D. S., & Hughes, B. L. (2013). Social threat and cognitive load magnify self-enhancement and attenuate self-deprecation. *Journal of Experimental Social Psychology*, 49(4), 706–711. <https://doi.org/10.1016/j.jesp.2013.02.017>
- Bem, D. J. (1972). Self-Perception Theory. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 6, pp. 1–62). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60024-6](https://doi.org/10.1016/S0065-2601(08)60024-6)

MOTIVATED SELF-LEARNING

- Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- Campbell, J. D., Trapnell, P. D., Heine, S. J., Katz, I. M., Lavalley, L. F., & Lehman, D. R. (1996). Self-concept clarity: Measurement, personality correlates, and cultural boundaries. *Journal of Personality and Social Psychology*, 70(1), 141–156. <https://doi.org/10.1037/0022-3514.70.1.141>
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models. In *Decision making, affect, and learning: Attention and performance XXIII* (1st ed., Vol. 23). Oxford University Press Oxford.
- Dorfman, H. M., Bhui, R., Hughes, B. L., & Gershman, S. J. (2019). Causal Inference About Good and Bad Outcomes. *Psychological Science*, 30(4), 516–525. <https://doi.org/10.1177/0956797619828724>
- Downey, G., & Feldman, S. I. (1996). Implications of rejection sensitivity for intimate relationships. *Journal of Personality and Social Psychology*, 70(6), 1327–1343. <https://doi.org/10.1037/0022-3514.70.6.1327>
- Dunning, D. (1999). A newer look: Motivated social cognition and the schematic representation of social concepts. *Psychological Inquiry*, 10(1), 1–11. https://doi.org/10.1207/s15327965pli1001_1
- Edwards, L. J., Muller, K. E., Wolfinger, R. D., Qaqish, B. F., & Schabenberger, O. (2008). An R² statistic for fixed effects in the linear mixed model. *Statistics in Medicine*, 27(29), 6137–6157. <https://doi.org/10.1002/sim.3429>

MOTIVATED SELF-LEARNING

Eisenberger, N. I., Inagaki, T. K., Muscatell, K. A., Byrne Haltom, K. E., & Leary, M. R. (2011).

The neural sociometer: Brain mechanisms underlying state self-esteem. *Journal of Cognitive Neuroscience*, 23(11), 3448–3455. https://doi.org/10.1162/jocn_a_00027

Elder, J., Cheung, B., Davis, T., & Hughes, B. (2020). Mapping the Self: A network approach for understanding psychological and neural representations of self-concept structure.

PsyArxiv. <https://doi.org/10.31234/osf.io/hj87w>

Frey, D., Stahlberg, D., & Fries, A. (1986). Information seeking of high- and low-anxiety subjects after receiving positive and negative self-relevant feedback. *Journal of*

Personality, 54(4), 694–703. <https://doi.org/10.1111/j.1467-6494.1986.tb00420.x>

Garrett, N., Sharot, T., Faulkner, P., Korn, C. W., Roiser, J. P., & Dolan, R. J. (2014). Losing the rose tinted glasses: Neural substrates of unbiased belief updating in depression. *Frontiers*

in Human Neuroscience, 8. <https://doi.org/10.3389/fnhum.2014.00639>

Gershman, S. J. (2015). Do learning rates adapt to the distribution of rewards? *Psychonomic*

Bulletin & Review, 22(5), 1320–1327. <https://doi.org/10.3758/s13423-014-0790-3>

Greenberg, J., & Pyszczynski, T. (1985). Compensatory self-inflation: A response to the threat to self-regard of public failure. *Journal of Personality and Social Psychology*, 49(1), 273–

280. <https://doi.org/10.1037/0022-3514.49.1.273>

Hackel, L. M., & Amodio, D. M. (2018). Computational neuroscience approaches to social cognition. *Current Opinion in Psychology*, 24, 92–97.

<https://doi.org/10.1016/j.copsyc.2018.09.001>

Hampson, S. E., Goldberg, L. R., & John, O. P. (1987). Category-breadth and social-desirability values for 573 personality terms. *European Journal of Personality*, 1(4), 241–258.

<https://doi.org/10.1002/per.2410010405>

MOTIVATED SELF-LEARNING

- Hepper, E., Hart, C. M., Gregg, A. P., & Sedikides, C. (2011). Motivated Expectations of Positive Feedback in Social Interactions. *The Journal of Social Psychology, 151*(4), 455–477. <https://doi.org/10.1080/00224545.2010.503722>
- Hepper, E., & Sedikides, C. (2012). Self-enhancing feedback. In R. Sutton, M. Hornsey, & K. Douglas (Eds.), *Feedback: The communication of praise, criticism, and advice* (No. 4; 4th ed., Vol. 11, Issue 4, pp. 43–56). Peter Lang.
<http://www.peterlang.com/index.cfm?event=cmp.ccc.seitenstruktur.detailseiten&seitenty p=produkt&pk=67761&concordeid=310512>
- Hopkins, A. K., Dolan, R., Button, K. S., & Moutoussis, M. (2021). A Reduced Self-Positive Belief Underpins Greater Sensitivity to Negative Evaluation in Socially Anxious Individuals. *Computational Psychiatry, 5*(1), 21–37. <https://doi.org/10.5334/cpsy.57>
- Hughes, B. L., & Beer, J. S. (2013). Protecting the self: The effect of social-evaluative threat on neural representations of self. *Journal of Cognitive Neuroscience, 25*(4), 613–622.
https://doi.org/10.1162/jocn_a_00343
- Hughes, B. L., & Zaki, J. (2015). The neuroscience of motivated cognition. *Trends in Cognitive Sciences, 19*(2), 62–64. <https://doi.org/10.1016/j.tics.2014.12.006>
- Jaeger, B. C., Edwards, L. J., Das, K., & Sen, P. K. (2017). An R² statistic for fixed effects in the generalized linear mixed model. *Journal of Applied Statistics, 44*(6), 1086–1105.
<https://doi.org/10.1080/02664763.2016.1193725>
- Kirby, D. M., & Gardner, R. C. (1972). Ethnic stereotypes: Norms on 208 words typically used in their assessment. *Canadian Journal of Psychology/Revue Canadienne de Psychologie, 26*(2), 140–154. <https://doi.org/10.1037/h0082423>

MOTIVATED SELF-LEARNING

- Koban, L., Schneider, R., Ashar, Y. K., Andrews-Hanna, J. R., Landy, L., Moscovitch, D. A., Wager, T. D., & Arch, J. J. (2017). Social anxiety is characterized by biased learning about performance and the self. *Emotion (Washington, D.C.)*, 17(8), 1144–1155.
<https://doi.org/10.1037/emo0000296>
- 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. <https://doi.org/10.1523/JNEUROSCI.3016-12.2012>
- Koul, A., Becchio, C., & Cavallo, A. (2018). Cross-Validation Approaches for Replicability in Psychology. *Frontiers in Psychology*, 9. <https://doi.org/10.3389/fpsyg.2018.01117>
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82(1), 1–26.
<https://doi.org/10.18637/jss.v082.i13>
- Lefebvre, G., Lebreton, M., Meyniel, F., Bourgeois-Gironde, S., & Palminteri, S. (2017). Behavioural and neural characterization of optimistic reinforcement learning. *Nature Human Behaviour*, 1(4), 1–9. <https://doi.org/10.1038/s41562-017-0067>
- Liebowitz, M. R. (1987). Social phobia. *Modern Problems of Pharmacopsychiatry*, 22, 141–173.
<https://doi.org/10.1159/000414022>
- Lockwood, P. L., Apps, M. A. J., & Chang, S. W. C. (2020). Is There a ‘Social’ Brain? Implementations and Algorithms. *Trends in Cognitive Sciences*, 24(10), 802–813.
<https://doi.org/10.1016/j.tics.2020.06.011>
- Lockwood, P. L., Apps, M. A. J., Valton, V., Viding, E., & Roiser, J. P. (2016). Neurocomputational mechanisms of prosocial learning and links to empathy. *Proceedings of the National Academy of Sciences*, 113(35), 9763–9768.

MOTIVATED SELF-LEARNING

- Lockwood, P. L., & Klein-Flügge, M. C. (2020). Computational modelling of social cognition and behaviour—A reinforcement learning primer. *Social Cognitive and Affective Neuroscience*, *nsaa040*. <https://doi.org/10.1093/scan/nsaa040>
- Markus, H., & Wurf, E. (1987). The dynamic self-concept: A social psychological perspective. *Annual Review of Psychology*, *38*, 299–337.
<https://doi.org/10.1146/annurev.ps.38.020187.001503>
- Mead, G. H. (1934). *Mind, Self, and Society from the Standpoint of a Social Behaviorist*. Chicago.
- Mehling, W. E., Acree, M., Stewart, A., Silas, J., & Jones, A. (2018). The Multidimensional Assessment of Interoceptive Awareness, Version 2 (MAIA-2). *PLoS ONE*, *13*(12).
<https://doi.org/10.1371/journal.pone.0208034>
- Müller-Pinzler, L., Czekalla, N., Mayer, A. V., Stolz, D. S., Gazzola, V., Keysers, C., Paulus, F. M., & Krach, S. (2019). Negativity-bias in forming beliefs about own abilities. *Scientific Reports*, *9*(1), 1–15. <https://doi.org/10.1038/s41598-019-50821-w>
- Nash, J. C., & Varadhan, R. (2011). Unifying Optimization Algorithms to Aid Software System Users: Optimx for R. *Journal of Statistical Software*, *43*(1), 1–14.
<https://doi.org/10.18637/jss.v043.i09>
- Niv, Y., Edlund, J. A., Dayan, P., & O’Doherty, J. P. (2012). Neural Prediction Errors Reveal a Risk-Sensitive Reinforcement-Learning Process in the Human Brain. *Journal of Neuroscience*, *32*(2), 551–562. <https://doi.org/10.1523/JNEUROSCI.5498-10.2012>
- Palminteri, S., Lefebvre, G., Kilford, E. J., & Blakemore, S.-J. (2017). Confirmation bias in human reinforcement learning: Evidence from counterfactual feedback processing. *PLOS Computational Biology*, *13*(8), e1005684. <https://doi.org/10.1371/journal.pcbi.1005684>

MOTIVATED SELF-LEARNING

- Pons, P., & Latapy, M. (2005). *Computing Communities in Large Networks Using Random Walks* (pInar Yolum, T. Güngör, F. Gürgen, & C. Özturan, Eds.; pp. 284–293). Springer.
https://doi.org/10.1007/11569596_31
- Pyszczynski, T., Greenberg, J., & LaPrelle, J. (1985). Social comparison after success and failure: Biased search for information consistent with a self-servicing conclusion. *Journal of Experimental Social Psychology*, 21(2), 195–211. [https://doi.org/10.1016/0022-1031\(85\)90015-0](https://doi.org/10.1016/0022-1031(85)90015-0)
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory* (Vol. 2, pp. 64–99). Appleton-Century-Crofts.
- Rodman, A. M., Powers, K. E., & Somerville, L. H. (2017). Development of self-protective biases in response to social evaluative feedback. *Proceedings of the National Academy of Sciences*, 114(50), 13158–13163. <https://doi.org/10.1073/pnas.1712398114>
- Rosenberg, M. (1965). Appendix D: Scales and Scores. In *Society and the Adolescent Self-Image* (pp. 305–320). Princeton University Press; JSTOR.
<https://doi.org/10.2307/j.ctt183pjhh.21>
- Schröder-Abé, M., Rudolph, A., Wiesner, A., & Schütz, A. (2007). Self-esteem discrepancies and defensive reactions to social feedback. *International Journal of Psychology*, 42(3), 174–183. <https://doi.org/10.1080/00207590601068134>
- Sharot, T., & Garrett, N. (2016). Forming Beliefs: Why Valence Matters. *Trends in Cognitive Sciences*, 20(1), 25–33. <https://doi.org/10.1016/j.tics.2015.11.002>

MOTIVATED SELF-LEARNING

Singelis, T. M. (1994). The Measurement of Independent and Interdependent Self-Construals.

Personality and Social Psychology Bulletin, 20(5), 580–591.

<https://doi.org/10.1177/0146167294205014>

Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature Centrality and Conceptual Coherence.

Cognitive Science, 22(2), 189–228. https://doi.org/10.1207/s15516709cog2202_2

Snyder, M. (1974). Self-monitoring of expressive behavior. *Journal of Personality and Social*

Psychology, 30(4), 526–537. <https://doi.org/10.1037/h0037039>

Somerville, L. H., Kelley, W. M., & Heatherton, T. F. (2010). Self-esteem modulates medial

prefrontal cortical responses to evaluative social feedback. *Cerebral Cortex (New York,*

N.Y.: 1991), 20(12), 3005–3013. <https://doi.org/10.1093/cercor/bhq049>

Soto, C. J., & John, O. P. (2017). The next Big Five Inventory (BFI-2): Developing and

assessing a hierarchical model with 15 facets to enhance bandwidth, fidelity, and

predictive power. *Journal of Personality and Social Psychology*, 113(1), 117–143.

<https://doi.org/10.1037/pspp0000096>

Spencer-Rodgers, J., Srivastava, S., Boucher, H. C., English, T., Paletz, S. B., & Peng, K.

(2015). The dialectical self scale. *Unpublished Manuscript*.

<https://doi.org/10.13140/RG.2.1.3695.5928>

Swann, W. B., & Hill, C. A. (1982). When our identities are mistaken: Reaffirming self-

conceptions through social interaction. *Journal of Personality and Social Psychology*,

43(1), 59–66. <https://doi.org/10.1037/0022-3514.43.1.59>

Swann, W. B., Rentfrow, P. J., & Guinn, J. S. (2003). Self-verification: The search for

coherence. In *Handbook of self and identity* (pp. 367–383). The Guilford Press.

MOTIVATED SELF-LEARNING

- Swann, W. B., Wenzlaff, R. M., & Tafarodi, R. W. (1992). Depression and the search for negative evaluations: More evidence of the role of self-verification strivings. *Journal of Abnormal Psychology, 101*(2), 314–317. <https://doi.org/10.1037/0021-843X.101.2.314>
- Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin, 103*(2), 193–210. <https://doi.org/10.1037/0033-2909.103.2.193>
- Vaidya, A. R., & Badre, D. (2020). Neural Systems for Memory-based Value Judgment and Decision-making. *Journal of Cognitive Neuroscience, 32*(10), 1896–1923. https://doi.org/10.1162/jocn_a_01595
- Vandellen, M. R., Campbell, W. K., Hoyle, R. H., & Bradfield, E. K. (2010). Compensating, Resisting, and Breaking: A Meta-Analytic Examination of Reactions to Self-Esteem Threat: *Personality and Social Psychology Review*. <https://doi.org/10.1177/1088868310372950>
- Vygotsky, L. (1978). *Mind in Society: Development of Higher Psychological Processes*. Harvard University Press. <https://doi.org/10.2307/j.ctvjf9vz4>
- Will, G.-J., Moutoussis, M., Womack, P. M., Bullmore, E. T., Goodyer, I. M., Fonagy, P., Jones, P. B., Rutledge, R. B., & Dolan, R. J. (2020). Neurocomputational mechanisms underpinning aberrant social learning in young adults with low self-esteem. *Translational Psychiatry, 10*(1), 1–14. <https://doi.org/10.1038/s41398-020-0702-4>
- Will, G.-J., Rutledge, R. B., Moutoussis, M., & Dolan, R. J. (2017). Neural and computational processes underlying dynamic changes in self-esteem. *ELife, 6*, e28098. <https://doi.org/10.7554/eLife.28098>

MOTIVATED SELF-LEARNING

Wilson, R. C., & Collins, A. G. (2019). Ten simple rules for the computational modeling of behavioral data. *ELife*, 8, e49547. <https://doi.org/10.7554/eLife.49547>

Wu, C. M., Schulz, E., & Gershman, S. J. (2020). Inference and Search on Graph-Structured Spaces. *Computational Brain & Behavior*. <https://doi.org/10.1007/s42113-020-00091-x>

Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science*, 12(6), 1100–1122. <https://doi.org/10.1177/1745691617693393>

Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005). Using Cognitive Models to Map Relations Between Neuropsychological Disorders and Human Decision-Making Deficits. *Psychological Science*, 16(12), 973–978. <https://doi.org/10.1111/j.1467-9280.2005.01646.x>

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. <https://doi.org/10.1093/scan/nsaa089>

MOTIVATED SELF-LEARNING

Figure 1. Network visualizations for full network and subset of network. Nodes sized according to outdegree centrality. (Top) Network visualization of 148 traits colored by community. (Bottom) Outdegree connections for a high outdegree centrality trait, “Friendly”.

MOTIVATED SELF-LEARNING

Figure 2. *Task schematic. During learning phase, participants observed trait and were allowed as much time as needed to self-evaluate. An orange square appeared around their selection. After their selection, the frame remained on-screen for the 0.20 second inter-stimulus interval (ISI). Valence of feedback was determined by the feedback probability that the trait's community was assigned. Meanwhile, the feedback response was determined according to set of rules based on participant response and feedback valence. Feedback appeared as a red square, with white text above the square denoting the discrepancy between the feedback and participant self-evaluation. After self-evaluating and receiving feedback on all 148 traits, participants proceeded to re-evaluation phase. They then observed and self-evaluated again on each trait.*

MOTIVATED SELF-LEARNING

Figure 3. Plot illustrating the differences between feedback probabilities at initial self-evaluations and re-evaluations, and how self-evaluations change within feedback probabilities from initial self-evaluations to re-evaluations. Self-evaluations are on y-axis, while on x-axis are timepoints representing initial self-evaluations while receiving feedback (T1) and re-evaluations after feedback (T2). See Figure S3 in Supplemental Materials for figure with raw data.

MOTIVATED SELF-LEARNING

Figure 4. Scatterplot of positive and negative learning rates for each participant. In the top left diagonal (blue) are participants with greater positive than negative learning rates (i.e. positively-biased), while in the bottom right diagonal (red) are participants with greater negative than positive learning rates (i.e. negatively-biased). Positive learning rates are on average greater than negative learning rates, while many negative learning rates are below 0.

Figure 5. Predicted values (i.e. estimated marginal effects) from residualized change re-evaluation models, while holding covariates constant, with confidence intervals of ± 1.96 standard error. Re-evaluations while controlling for initial self-evaluations on y-axis. Interaction demonstrating that prediction errors contributes less to self-evaluations for traits higher in outdegree centrality. Outdegree centrality on the x-axis. Continuous variable of prediction error is split into three levels for visualization of interaction. See Figure S8 in Supplemental Materials for figure plotted with raw data.

MOTIVATED SELF-LEARNING

Figure 6. *Positive learning rate correlations with individual differences. Pearson coefficient magnitude on x-axis. Individual differences on y-axis. Correlations displayed in ascending order. From left to right: Rejection Sensitivity, (Liebowitz) Social Anxiety, Depression (BDI), Neuroticism, Dialectical Self-Views, Interdependent Self-Construal, Openness, Social Status, Agreeableness, Interceptive Awareness, Independent Self-Construal, Self-Monitoring, (Rosenberg) Self-Esteem, Extraversion, Conscientiousness, Self-Concept Clarity. Asterisks represent FDR-corrected $p < .05$. See Figure S12 in Supplemental Materials for panel figure of scatterplots depicting associations with $p < .10$ uncorrected. See Table S11 for table of all correlations.*