

Context-Dependent Learning in Social Interaction: Trait Impressions Support Flexible Social Choices

Leor M. Hackel¹, Peter Mende-Siedlecki², Siri Loken³, and David M. Amodio^{3, 4}

¹ Department of Psychology, University of Southern California

² Department of Psychological and Brain Sciences, University of Delaware

³ Department of Psychology, New York University

⁴ Department of Psychology, University of Amsterdam

How do humans learn, through social interaction, whom to depend on in different situations? We compared the extent to which inferred trait attributes—as opposed to learned reward associations previously examined as part of feedback-based learning—could adaptively inform cross-context social decision-making. In four experiments, participants completed a novel task in which they chose to “hire” other players to solve math and verbal questions for money. These players varied in their trait-level competence across these contexts and, independently, in the monetary rewards they offered to participants across contexts. Results revealed that participants chose partners primarily based on context-specific traits, as opposed to either global trait impressions or material rewards. When making choices in novel contexts—including determining who to choose for social and emotional support—participants generalized trait knowledge from past contexts that required similar traits. Reward-based learning, by contrast, demonstrated significantly weaker context-sensitivity and generalization. These findings suggest that people form context-dependent trait impressions from interactive feedback and use this knowledge to make flexible social decisions. These results support a novel theoretical account of how interaction-based social learning can support context-specific impression formation and adaptive decision-making.

Keywords: social cognition, learning, reward, context, generalization

Supplemental materials: <https://doi.org/10.1037/pspa0000296.supp>

Human social behavior is rooted in direct interpersonal interaction, and the feedback we receive from others through interaction provides a basis for the impressions we form of them. Yet, interactions with a person often take place across a variety of contexts, and the way a person acts in one context can differ dramatically from how they act in another. For instance, a friend who makes great conversation at a party may struggle to listen supportively after a bad day, and a colleague who offers excellent math advice may suggest poor advice about a social conflict. To effectively navigate social relationships, people must learn to choose the right partner in the appropriate context—for instance, learning whom to seek for social support as opposed to financial advice. In addition, people must be able to generalize this learning to novel situations, since they cannot

directly experience every scenario in which they might interact with a person. This ability is key to survival in complex human societies: context-sensitive social decisions afford greater access to resources embedded in social networks (Bendtsen et al., 2016; Cornwell & Cornwell, 2008; Maisel & Gable, 2009; Morelli et al., 2015, 2017; Shrout et al., 2006).

How do humans accomplish this complex form of interactive social learning, which involves learning from feedback from other humans across different contexts? Although much is known about nonsocial forms of feedback-based (i.e., reinforcement) learning across contexts (see Gershman, 2017, for review), the question of how people learn about social partners in different contexts through trial-and-error feedback in interactions remains unstudied. Here, we

Leor M. Hackel  <https://orcid.org/0000-0002-7959-0642>

Peter Mende-Siedlecki  <https://orcid.org/0000-0001-5761-0182>

Siri Loken  <https://orcid.org/0000-0003-1697-8815>

David M. Amodio  <https://orcid.org/0000-0001-7746-0150>

De-identified data for the experiments reported in this manuscript have been made available at: https://osf.io/496rn/?view_only=1d5f069aa26b421f9838e7328ba8f6a6.

Leor M. Hackel played a lead role in formal analysis and equal role in conceptualization, project administration, methodology, writing of original draft, and writing of review and editing. Peter Mende-Siedlecki played an

equal role in conceptualization, methodology, writing of original draft, and writing of review and editing. Siri Loken played supporting role in writing of original draft and equal role in project administration. David M. Amodio played an equal role in conceptualization, methodology, writing of original draft, and writing of review and editing.

Correspondence concerning this article should be addressed to Leor M. Hackel, Department of Psychology, University of Southern California, 3620 South McClintock Avenue, Los Angeles, CA 90089, United States or David M. Amodio, Department of Psychology, New York University, 6 Washington Pl, New York, NY 10003, United States. Email: lhackel@usc.edu or david.amodio@nyu.edu

integrate approaches from social cognition and reinforcement learning to investigate how people form impressions through interaction that can guide context-specific social choices.

Reinforcement Learning in Social Contexts

In social interactions, we learn about others through their feedback on our actions—a trial-and-error process that characterizes reinforcement learning. While reinforcement is critical to interactive learning, it has been studied primarily in nonsocial contexts (Daw et al., 2011; Sutton & Barto, 1998; Thorndike, 1911). In reward-based reinforcement learning, people learn to associate actions with rewarding outcomes and to repeat rewarded behaviors. (Similarly, in the context of negative reinforcement, people learn to associate actions with punishment and to avoid punished behaviors.) Recent research has examined this form of learning in social contexts and has shown that rewards similarly reinforce behavior in social interactions (Bhanji & Delgado, 2014; FeldmanHall et al., 2017; Hackel et al., 2015; Lindström et al., 2014).

How might reward learning support context-dependent choices? As an example, one might learn to associate another person with rewarding outcomes in one domain (e.g., help fixing a computer) but not in another (e.g., help editing a manuscript). These rewarding or nonrewarding outcomes may reinforce one's choice of interaction partner in a given context, such that the same partner is chosen again in that context in the future. Indeed, the ability to make context- (or state) dependent choices is central to formal models of reinforcement learning (Dayan & Niv, 2008; Gershman et al., 2015; Sutton & Barto, 1998) and has been found across people and animals (Bouton & Todd, 2014; Schuck et al., 2016; Trask & Bouton, 2014).

By receiving reward feedback, people could thus learn to approach another person in some contexts and avoid that person in different contexts. Moreover, people and animals often generalize learning about one stimulus to similar novel stimuli (Honig & Urcioli, 1981; Shepard, 1987; Soto et al., 2014). If people generalize learning from reward feedback, they might choose an interaction partner in a novel context (e.g., advice on writing an email) based on the rewards that partner offered in a similar context (e.g., advice on writing a manuscript).

However, while reward reinforcement provides a mechanism for interaction-based learning, its utility for cross-context social learning may be limited: because reward value is encoded as a generic, content-less metric (i.e., a common scale of value for different kinds of goods; Levy & Glimcher, 2012), it is unlikely to capture the nuance required for the flexible choice of social partners who may vary in their behavior across context.

Trait-Based Learning

Social psychological research on impression formation has traditionally focused on the learning of trait impressions, in contrast to the focus on reward value in reinforcement learning. Traits offer a highly nuanced and variegated characterization of an individual that affords much more precise predictions about a person's behavior in specific situations. Although major theories of trait impressions tend to focus on global, context-spanning inferences (e.g., Fiske et al., 2007; Olivola et al., 2014; Rosenberg et al., 1968; Tavares et al., 2015; Wojciszke, 2005), there is also substantial evidence for context-dependent impressions (Friesen & Kammrath, 2011;

Gawronski & Cesario, 2013; Kammrath et al., 2005; Plaks & Higgins, 2000; Rydell et al., 2009; Shoda et al., 1993; Shoda & Mischel, 1993). Notably, traits can also imply the availability of rewards; for instance, a kind individual is likely to provide us with rewarding experiences. However, traits differ from material rewards in that they provide abstract concepts that let us predict different kinds of rewards across different situations; a kind individual may provide us with a compliment or a gift, but they won't necessarily provide us with rewarding stock advice.

Hence, in contrast to reward value, trait inferences appear to provide the nuance and situational specificity needed for adaptive contextualized learning and decision-making. A limitation of the trait approach, however, is that existing models of trait inference focus almost exclusively on passive, noninteractive forms of learning, such as those based on behavioral observation, behavioral descriptions, or conceptual associations (Asch, 1946; Hastie, 1980; Heider, 1958; Jones & Nisbett, 1987; Kelley, 1967; Rosenberg et al., 1968; Uleman & Kressel, 2013; Winter & Uleman, 1984). These forms of trait inference rely on concept-based learning processes, such as semantic learning of beliefs and associations, which function very differently from the instrumental learning mechanism involved in interactive learning (Amodio, 2019; Amodio & Berg, 2018).

Despite the traditional focus on passive modes of trait inference, recent research has begun to show that trait impressions can also be formed through direct social interactions via mechanisms of reinforcement learning (e.g., Boorman et al., 2013; Hackel et al., 2015, 2020). In a study by Hackel et al. (2015), participants interacted with players in a money sharing game with the goal to learn about the players and their reward value. These players shared amounts that were either large or small in absolute terms while, independently, representing a large or small proportion of the player's total endowment. With this design, participants could learn from the reward value associated with a player, based on positive or negative outcomes, or, independently, the player's generosity—a trait-like attribute that involves the inference of a person's disposition from their behavior (Heider, 1958).

Results indicated that participants simultaneously encoded each player's reward value, based on the absolute amount, and trait-level generosity, based on the proportional amount, as revealed in behavioral choice patterns and neural activity. Although behavioral economic models suggest that reward value alone should drive choice preferences, an analysis that considered the joint effects of reward-based and trait-based learning found that choices were more strongly guided by traits. Moreover, while both forms of learning were associated with prediction error signals in ventral striatum, a hallmark of feedback-based reward learning (Garrison et al., 2013), only trait learning was additionally associated with regions implicated in social impression updating (Mende-Siedlecki, Baron, & Todorov, 2013; Mende-Siedlecki, Cai, & Todorov, 2013; Mende-Siedlecki & Todorov, 2016).

This research also demonstrated that, like reward learning, the feedback-based learning of traits may be a domain-general process: participants formed both reward-based and trait-based associations when interacting with humans as well as with nonhuman objects (e.g., slot machines), both of which may be characterized in terms of their absolute reward value and their proportional generosity (Baetens et al., 2017; Spunt & Adolphs, 2015; Waytz et al., 2014). Yet, while these learning processes may be domain-general, the process of inferring traits through interactive feedback is particularly

relevant to social relations (Heider, 1958). Indeed, subsequent work revealed that people rely more on trait- than reward-based information when interacting with humans in comparison to interacting with slot machines (Hackel et al., 2020). Thus, although trait-based learning is not unique to social interactions, people appear to rely on it more heavily when learning about social than nonsocial agents.

Together, these recent findings suggest that trait inferences can indeed be made through direct interactive learning (i.e., reinforcement learning). Importantly, given our present questions, they raise the possibility that, in contrast to reward-based learning, interactive trait learning can support context-specific impression formation and context-appropriate social decision-making.

Effects of Interactive Reward and Trait Learning on Flexible Social Choices

How might interaction-based trait impressions support context-dependent choice? Trait concepts provide nuanced, high-level information about a stable characteristic, such as one's competence in computer programming as opposed to creative writing. This conceptual specificity afforded by trait inference, rooted in semantic knowledge structures (Amodio, 2019), lends itself to flexible, context-specific choices. That is, people have rich conceptual knowledge of how different traits relate to one another and to specific behaviors; for instance, a person who shares resources is "cooperative," and "cooperative" is distinct from "assertive" but similar to "trustworthy" (Stolier et al., 2020). This knowledge permits people to form meaningful inferences and predictions about an individual's behavior across contexts that differ in the specific expression of a trait—for instance, understanding why a partner was helpful in tasks that required the sharing of resources but not in tasks requiring negotiation. People could estimate the value of interacting with a partner in a specific setting based on this fine-grained trait knowledge.

In turn, this possibility suggests a new prediction about how people generalize from past interaction experiences to choose partners in novel social contexts: People might choose partners based on the similarity of traits required by new and previous contexts. For instance, imagine a colleague who offers excellent help editing a manuscript. Beyond generalizing to contexts that appear similar (e.g., help editing an email), people may generalize to contexts that require *traits* perceived as similar (e.g., offering empathy, to the extent that a perceiver views socioemotional skills as similar to writing skills), even if the contexts bear little surface similarity to each other.

Trait knowledge can therefore serve as a *conceptual map* that allows people not only to make sense of context-dependent feedback during learning but also to generalize in new contexts. In contrast, reward feedback offers a common metric of value for different kinds of outcomes, without providing the same conceptual specificity that allows people to make sense of context-dependent feedback during learning. For this reason, reward-based associations are less likely to generalize adaptively across context.

Overview of Studies

In the present research, we investigated how people learn about social partners through both reward feedback and social trait feedback during interactions in different contexts. That is, we

measured the extent to which social choices are guided by *context dependent* as opposed to *global* learning of traits and rewards. We hypothesized that trait learning affords greater context-dependence than reward learning in social interactions, such that participants would choose interaction partners as a function of their distinct competencies in different contexts. We tested this core hypothesis in four studies.

In addition, in Studies 3 and 4, we examined participant choices and social preferences in novel contexts, including preferences for empathy and social support. We hypothesized that participants would generalize value based on the conceptual similarity of traits required in old and novel contexts. In this manner, trait learning would afford greater flexibility when applying learning to choices in novel contexts within social interactions.

In each study, participants learned about other individuals in a money sharing task, modeled after previously established reinforcement learning paradigms (Hackel et al., 2015). Although real-life social interaction is often very complex, this task captures the essential feature of active, interaction-based learning: the process of one individual engaging with another person and receiving feedback from that person. This form of learning differs from previously studied passive social learning processes, such as reading about another person or observing their behavior, which do not involve feedback on one's action—a crucial distinction in terms of underlying learning mechanism and implications for decisions and behavior (e.g., Foerde et al., 2006; Poldrack et al., 2001). Hence, this task provided a valid means for testing our hypotheses regarding the roles of reward-based and trait-based in interactive social learning across contexts.

Study 1

In Study 1, participants played a game in which they learned about other players through their active choices and feedback across repeated interactions in two different contexts (in domains assessing math vs. verbal ability). In each context, feedback indicated (a) the absolute reward value of choosing a particular player on a specific trial as well as (b) the player's trait-level competence, as revealed by the proportion of available points they earned. By independently varying the reward value and trait competence associated with each player in each context, this design permitted us to test the degree to which each aspect of feedback—reward and trait—was learned and applied to decisions across contexts. We expected participants to learn from both the rewards and traits associated with each player but hypothesized that cross-context interactive learning and decision-making would rely more heavily on trait-based than reward-based inferences.

Method

Overview

We adapted a previously validated task to distinguish reward and trait learning (Hackel et al., 2015). Participants played a game in which they learned about four people who had ostensibly participated in prior sessions. Participants were told these other players had answered verbal and math questions from the Graduate Record Examination (GRE) in order to win points that would be converted to money. Participants were further told that GRE questions varied

in the number of points they were worth, and that players had won a *proportion* of these points on each trial based on how quickly and accurately they answered each question. Hence, the absolute amount of points associated with a player on a given trial (e.g., reward value) varied independently of the proportion they earned (indicating their competency in answering the question). This design allowed us to independently manipulate the reward value and trait implications of each instance of feedback.

Critically, players were presented in two different contexts: a “math” context, representing responses to math questions, and a “verbal” context, representing responses to verbal questions. Thus, participants’ choices could reflect the average rewards a player provided in each context, their competencies in each context, or both. This design allowed us to distinguish context-dependent and global learning of rewards and traits.

We report how we determined our sample size, all data exclusions, all manipulations, and all measures included in each study. Sample size for each study was determined before any data analysis. De-identified data from each study have been made available at: https://osf.io/496rn/?view_only=1d5f069aa26b421f9838e7328ba8f6a6.

Participants

Fifty undergraduate students from New York University (10 males, 40 females, $M_{\text{age}} = 18.94$, $SD = .91$) participated in Study 1 for course credit. Sample size was determined a priori as a minimum of 50 participants, based on past work using a similar task (within-participants design with many trials per participant) which produced large effect sizes (Hackel et al., 2015). The study was approved by the New York University Committee on Activities Involving Human Subjects, and informed consent was collected from all participants. Although this sample was relatively homogeneous in terms of age, past work has observed similar findings across student and online samples when using similar tasks (Hackel et al., 2015, 2020; Hackel & Zaki, 2018).

Stimuli

Participants viewed avatars representing four previous players. Gender and race were held constant across the four avatars to avoid any cues to social group membership that could influence social judgments. For each participant, the four avatars were randomly assigned to the four competence/reward levels (see Table S1), ensuring that variability in visual features (e.g., hairstyle) was randomly distributed across player types and would not influence effects of interest.

Players varied independently on math and verbal ability—some were high on one but low on the other, high on both, or low on both—as reflected by the average proportion of points they accrued in each domain (Table S1). Each player was also associated with a unique combination of reward levels (i.e., high or low average amount earned) on verbal and math trials; math reward values and verbal reward values were also orthogonal to one another. Finally, rewards in a given context (e.g., verbal rewards) were orthogonal to competence in the same context (e.g., verbal ability; see Table S1 in Supplemental Materials for average trait and reward values in each context).

In order to ensure sufficient learning without excessive difficulty, the target set was restricted to four individuals. This required an orthogonalization scheme in which the pairs of variables listed above were independent of one another within subjects and the remaining pairs of variables were rendered orthogonal between subjects, with subjects randomly assigned to have a positive or negative correlation between these variables (see Table S1). Hierarchical analyses including data from all participants could therefore distinguish effects of reward learning and trait learning.

Procedure

Upon arriving at the experiment site, participants were told that they would be assigned to one of two roles. In a Player A role, they would answer GRE questions to win points worth money, based on how quickly and accurately they answer. In a Player B role, they would learn about the responses of previous Player A participants and win points as a result of how Player A participants performed. In reality, all participants were assigned to the Player B role through a rigged drawing, such that participants would learn from the behaviors of past players (see Supplemental Material for complete instructions text).

This task included a training phase and a test phase. During the training phase, participants could learn about players through their choices and players’ feedback across repeated interactions. This phase allowed us to examine the process of feedback-based learning across contexts. The test phase was designed to assess decision outcomes; participants made choices without feedback and their choice patterns yielded a fine-grained assessment of the reward-based and trait-based choice.

Training Phase

In the training phase of the task, participants were told that players had answered a verbal or math GRE question on each round of a prior session and won a proportion of available points based on the speed and accuracy of their response. To prevent the belief that point pool amounts were related to question difficulty or that players were more strongly motivated by larger points pools, participants were told that the point pool was set by a computer algorithm on each round, in a manner unrelated to the difficulty of the questions, and that the point pool was unknown to each prior participant before answering. These instructions were again emphasized in a post-instruction comprehension quiz delivered by the experimenter.

Participants were informed that, on each round, they would see two of the four players and pick one of them to “hire” for the round. They were informed that choices would be followed by feedback indicating the number of points that player earned as well as the point pool available to the player on that round. Moreover, participants were told they would receive the number of points earned by their chosen player on each round, allowing them to accumulate points that would be exchanged for money at the end of the experiment. In other words, participants received money based on the amount earned by the player they chose on each round.

The task instructions referred equally to variability in rewards and traits associated with players. Additionally, instructions made it clear to participants that their explicit goal was to earn money. Finally, this task has been shown to support learning on the basis of both reward and trait information (Hackel et al., 2015, 2020). Hence,

the task and instructions were designed to ensure equal opportunities to learn from reward value and trait information.

Participants completed 168 trials of the learning phase, including 84 verbal trials and 84 math trials. Verbal and math trials were pseudorandomly interleaved in a different order for each participant, across three blocks of 56 trials (with short rest breaks in between to avoid fatigue). At the start of each trial, the word “Math” or “Verbal” was displayed for 2 s to indicate the upcoming trial type (Figure 1). Immediately following the cue, participants viewed two avatars; each possible pair of avatars was encountered 14 times in each context. Each avatar was equally likely to appear on each side of the screen. During a subsequent 2-s choice epoch, participants indicated which avatar they preferred to hire by pressing one of two response keys. Immediately following choice, feedback was displayed, for 3 s, indicating the number of points the selected participant earned

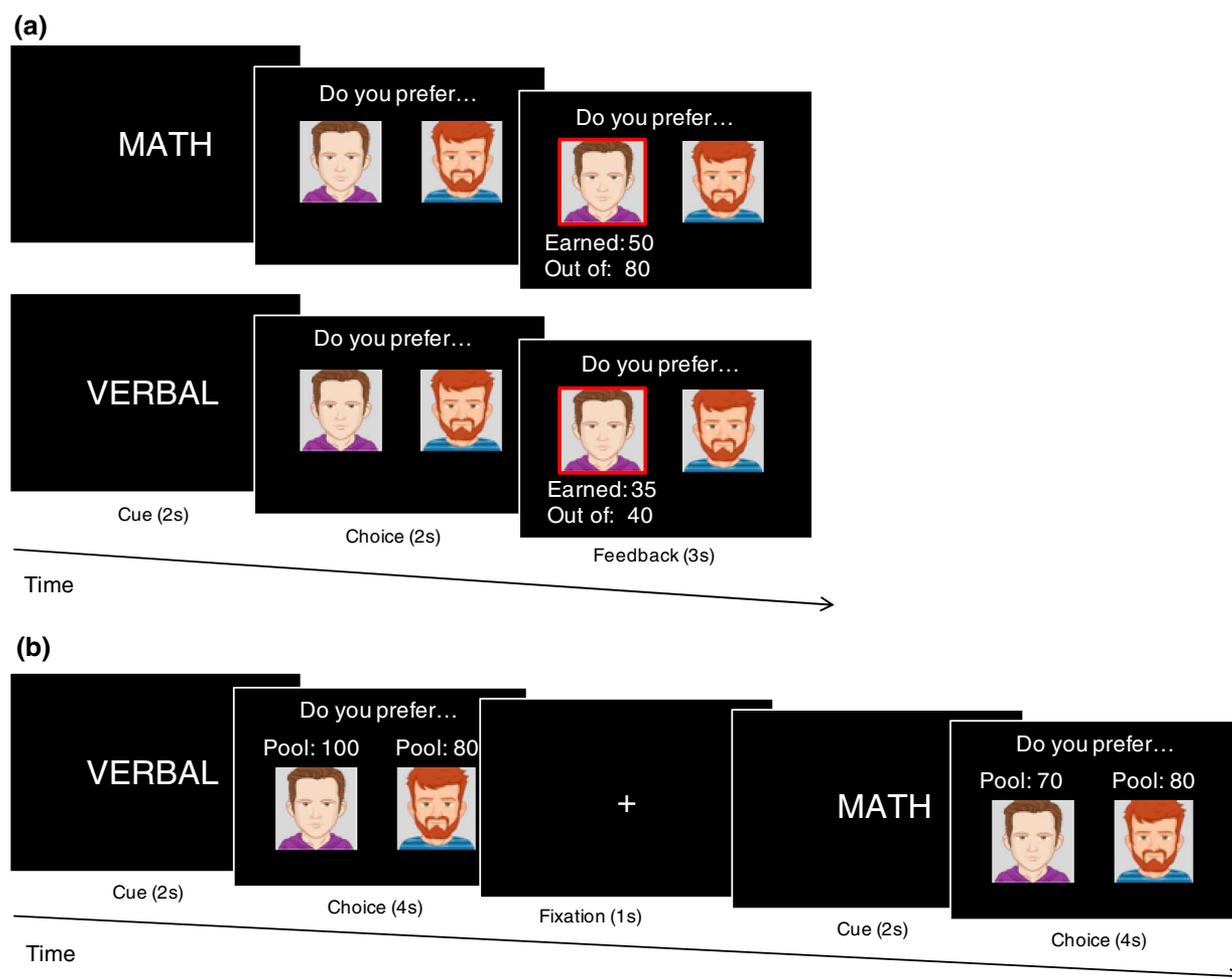
(labeled “earned”) as well as the total number of points available to the player on that trial (labeled “out of”).

The reward and competence values displayed during feedback were generated using average values for the chosen player (displayed in Table S1) plus noise with $SD = 0.15$ for competence and $SD = 7.5$ for reward. These noise distributions were equivalent given the difference in scaling for average trait and reward values, even though the raw values of the standard deviations differed, and therefore equated for learning difficulty. The task was implemented using the Psychtoolbox for Matlab (Kleiner et al., 2007; Pelli, 1997).

Test Phase

In the test phase portion of the task, participants were informed that the point pool associated with each player would now be shown

Figure 1
Schematic of Learning and Choice Task



Note. (a) In the learning phase, participants made choices to “hire” four ostensible earlier participants who completed verbal and math GRE questions to win points worth money; participants won any points earned by the player they hired. A cue before each trial revealed the upcoming learning context (verbal or math); contexts were interleaved (Studies 1–3) or blocked (Study 4). Feedback indicated the number of points earned by the partner (and thus reward value to the participant), as well as the proportion of available points earned (indicating the competence of the partner). (b) In the test phase, participants chose a partner while viewing the point pool available for partners to earn. No feedback was provided. GRE = Graduate Record Examination. See the online article for the color version of this figure.

above each picture before each choice so they can use this information to make decisions. In addition, participants were informed their winnings for this phase would be displayed only after completion of the block rather than after each choice. This phase allowed a test of prior learning without any further adaptation to feedback.

Participants completed 192 trials of the test phase, including 96 verbal trials and 96 math trials. As in the training phase, verbal and math trials were pseudorandomly interleaved. Trials were divided into four blocks of 48 trials.

At the start of each test phase trial, the word “Math” or “Verbal” again appeared on screen for 2 s to indicate trial type. Following the cue, participants viewed two of the four avatars; each pair of avatars was seen 16 times in each context. Each avatar was equally likely to appear on each side of the screen. Above each avatar, a point pool was displayed representing the total number of points available for the question being answered. To determine point pools, a random integer from 10 to 100 was generated as the pool value for one of the players. Next, this amount was multiplied by one of seven ratios (0.33, 0.67, 0.9, 1, 1.11, 1.5, 3) to determine the point pool for the second player. The ratios were designed to be symmetric around 1, allowing fine-grained expression of competence knowledge (e.g., knowing to choose a target who has 100 points available and earns 80% over a target who has 150 points available and earns 40%). Each avatar pair appeared in each context four times at a 1:1 ratio and two times at each other ratio. Extra trials at the 1:1 ratio because they were particularly informative about choices rooted in prior learning, given the equivalent point pools. Participants made choices in a 4-s decision epoch. No further feedback was provided in order to prevent further learning.

Posttask Ratings

Immediately following task completion, participants completed posttask ratings to test whether their learning transferred to preferences beyond the economically framed task feedback. First, they were asked how much they would like to be paired with each of the players encountered in the task on an assignment in a Statistics class and how much they would like to be paired with each of the players on an assignment for an English class. Next, to test whether the task did indeed generate conceptual trait impressions, participants rated each of the previously encountered players on degree of verbal ability, math ability, and overall intelligence. Each set of ratings was completed using 7-point Likert-type items (1 = *not at all*, 7 = *very much*). Finally, for exploratory purposes, participants completed a measure of social self-efficacy (Sherer et al., 1982) which was unrelated to behavior in the task and is not discussed further. Participants then completed demographic measures, were paid their performance-based bonus winnings, debriefed, and thanked.

Results

Learning

To determine whether participants learned in a context-dependent manner, we tested the degree to which reward and trait feedback from a given player shaped future decisions to hire that player, within and across contexts during the learning phase. To this end, we fit learning phase data to a lagged, mixed-effects logistic regression. This regression predicted the log odds that a participant stays with

the most recently chosen player of the two shown onscreen on a given trial (coded as stay = 1, switch = 0). Predictors included (1) reward feedback most recently received from that player (i.e., amount received), (2) competence feedback most recently received from that player (i.e., proportion received), (3) a variable indicating whether the current trial matched (1) or differed from (−1) the previous context, and (4 & 5) interactions of the context variable with reward and trait feedback. This analysis therefore used a regression framework to examine whether feedback reinforces choices—that is, whether higher rewards and higher competence lead participants to stay with the same player when next available, as a function of context. In past work, lagged regression analyses have been used as a model-free approximation to reinforcement learning models (see Doll et al., 2015; Otto et al., 2013).

In the regression analysis, continuous predictors were standardized to *z*-scores within participants to allow meaningful comparison between variables, and we included fixed effects and random variances for the intercept and all slopes. Analyses were performed using the *lme4* and *lmerTest* packages for R (Bates et al., 2015; Kuznetsova et al., 2016; R Development Core Team, 2016). To compare the contributions of relevant and irrelevant rewards and traits, we contrasted coefficients against one another using the *doBy* package for R (Højsgaard & Halekoh, 2016). All coefficients and contrasts are reported in Table S3.

Results indicated that competence feedback reinforced choices globally (i.e., across contexts), as revealed by a main effect of competence, $b = .45$, $SE = .04$, $z = 11.91$, $p < .001$. However, this effect was qualified by a Competence \times Context interaction, $b = .27$, $SE = .04$, $z = 7.49$, $p < .001$. Simple effects analysis indicated that competence more strongly reinforced choices within the same context, $b = .72$, $SE = .05$, $z = 13.42$, $p < .001$, and more weakly reinforced choices in a different context, $b = .18$, $SE = .05$, $z = 3.62$, $p < .001$. That is, if a player displayed high competence on a math trial, participants were very likely to choose that player again on a subsequent math trial and only somewhat likely to choose that player again on a subsequent verbal trial. Hence, trait learning occurred in a highly context-dependent manner.

Strikingly, we found no evidence that rewards reinforced choices, whether overall, $b = .02$, $SE = .04$, $z = .65$, $p = .51$, or in a context-dependent manner, $b = .01$, $SE = .0035$, $z = .36$, $p = .72$. Thus, neither *context-dependent* nor *global* experiences of prior reward guided social decision-making. Moreover, in a linear contrast of coefficients, the effect of trait competence was significantly greater than the effect of prior reward values ($\chi^2 = 68.50$, $p < .0001$).

Computational Model of Learning

The regression analyses approximated a reinforcement learning model while using a conventional linear model framework. To complement these analyses, we next fit a hierarchical computational model of learning to fully modeled the nonlinear dynamics of learning (Hackel & Amodio, 2018). We fit behavior during the learning phase to an adapted reinforcement learning model that hybridizes reward and trait learning, based on prior work (Hackel et al., 2015). Broadly, this model assumes people learn the reward value associated with each player in each context and the competence of each player in each context (a total of four values). The model then allows integration of these values via a choice weight (β) for each one. These choice weights provide insight into whether

each set of values does, in fact, guide choice. We fit values produced by this model to all participants' data using a hierarchical approach, allowing us to test whether each choice weight was significantly different from zero on average (using Wald tests; see Table S2 for all coefficients).

Unlike prior work (Hackel et al., 2015), the present model allowed separate learning of traits and reward values across two contexts. That is, the model separately learned reward values (QV) and competence values (CV) for the verbal context, and reward values (QM) and competence values (CM) for the math context. On every trial t of the learning phase, the model assumes that participants update reward values and competence values for the chosen target within the currently relevant context, according to Equations 1 and 2:

$$Q_{\text{relevant},t} = Q_{\text{relevant},t-1} + \alpha \delta_{R,\text{relevant},t}, \quad (1)$$

$$C_{\text{relevant},t} = C_{\text{relevant},t-1} + \alpha \delta_{C,\text{relevant},t}, \quad (2)$$

where Q_{relevant} is set to QV on verbal trials and QM on math trials, C_{relevant} is set to CV on verbal trials and CM on math trials, α is a free parameter representing a learning rate, $\delta_{R,\text{relevant},t}$ represents a reward prediction error for the relevant context, and $\delta_{C,\text{relevant},t}$ represents a competence prediction error for the relevant context. We fit one learning rate to both reward and competence based on prior work (Hackel et al., 2015); this feature reduces the number of free parameters and avoids trade-offs between learning rates and choice weights, thus stabilizing the model. Prediction errors were defined as the difference between values received and values expected for reward and competence, according to Equations 3 and 4:

$$\delta_{R,\text{relevant},t} = \text{Reward}_t - Q_{\text{relevant},t-1}, \quad (3)$$

$$\delta_{C,\text{relevant},t} = \text{Competence}_t - C_{\text{relevant},t-1}. \quad (4)$$

Reward was defined as the number of points received, and competence was defined as the proportion of available points earned. During choices, the model allowed integration of competence and reward via different choice weights in a softmax choice function, according to Equation 5:

(See below)

where $p_{i,t}$ is the probability of choosing option i (of j options) on trial t , and each β is a separate choice weight for the corresponding Q and C values. Thus, the model allowed separate choice weights for the relevant reward values, irrelevant reward values, relevant competence values, and irrelevant competence values. On math trials, QM and CM were coded as *relevant*, while QV and CV were coded as *irrelevant*, and on verbal trials, QV and CV were coded as *relevant*, while QM and CM were coded as *irrelevant*. Before fitting this model, we standardized reward and competence feedback (to z -scores) for each participant in order to equate the scaling of the two variables and allow meaningful comparisons of choice weights for each. Both reward expectations and competence expectations were initialized to zero to correspond to participants' initial expectations based on task instructions.

Reinforcement learning models are commonly fit separately for each participant, and parameters fit across participants are then aggregated into summary statistics (e.g., mean or median) for analysis. However, in the present work, our analyses relied on a between-participants counterbalancing: two variables were positively correlated for half the subjects and negatively correlated for the other half of subjects, as described above. This design precluded single-subject models, since these models would be unable to dissociate the effects of two correlated variables within a participant. Therefore, we used a hierarchical modeling approach (Daw, 2011). This model assumes that each participant's parameters are drawn from a population distribution of parameters. The model directly fits the population distribution using data from all subjects, analogous to mixed-effects regression. That is, each choice parameter is assumed to have a population average (akin to a fixed effect) and a variance (akin to a random effect).

As in mixed-effects regression, standard errors for the average parameters can be computed (Daw, 2011; see Supplemental Materials). We used these standard errors to construct Wald tests in order to test if choice parameters were significantly different from zero (see Table S2). In addition to testing the choice parameters against zero, we contrasted parameters against one another—for example, comparing the choice weight for competence in the relevant context against the choice weight in the irrelevant context—using linear contrast of coefficients (see Supplemental Materials).

Consistent with the regression analyses, results indicated that competence feedback strongly reinforced choices within the same context, $\beta = 1.68$, $SE = .11$, $z = 15.00$, $p < .0001$, and more weakly reinforced choices in a different context, $\beta = .69$, $SE = .13$, $z = 5.28$, $p < .0001$. That is, if a player displayed high competence on math trials, participants were very likely to choose that player again on a subsequent math trial, and only somewhat likely to choose that player again on a subsequent verbal trial. The difference between these coefficients was itself significant, $z = 4.88$, $p < .0001$, indicating that context-sensitive learning was stronger than global learning. Again, we found no evidence that rewards reinforced choices, whether in the same context, $\beta = -.08$, $SE = .10$, $z = -.83$, $p = .41$, or in a different context, $\beta = .14$, $SE = .10$, $z = 1.40$, $p = .16$. Thus, neither *context-specific* nor *global* experiences of prior reward guided social decision-making. Moreover, in a linear contrast of coefficients, the effects of trait competence—across relevant and irrelevant contexts—were significantly greater than the effects of prior reward values ($z = 14.63$, $p < .0001$). Thus, the results of this computational modeling approach were consistent with those obtained using a regression approach.

Test Phase

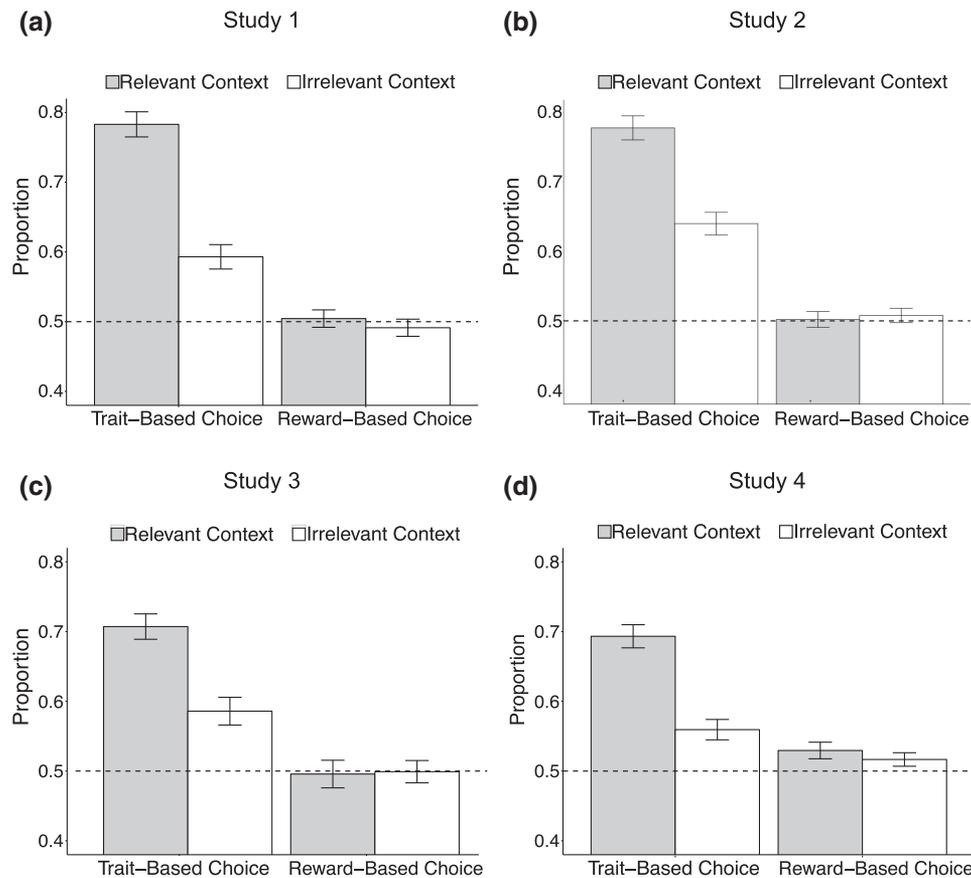
Similar results were observed in the test phase, in which participants had to apply prior learning to make fine-grained decisions in the absence of feedback. We fit a mixed-effects logistic regression that predicted the log odds of choosing the player on the right side of the screen (arbitrarily chosen), as a function of the difference

$$p_{i,t} = \frac{\exp(\beta_{Q,\text{relevant}} \times Q_{\text{relevant},i,t} + \beta_{Q,\text{irrelevant}} \times Q_{\text{irrelevant},i,t} + \beta_{C,\text{relevant}} \times C_{\text{relevant},i,t} + \beta_{C,\text{irrelevant}} \times C_{\text{irrelevant},i,t})}{\sum_j \exp(\beta_{Q,\text{relevant}} \times Q_{\text{relevant},j,t} + \beta_{Q,\text{irrelevant}} \times Q_{\text{irrelevant},j,t} + \beta_{C,\text{relevant}} \times C_{\text{relevant},j,t} + \beta_{C,\text{irrelevant}} \times C_{\text{irrelevant},j,t})}, \quad (5)$$

between the two players (right–left) in (1) point pools, (2) competence in the relevant context, and (3) competence in the irrelevant context, (3) reward value in the relevant context, and (4) reward value in the irrelevant context. (On “verbal” trials, verbal ability was coded as relevant and math ability was coded as irrelevant; on “math” trials, the reverse was true.) We used the true underlying competence and reward values of each player as predictors so that this analysis would be independent of any assumptions derived from the learning model. Continuous predictors were again standardized (within participant, to z -scores) to allow meaningful comparison between variables, and we included fixed effects and random variances for the intercept and all slopes. All coefficients and contrasts are reported in Table S4.

As in the learning phase, participants were more likely to choose players who were competent (as opposed to incompetent) in the relevant context, $b = 1.78$, $SE = .18$, $z = 9.83$, $p < .001$, while also showing a weaker tendency to choose players who were competent (as opposed to incompetent) in the irrelevant context, $b = .58$, $SE = .12$, $z = 4.76$, $p < .001$ (Figure 2a). The difference between these coefficients was significantly different from zero, $\chi^2 = 30.38$, $p < .001$, indicating that context-dependent choice was stronger than global choice of competent targets. However, as in the learning phase, test phase choices showed no influence of prior reward values in either the relevant context, $b = .07$, $SE = .08$, $z = .91$, $p = .37$, or the irrelevant context, $b = -.02$, $SE = .08$, $z = -.20$, $p = .84$. Moreover, the effects of trait competence were significantly greater

Figure 2
Test Phase Choices



Note. The plot shows the proportion of test phase choices for which participants chose a partner with higher competence in the relevant context, higher competence in the irrelevant context, higher reward value in the relevant context, and higher reward value in the irrelevant context, split by study. Each index is computed independently of the others (see Supplemental Methods). Error bars show SEM, adjusted for within-subject comparisons (Morey, 2008). Dashed lines indicate chance. (a–c) In Studies 1–3 ($N = 50$, $N = 50$, $N = 51$), in which contexts were interleaved during learning, participants chose partners primarily based on trait-level competence in the relevant context, less so based on competence in the irrelevant context, and not at all based on reward values from either context. (d) In Study 4 ($N = 85$), in which contexts were blocked instead of interleaved during learning, participants chose partners primarily based on competence in the relevant context, less so based on competence in the irrelevant context, weakly based on reward value in the relevant context, and nonsignificantly based on reward values in the irrelevant context. SEM = standard error of the mean.

than the effects of prior reward values, $ps \leq .001$ (see Table S4 for all coefficients and contrasts).

Posttask Ratings

To determine whether participants applied their learning to noneconomic preferences outside the main task, and to determine whether participants did form conceptual representations of traits, we analyzed explicit ratings of partner preferences, intelligence, and verbal and math ability using Generalized Estimating Equations (GEE; Liang & Zeger, 1986). We did not use mixed-effects regression because each participant made one only rating for each player in each context, which produced too few data points to fit random slopes. Instead, GEE fits a marginal regression model that accounts for within-subject dependencies. (For linear models, GEE and mixed models provide equivalent interpretations; Fitzmaurice et al., 2004.) We fit the models using the geepack package for R (Halekoh et al., 2006). The model predicted partner preference and ability ratings for each player as a function of that player's (1) relevant competence, (2) irrelevant competence, (3) average reward value in the relevant context, and (4) average reward value in the irrelevant context. (For the Statistics assignment and ratings of math ability, math competence was coded as "relevant," whereas verbal competence was coded as "irrelevant"; for the English assignment and ratings of verbal ability, the opposite coding was used.) Again, we used true underlying values as predictors, and all predictors were standardized to z -scores within-subject before analysis. All coefficients and contrasts are reported in Tables S5–S7.

When asked to rate their preference for partners in academic assignments in a Statistics versus English course, participants showed a similar pattern of trait-based preference as in the choice task. Participants' ratings relied primarily on the relevant competency, $b = 1.40$, $SE = .11$, $\chi^2 = 169.88$, $p < .001$, secondarily on the irrelevant competency, $b = .48$, $SE = .08$, $\chi^2 = 32.85$, $p < .001$, and not at all on previously experienced reward values in the relevant context, $b = .02$, $SE = .07$, $\chi^2 = .11$, $p = .74$, or the irrelevant context, $b = .08$, $SE = .06$, $\chi^2 = 1.85$, $p = .17$ (Table S5). Moreover, when asked to rate the abilities of each player, participants provided context-dependent impressions of verbal and math ability. That is, they relied more on the relevant ability, $b = .40$, $SE = .12$, $\chi^2 = 11.35$, $p < .001$, than the irrelevant ability, $b = .11$, $SE = .10$, $\chi^2 = 1.16$, $p = .28$, as revealed in a linear contrast of coefficients, $\chi^2 = 5.20$, $p = .02$ (Table S6). When asked to rate impressions of overall intelligence, participants relied to a similar degree on learning from the verbal context, $b = .88$, $SE = .10$, $\chi^2 = 84.59$, $p < .001$, and the math context, $b = .91$, $SE = .07$, $\chi^2 = 149.86$, $p < .001$ (Table S7). Thus, participants formed both context-dependent and global impressions of intelligence, yet context-dependent learning was more predictive of choice.

Discussion

In Study 1, we examined the extent to which people rely on reward and trait information during interaction-based social learning across contexts. We found that trait-based information was learned in a context-dependent manner, above and beyond participants' global (i.e., cross-context) trait impression. By contrast, there was no evidence for the influence of reward-based information in either global or context-dependent learning. That is, participants primarily

hired workers skilled in math to answer math problems and workers skilled in verbal ability to answer verbal problems, demonstrating context-dependence. A similar pattern was observed in decision-making: in the test phase, participants' interaction choices reflected context-specific trait learning, beyond effects of global trait learning, but did not reflect any influence of reward learning.

Although reward-based learning was not predicted to be sensitive to context, the lack of any reward-based learning in this study was surprising. Indeed, simulations of a similar task in past work demonstrated that selections based solely on reward value is the optimal strategy for earning money during learning (Hackel et al., 2015), and thus if anything, the task structure favored an emphasis on reward learning. Current models of reward reinforcement learning suggest that reward associations should be prioritized in learning and decision-making, but it is possible that reward information lacks the nuance needed for adaptive cross-context social decisions. We return to this issue in Study 4.

Participants' explicit trait ratings of players indicated that trait inferences apparent in participant's choice behaviors corresponded to the kind of conceptual trait inference more commonly studied in social cognition research. Indeed, players associated with high-competence feedback were rated as being more intelligent, and these ratings showed context specificity. Hence, the trait competence information conveyed through interaction and feedback related to participants' conceptual trait representations, supporting our interpretation of this process as supporting trait inference.

Finally, we showed that context-specific trait inferences formed through interaction and feedback generalized to noneconomic social decisions. That is, players whose feedback indicated higher competence in math were more likely to be approached for help with statistics homework, whereas players whose feedback indicated higher competence in verbal ability were more likely to be approached for help with English homework. This pattern adds converging evidence for the context specificity of feedback-based trait inference and begins to demonstrate the functional utility of this form of interactive social learning.

Together, these findings provide initial evidence for our hypothesis that trait-based feedback, in comparison with reward feedback, more strongly supports the ability to form context-specific impressions through direct social interaction and to apply them to interaction choices.

Study 2

Study 2 was conducted to address two goals: First, we sought to replicate the results of Study 1. Second, we sought to address a potential limitation of Study 1 whereby the instructions may have led participants to expect greater context specificity in players' traits than rewards. Specifically, the descriptions of players' verbal and math GRE questions may have implied that competence is likely to vary across these two test contexts, whereas the instructions did not provide a corresponding rationale for contextual variation in rewards. Although Study 1 instructions did not explicitly refer to contextual variation in either type of feedback, participants may nonetheless have formed different expectations of variability in rewards and traits. Thus, in Study 2, instructions were framed to explain and emphasize the context specificity of reward feedback. Despite an emphasis on context-specific rewards, we again

predicted stronger context-based effects for traits in learning and choice.

Method

Participants

Fifty undergraduate students from the University of Delaware (33 females, 17 males, $M_{\text{age}} = 19.14$, $SD = 1.70$) participated in Study 2 for course credit. Sample size was determined a priori as a minimum of 50 participants, based on Study 1. The study was approved by the University of Delaware Institutional Review Board, and informed consent was collected from all participants.

Stimuli

Stimuli were identical to Study 1.

Procedure

The procedure was identical to that of Study 1—participants completed a learning phase, a test phase, and posttask ratings of each target—with two exceptions.

First, the learning phase instructions were altered to provide an explanation of why rewards could vary across contexts. In addition to the instructions described in Study 1, participants were told that points had been assigned to players on each round by either a blue slot machine that paid out many points on average or a red slot machine that paid out few points on average. Participants were told that each player was randomly assigned the blue or red slot machine for each context (verbal and math), and that players always received points from that assigned slot machine within a given context. These instructions explicitly noted that “some Player A participants got bigger point pools for one type of question than the other.” The instructions thus offered a rationale for why point pools—and therefore reward outcomes—could vary systematically across contexts.

Second, three exploratory posttask ratings were added to those collected in Study 1. In addition to measuring context-dependent impressions of competence (verbal ability, math ability) and global impressions of competence (intelligence), as in Study 1, we further examined whether participants formed context-dependent and global impressions from reward feedback. First, we evaluated whether participants explicitly formed impressions of target point pools across settings. Participants were reminded that each player’s point pools were assigned by a “good” or “bad” slot machine, and that these assignments varied between Math and Verbal questions; as a result, they were told could think of these point pools in terms of “wealth,” given that some players had more points at stake during one type of question or the other. Participants rated how “wealthy” each player was in each context. If participants formed context-specific impressions from reward feedback, they should rate players as “wealthier” in contexts in which they had larger point pools. Next, to evaluate whether participants formed *global* impressions based on reward feedback, participants rated how lucky the player was overall, using 7-point Likert scales ranging from 1 (*not at all*) to 7 (*very much*). If participants formed overall conceptual impressions due to reward feedback, they should rate players as “luckier” based on rewards accrued in both contexts. Finally, participants also rated the generosity of each player, using the same scale, in order to

examine whether halo effects (e.g., inferring generosity from competent targets) might be bounded by context (see Supplemental Materials). To compensate for the time added by making these ratings, we removed the ratings of partner desirability for different types of academic assignments.

Results

Learning

As in Study 1, we tested the degree to which reward and trait feedback from a given player shaped future decisions to hire that player, within and across contexts. We again fit a mixed-effects logistic regression predicting the likelihood participants stay with a previously chosen player, given the competence previously displayed, reward previously received, and whether the context was the same or different (Table S3), along with the computational model (Table S2).

Replicating the results of the regression model in Study 1, we observed a main effect of competence, $b = .47$, $SE = .04$, $z = 10.83$, $p < .001$, qualified by the predicted Competence \times Context interaction, $b = .21$, $SE = .05$, $z = 4.22$, $p < .001$, which indicated that competence feedback more strongly influenced choices within the same context than the alternative context. Simple effects analysis revealed that competence strongly reinforced choices within the same context, $b = .68$, $SE = .07$, $z = 10.16$, $p < .001$, but more weakly reinforced choices in the other context, $b = .26$, $SE = .07$, $z = 3.97$, $p < .001$.

Analysis of reward feedback effects again produced no evidence that rewards reinforced choices, $b = .05$, $SE = .04$, $z = 1.24$, $p = .22$, or that context moderated the effect of reward feedback, $b = -.01$, $SE = .03$, $z = -.42$, $p = .67$. Thus, even when the instructions provided participants with an explanation for why reward feedback might differ across contexts, participants learned from traits—but not rewards—in a context-dependent manner.

Computational Mode of Learning

Consistent with the regression results, the computational model similarly indicated that competence feedback strongly reinforced choices within the same context, $\beta = 1.57$, $SE = .12$, $z = 13.04$, $p < .0001$, and more weakly reinforced choices in a different context, $\beta = .73$, $SE = .08$, $z = 9.36$, $p < .0001$. The difference between these coefficients was itself significant, $z = 6.16$, $p < .0001$, demonstrating that context-sensitive learning was stronger than global learning. We did not observe evidence that rewards reinforced choices in the same context, $\beta = .11$, $SE = .07$, $z = 1.57$, $p = .12$, although choices were reinforced by rewards in a different context, $\beta = .29$, $SE = .09$, $z = 3.32$, $p < .001$. Given that no other analyses detected this effect, that it was restricted only to the irrelevant context, and that the effect size was still far smaller than the effects of competencies, we are cautious in interpreting this result. Moreover, in a linear contrast of coefficients, the effects of trait competence—across relevant and irrelevant contexts—were significantly greater than the effects of prior reward values ($z = 8.50$, $p < .0001$).

Test Phase

Test phase choices were analyzed using the same mixed-effects logistic regression analysis reported in Study 1. Results once again

replicated those of Study 1: Participants were more likely to choose players who were competent (as opposed to incompetent) in the relevant context, $b = 1.58$, $SE = .14$, $z = 11.63$, $p < .001$, while also showing a weaker tendency to choose players who were competent (as opposed to incompetent) in the irrelevant context, $b = .77$, $SE = .12$, $z = 6.67$, $p < .001$ (Figure 2b). The difference between these coefficients was significantly different from zero, $\chi^2 = 21.17$, $p < .001$, indicating that context-dependent choice was stronger than global choice of competent targets. Yet, test phase choices again showed no influence of prior reward values in the relevant context, $b = .01$, $SE = .08$, $z = .15$, $p = .88$, or the irrelevant context, $b = .06$, $SE = .08$, $z = .68$, $p = .50$. Again, each effect of trait competence was significantly greater than the corresponding effect of prior reward values, $ps \leq .001$ (see Table S4 for all coefficients and contrasts).

Posttask Ratings

To analyze explicit ratings of partner ability, we used the same GEE regression approach described in Study 1. All coefficients and contrasts are reported in Table S6 and S7.

We found identical patterns of results as in Study 1. When asked to rate the abilities of each player, participants again relied more on the relevant ability, $b = 1.42$, $SE = .10$, $\chi^2 = 206.51$, $p < .001$, than the irrelevant ability, $b = .55$, $SE = .09$, $\chi^2 = 40.38$, $p < .001$, as revealed in a linear contrast of coefficients, $\chi^2 = 39.90$, $p < .001$ (Table S6). When asked to rate impressions of overall intelligence, participants again relied to a similar degree on learning from the verbal context, $b = .54$, $SE = .13$, $\chi^2 = 17.64$, $p < .001$, and the math context, $b = .72$, $SE = .12$, $\chi^2 = 38.86$, $p < .001$ (Table S7).

In contrast, participants did not form conceptual impressions based on reward feedback. Context-specific impressions of player wealth in the task did not vary based on reward feedback in each context, $b = -.05$, $SE = .08$, $\chi^2 = .53$, $p = .47$, and overall impressions of luck did not vary based on reward feedback in the verbal context, $b = -.02$, $SE = .13$, $\chi^2 = .02$, $p = .89$, or reward feedback in the math context, $b = .02$, $SE = .11$, $\chi^2 = .02$, $p = .88$. Thus, participants generated conceptual trait representations based on competence feedback but not reward feedback; this held true across context-dependent representations (impressions of verbal vs. math ability as opposed to verbal vs. math wealth) and global representations (impressions of overall intelligence as opposed to overall luck).

Discussion

Study 2 replicated the results of Study 1: Interactive trait feedback regarding player competence supported context-specific social learning and choice, whereas reward feedback did not. Furthermore, the context-specific effect of trait feedback evident in choice behavior was reflected in trait intelligence ratings of players, again supporting a link between reinforcement learning and conceptual representation of trait impressions.

In addition, Study 2 addressed the possibility that participants implicitly expected trait feedback to be more context specific than reward feedback. To this end, the instructions provided an explicit rationale for why reward values could vary across contexts, rendering expectations about rewards and traits more equivalent. Despite these modified instructions, Study 2 closely replicated Study 1 results, demonstrating again that participants' impressions and

decisions relied primarily on context-dependent trait impressions, with no evidence of reward-based impressions.

Study 3

An important, yet unstudied function of feedback-based trait inference is that, unlike reward, it may have the capacity to generalize appropriately to social decisions in novel situations. Having observed context-dependent learning of trait knowledge through instrumental feedback, we next asked how this learning may be flexibly expressed in new contexts. If participants used simple rules to map contexts to responses (e.g., "pick the person in the blue shirt on math trials"), then they should not be able to generalize learning to relevant novel contexts. In contrast, if trait learning represents a flexible form of knowledge rooted in abstract social cognitive knowledge structures—consistent with the explicit trait ratings observed in Studies 1 and 2—then they should be able to apply learning to novel contexts (e.g., "the person who is skilled at math may also be skilled at science"). Prior research has shown that humans and nonhuman animals often generalize prior learning to novel situations based on the similarity of stimuli (Honig & Urcuioli, 1981; Shepard, 1987). Here, we focused on people's ability to generalize trait learning to novel situations that require conceptually similar abilities, reflecting a form of generalization rooted in cognitive structure (Tenenbaum et al., 2011).

To test generalization of trait knowledge, participants in Study 3 completed additional test phase choices beyond the verbal and math trials reported above. Participants were told that the players had also answered questions regarding six additional skill and knowledge areas (History, Arts, Pattern Recognition, Science, Sports, and Logic). These abilities were selected to represent varying degrees of similarity to math and verbal ability, as determined in an independent pretest.

Method

Pretest Ratings

Thirty participants were recruited on website Mechanical Turk to provide ratings of similarity between various abilities and math/verbal ability. Participants viewed a list of 17 abilities or knowledge areas and rated the similarity of each to verbal and math ability on a scale ranging from 0 (*not at all similar*) to 100 (*extremely similar*). Based on these ratings, a subset of six abilities were chosen for use in the test phases of Studies 3 and 4, including two abilities rated as more similar to verbal than math (History, Arts), two abilities rated as more similar to math than verbal (Science, Pattern Recognition), one rated as similar to neither (Sports), and one rated as similar to both (Logic). For mean ratings, see Table S8.

Participants

Fifty-one New York University undergraduate students (22 males, 29 females, $M_{\text{age}} = 19.86$, $SD = 1.50$) completed Study 3 for course credit. Sample size was determined as in Studies 1 and 2. The study was approved by the New York University Committee on Activities Involving Human Subjects, and informed consent was collected from all participants.

Learning Phase

The learning phase followed an identical protocol to that of Study 1.

Test Phase

Instructions were identical to those of Study 1 with one addition: Participants were informed that Player A participants had also completed questions testing new skills or knowledge areas, including history, science, arts, logic, pattern recognition, and sports. Participants were told they would see rounds featuring these abilities and were asked to do their best to choose on each round.

Participants then completed 192 trials of the test phase, comprising 24 each in the verbal context, math context, and six novel contexts (history, science, arts, logic, pattern recognition, and sports). Trials from each of the eight contexts were pseudorandomly interleaved in different orders for each participant and divided into four blocks of 48 trials. For each trial, a given pair of targets at a given point ratio was randomly assigned to a context.

At the start of each trial, a context cue (e.g., HISTORY) appeared on the screen for 2 s. Following the cue, participants viewed two of the four players; each possible pair of avatars was encountered four times in each context. Each player was equally likely to appear on each side of the screen. As in Study 1, a point pool was displayed above each avatar representing the total number of points available for the question being answered, determined in the same manner.

Posttask Ratings

Participants completed the same posttask ratings of players as in Study 1, and also rated how much they would want to turn to each player for social support and advice (see Supplemental Materials). For exploratory purposes, participants also completed measures of everyday interaction quality (Chiu et al., 1995) and social support appraisals (Vaux et al., 1986), since these measures have been suggested to relate to context-dependent social impressions in past work (Cheng et al., 2001; Chiu et al., 1995). These measures did not relate to task behavior and are not discussed further.

Results

Learning Phase and Test Phase Choices in Verbal/Math Contexts

Analysis of learning phase data and test phase data involving math and verbal contexts were conducted as in Studies 1 and 2. Results from regression and computational analysis closely replicated the results of Studies 1 and 2 and, for brevity, are reported in the supplement (illustrated in Figure 2c and reported in Tables S2–S7).

Novel Test Phase Choices

To address our main question of whether participants generalized learning to novel contexts in the test phase based on similarity to the original math and verbal contexts, we examined choices in novel contexts. Mixed-effects logistic regression was used to predict whether participants chose players as a function of each player's

competencies and reward values from the verbal and math contexts, as well as the similarity of each novel context to verbal and math (as rated in independent pretesting). In this analysis, we could not collapse abilities and reward values from each context into regressors indicating “relevant” and “irrelevant” values, since all contexts were novel. Instead, we separately entered the difference in value between players in each choice pair in (1) point pools, (2) math competence, (3) verbal competence, (4) reward value on math trials, (5) reward value on verbal trials, (6) similarity of the current context to math, (7) similarity of the current context to verbal, and (8–15) interactions of each reward value and each competency with each of the similarity regressors. These interactions indicated the extent to which relying on values learned in a given context depends on the similarity of the learning context to the current context. We initially included all possible random variances, but the model failed to converge. Therefore, we iteratively removed the smallest random slope until the model converged (Barr et al., 2013). This procedure led us to remove eight random variance terms. All other analysis procedures were identical to Study 1, including analysis of posttask ratings.

This analysis revealed first that *global* trait impressions contributed to choices across new contexts, as demonstrated by main effects of verbal competence, $b = .62$, $SE = .09$, $z = 6.74$, $p < .0001$, and math competence, $b = .83$, $SE = .12$, $z = 6.95$, $p < .0001$. However, the hypothesized effect of context similarity also emerged: In contexts more similar to math, participants were more likely to choose math-skilled players, $b = .37$, $SE = .08$, $z = 4.37$, $p < .0001$, and in contexts rated as more similar to verbal, participants more likely to choose verbal-skilled players, $b = .10$, $SE = .03$, $z = 2.96$, $p = .003$. These effects were selective, consistent with a double dissociation: Similarity to the verbal domain did not influence reliance on math ability, $b = -.01$, $SE = .03$, $z = -.36$, $p = .72$, and similarity to the math domain did not influence reliance on verbal ability, $b = -.01$, $SE = .05$, $z = -.24$, $p = .81$. A contrast of coefficients comparing generalization from relevant versus irrelevant domains directly supported a double dissociation, $\chi^2 = 18.20$, $p < .0001$. In other words, trait learning was flexibly and selectively applied to novel contexts based on the similarity of the abilities required in old and novel contexts (Figure 3a). By contrast, learned reward values again had no effects on choice overall, $ps > .36$, or as a function of similarity to original contexts, $ps > .31$ (see Table S9 for all coefficients).

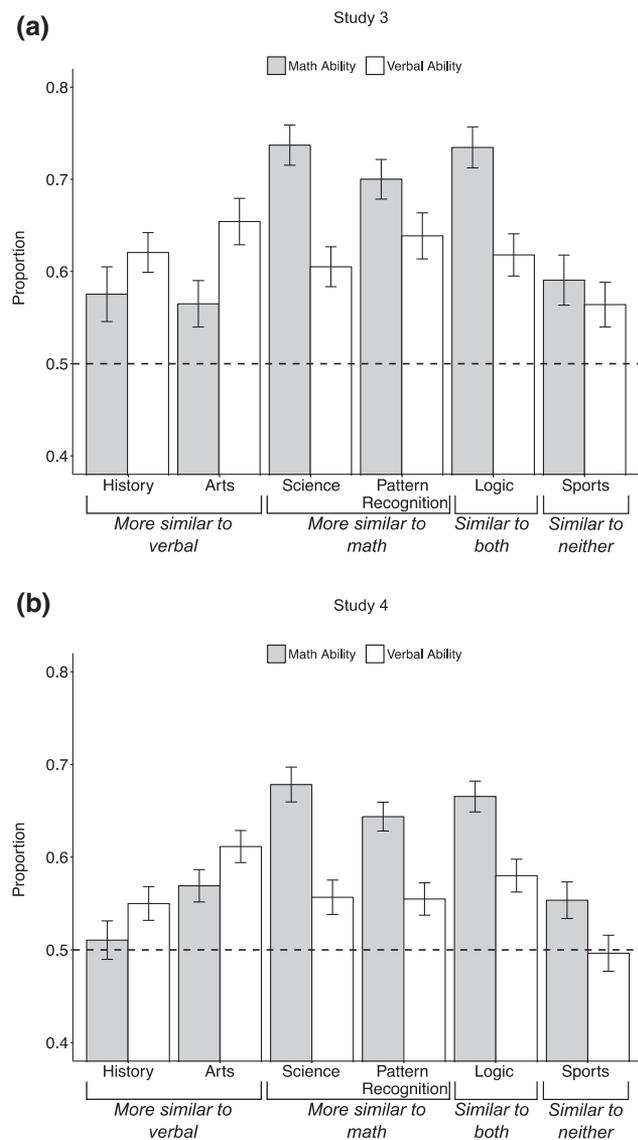
Posttask Ratings

All results related to posttask ratings of partner choice and impressions replicated those reported in Study 1 (Tables S5–S7).

Discussion

We proposed that trait knowledge permits the generalization of past experiences to decisions in novel situations because traits, unlike rewards, can relate to one another semantically in a conceptual map (e.g., Stoller et al., 2020). Thus, in Study 3, we tested the hypothesis that people generalize the context-dependent trait information they learn through interaction and feedback to their choices of partners in novel contexts. Consistent with this hypothesis, we found that participants generalized trait knowledge based on the similarity of traits required in both the initial and novel contexts

Figure 3
Trait-Based Choice in Novel Contexts That Varied in Similarity to Math and Verbal Ability (as Rated by an Independent Sample)



Note. The y-axis indicates the proportion of times participants chose players who were more competent at a given ability (e.g., verbal) than an alternative player, collapsing over the other ability (e.g., math) and collapsing over reward values, in (a) Study 3 and (b) Study 4. Error bars show SEM, adjusted for within-subject comparisons (Morey, 2008). Dashed line indicates chance performance. SEM = standard error of the mean.

(e.g., choosing someone with strong verbal ability to answer a history question). By contrast, reward learning had no discernible impact on choices within the original or novel contexts. This finding reveals an important function of interaction-based trait inference: the capacity to support flexible cross-context decision-making in novel situations. Although prior models of reinforcement learning focus only on reward feedback, we found that reward-based learning could not support appropriate cross-context decision-making.

In addition, Study 3 replicated the pattern of interaction-based learning and choice observed in Studies 1 and 2, such that participants relied primarily on trait-indicating feedback to form context-specific impressions and make context-appropriate choices.

Study 4

Studies 1–3 supported our main hypothesis that context-dependent reinforcement learning relies primarily on traits inference, as compared with reward learning. However, the near complete lack of reward learning effects was intriguing. In prior research in which learning occurred within a single context, social decisions were influenced by both traits and rewards (Hackel et al., 2015, 2020). This discrepancy with past work raises the possibility that reward associations are encoded in environments that offer more consistent experience, such as when interacting with the same people repeatedly in the same context, but that reward associations may be difficult to learn in more complex environments marked by rapid contextual changes. In Studies 1–3, verbal and math trials were randomly interleaved during the learning phase. Although participants were robustly able to track trait-based contingencies across contexts in a trial-by-trial manner, they were unable to track reward contingencies, even though the statistical distributions of reward and trait outcomes were comparably discriminable. To explore this possibility, in Study 4, we used a blocked learning design: Participants learned about players first in the math context and then in the verbal context, or vice versa, with the order of contexts counter-balanced across participants. This design tested whether rewarding outcomes play a role in social choices when experiences of reward are more consistent, and it provided a more stringent test of the hypothesis that trait learning is more context-sensitive than reward learning.

In Study 4, we additionally probed whether generalization of trait learning would extend to preferences for social support—a type of preference even further removed from the original learning contexts. In doing so, we provided a stronger test of the flexibility of trait learning and tested whether our model could predict more naturalistic social preferences that relate to individual well-being (Morelli et al., 2015; Shrout et al., 2006). In our initial independent pretest, participants judged social and emotional skills to be highly similar to verbal ability ($M = 76.24$, $SD = 24.77$) but not math ability ($M = 23.45$, $SD = 24.65$), $t(28) = 9.30$, $p < .001$, $d = 2.14$. Therefore, in Study 4, we tested whether context-dependent trait learning would predict preferences in situations that required empathy, social problem-solving, and analytic skill.¹

Following the learning task, Study 4 participants were asked to report the degree to which they would turn to each player for support in three different situations involving either emotional social

¹ In Study 3, we had asked participants how much they would want to turn to each player for support and advice after a fight with a friend. In those ratings, participants did not differentiate between verbal and math ability; they desired competent partners overall, across verbal and math ability (see Supplemental Analyses). However, that question was broadly phrased; participants could have imagined seeking empathy, practical advice and analysis, or some combination of the above. Therefore, in Study 4, we created more detailed scenarios designed to elicit different needs to different degrees. These scenarios were constructed and validated through pilot testing, which offered initial evidence that the scenarios evoked more context-dependent preferences for verbal and math abilities (see Supplemental Materials).

support, social problem-solving, or practical advice requiring analytic skill. We reasoned that empathy and social problem-solving require high socioemotional skill, whereas analytic advice would primarily require abilities viewed as more similar to math; pilot testing confirmed this to be the case (see Supplemental Materials). Ratings were made using a 7-point Likert-type scale.

Method

Participants

Eighty-five New York University undergraduates (62 females, 23 males, $M_{\text{age}} = 19.56$, $SD = 2.05$) completed the study for course credit. We recruited a larger sample for Study 4 because we wanted to test individual difference correlations between task behavior and posttask ratings (see Supplemental Materials). On the basis of a power analysis assuming a moderate correlation typical of individual difference effects ($r = .3$), we aimed for a sample of at least 84 participants to achieve 80% power. The study was approved by the New York University Committee on Activities Involving Human Subjects, and informed consent was collected from all participants.

Learning Phase

The learning phase was similar to Studies 1–3, except that instead of presenting verbal and math trials in pseudorandomly interleaved order, we presented the two contexts in different blocks. A single context (e.g., verbal) was presented over 84 trials (broken into two blocks), followed by the other context (e.g., math) for another 84 trials (also broken into two blocks). Context order was counter-balanced across participants.

Test Phase

The test phase followed an identical protocol to that of Study 3.

Posttask Ratings

Participants completed the same posttask measures as in Studies 1 and 3. First, however, participants also saw three hypothetical social scenarios presented in randomized order (see Supplemental Materials for the full text of each scenario). In an “empathy” scenario, participants imagined that a boss or professor was rude and unfair to them, and that they wanted to vent to someone without receiving advice. Participants rated how much they would want to turn to each player for empathy using 7-point Likert-type items (1 = *not at all*, 7 = *very much*).

In a “social problem-solving” scenario, participants imagined they had two close friends who had ended a romantic relationship and were no longer speaking. Participants were further asked to imagine that they were having difficulty managing relationships with each friend, that they were getting ready to throw a party, and that they were worried about whom to invite. Participants rated how much they would want to turn to each player for advice, using the same scale.

In a “moving” scenario, participants imagined that they were moving to a new apartment and having a hard time figuring out how to manage the finances and organization of their move. Participants rated how much they would want to turn to each player for advice, using the same scale.

Participants also rated how well they thought they would get along with each player overall, to explore broader social preferences removed from specific contexts, and their own identification with verbal and math ability. (Since these measures did not relate to our primary questions regarding context-dependence, these results are described in Supplemental Materials.)

Results

Reward Learning in Consistent Environments

To examine how reward feedback during the learning phase influenced decision-making, we fit test phase behavior to the same mixed-effects logistic regression described in Studies 1–3 (Figure 2d). Despite the use of a blocked design, the pattern of trait inference replicated Studies 1–3: During math and verbal trials, participants relied on competence in the relevant context during the test phase, $b = .95$, $SE = .10$, $z = 9.27$, $p < .0001$, and, to a lesser extent, in the irrelevant context, $b = .33$, $SE = .08$, $z = 4.00$, $p < .0001$. Again, choice effects based on competence were stronger for relevant as opposed to irrelevant contexts, $\chi^2 = 22.70$, $p < .0001$.

Unlike Studies 1–3, however, we observed a weak effect of reward feedback in the relevant context, $b = .14$, $SE = .05$, $z = 2.51$, $p = .01$. The effect of rewards in the irrelevant context was not significant, $b = .06$, $SE = .04$, $z = 1.46$, $p = .15$ —but the coefficients for relevant and irrelevant contexts did not differ significantly from one another, $\chi^2 = 1.00$, $p = .46$. Although these results may suggest the possibility of context-sensitivity in the reward domain, they do not offer conclusive evidence for context-sensitivity to rewards. This pattern of results was mirrored in the computational modeling of learning, which showed evidence that rewards reinforced behavior to a small degree in both the relevant context, $b = .12$, $SE = .04$, $z = 2.75$, $p = .006$, and marginally in the irrelevant context, $b = .16$, $SE = .08$, $z = 1.93$, $p = .05$, but did not show evidence for context specificity (Table S2).² In other words, people did learn from rewards when they experienced stable contingencies—that is, they repeated interaction choices that led to rewarding outcomes—but this learning was not selective to context.

To directly test our broad hypothesis that trait learning is applied with greater context-dependence than reward learning, we conducted a linear contrast of coefficients comparing the influence of competence and rewards in the test phase as a function of context relevance ([Relevant Competency – Irrelevant Competency] > [Relevant Reward Value – Irrelevant Reward Value]). This contrast indicated significantly greater context-dependence of competence than rewards, $\chi^2 = 13.82$, $p = .0002$. Thus, although participants’ choices in the test phase were influenced by previously experienced rewards following a blocked learning design, participants nevertheless showed more robust context-dependence of traits than rewards. In addition, participants relied more overall on traits learned from each context than reward associations learned from each context, $ps \leq .005$ (Table S4). This pattern held when examining posttask ratings, which revealed small main effects of reward value in addition to replicating the trait-based preferences and impressions

² The lagged regression approach was not suitable to provide a test of context-dependent learning in Study 4; given that contexts were presented in separate blocks, a lagged regression could not distinguish learning from the same versus a different context, whereas analyses of the test phase and computational modeling could do so.

reported in previous studies (Tables S5–S7). These findings demonstrate that trait learning via interaction and feedback was more sensitive to context—even in a situation that offered greater consistency and promoted a degree of reward-based behavior.

In test phase trials featuring novel contexts, we repeated the analysis strategy described in Study 3 to test for generalization of learning to similar abilities. (Following the procedure described in Study 3, we removed 3 random-effect parameters to facilitate model convergence.) These results replicated the generalization results observed in Study 3 (Figure 3b, Table S9): Players who were more competent in math or verbal domains were more likely to be chosen in conceptually similar novel contexts, respectively.

Context-Sensitive Trait Learning Predicts Novel Social Preferences

To extend our analysis of generalization beyond the academic domain, we next examined the generalization of learning to preferences in three social scenarios designed to rely on abilities seen as similar to verbal or math. Posttasks ratings were analyzed using GEE, as in prior studies (see Table S10 for all coefficients). Preference in each scenario served as the dependent variable, and regressors included competence and reward value from each context. We expected that the “empathy” scenario would require socioemotional skills viewed as similar to verbal competence; the “moving” scenario would require analytic skills viewed as similar to math competence; and the “social problem-solving” scenario would rely on both, with a heavier emphasis on socioemotional skills.

Results revealed that in the empathy scenario, preference for seeking support was predicted by verbal ability, $b = .40$, $SE = .08$, $\chi^2 = 27.82$, $p < .001$, but not math ability, $b = -.02$, $SE = .08$, $\chi^2 = .08$, $p = .78$. In the social problem-solving scenario, preference for seeking support was predicted by verbal ability, $b = .51$, $SE = .09$, $\chi^2 = 34.75$, $p < .001$, and, more weakly, by math ability, $b = .20$, $SE = .09$, $\chi^2 = 4.45$, $p = .04$; the coefficient for verbal ability was significantly stronger in a contrast of coefficients, $\chi^2 = 7.87$, $p = .005$. Finally, in the “moving” scenario, preference for seeking support was more strongly predicted by math ability, $b = .74$, $SE = .10$, $\chi^2 = 53.98$, $p < .001$, than by verbal ability, $b = .52$, $SE = .07$, $\chi^2 = 57.12$, $p < .001$; a contrast of coefficients for math versus verbal ability was marginally significant, $\chi^2 = 3.2$, $p = .07$. This pattern of results was consistent with our trait-based generalization hypothesis: Participants relied primarily on verbal ability when rating partners for empathy; relied primarily on verbal ability and secondarily on math ability when rating partners for social problem-solving; and relied primarily on math ability and secondarily on verbal ability for organizational and financial advice.

To directly compare the degree to which participants’ preferences generalized from each context, we computed difference scores indicating their relative reliance on verbal versus math competence in each scenario. First, for each scenario, we computed a measure of “verbal sensitivity” as ratings for *high verbal ability players—low verbal ability players*, collapsing across math ability, and a “math sensitivity” score as ratings for *high math ability players—low math ability players*, collapsing across verbal ability. Then, we computed a difference score (verbal sensitivity—math sensitivity) as an index of “verbal prioritization” for each scenario. This score indicates relative priority given to players specifically skilled at verbal as opposed to math within a given scenario. We entered this score from

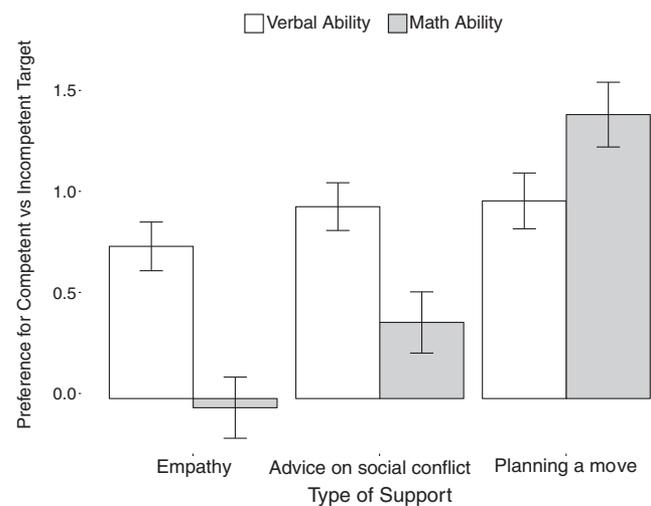
each scenario into a repeated measures analysis of variance (ANCOVA) with scenario type as a factor (empathy, social problem-solving, moving advice), providing a test of whether relative reliance on verbal over math ability differed between scenarios. A corresponding procedure was used to examine sensitivity to math rewards and verbal rewards in the social scenarios.

Results revealed that participants used trait inferences to choose partners flexibly based on the similarity of previous contexts to relevant social needs. A comparison of reliance on verbal as opposed to math ability for each scenario reflected prioritization of verbal ability for the empathy scenario, $M = .83$, $SD = 1.96$, $t(83) = 3.90$, $p < .001$, $d = .43$, and the social problem-solving scenario, $M = .60$, $SD = 2.07$, $t(83) = 2.64$, $p = .01$, $d = .29$, but prioritization of math ability for the moving scenario, $M = -.48$, $SD = 2.29$, $t(83) = -1.90$, $p = .06$, $d = .21$. This pattern of differences was supported by a repeated measures ANOVA testing the difference between verbal and math prioritization for each scenario, $F(2, 168) = 10.42$, $p < .001$, $\eta_p^2 = .11$ (Figure 4a). These results support the proposal that people flexibly apply trait learning to social choices in novel situations, based on the similarity between old and new contexts. By contrast, reward learning did not show a similar pattern of generalization, as revealed in a corresponding ANOVA, $F(2, 168) = .38$, $p = .68$, $\eta_p^2 = .005$ (Table S10, Figure S1). These results suggest that the tendency to learn context-dependent traits may underlie the ability to make flexible social decisions (see also Supplemental Analyses for individual differences).

Discussion

Study 4 clarified and expanded our understanding of reward learning and trait learning in cross-context interactions in two major ways. First, we explored the possibility that the lack of reward learning effects in Studies 1–3 reflected a limitation of reward

Figure 4
Posttask Preferences for Social Support in Study 4



Note. Preference for seeking social support in different scenarios generalized as a function of their context-specific trait competence in math versus verbal domains. Error bars show SEM, adjusted for within-subject comparisons (Morey, 2008). SEM = standard error of the mean.

learning to track rapid changes in context. In support of this account, the blocked context design of Study 4 yielded evidence for reward-based reinforcement learning. Nevertheless, even in a blocked design, there was no clear evidence for context-dependence in reward learning. Second, Study 4 expanded the theoretical reach of context-dependent trait generalization observed in Studies 1–3, which focused on academic domains, to include preferences for seeking social support. Participants generalized the trait inferences formed through direct interaction and feedback based on verbal and math competence to preferences for socioemotional support corresponding to these competencies: They chose partners for empathy based primarily on their verbal ability, but partners for financial advice based on their math ability. Moreover, this pattern of generalization emerged despite a more subtle test in comparison with Study 3; in Study 4, scenarios were not explicitly labeled with different abilities but instead varied implicitly in the extent to which different abilities were required (based on pretesting). Furthermore, these scenarios involved noneconomic social preferences far removed from the points-based competence game where learning occurred. This pattern of generalization was not observed for reward-based learning. Thus, again, we found that while prior models of interactive learning focus on reward association, it is the trait information inferred from feedback that accommodates context specificity and the generalization to novel socially relevant domains.

General Discussion

The impressions we form of others often emerge from the direct interactions we have with them across a variety of situations. Yet previous research in social psychology has not explored the mechanisms through which we form impressions of others through interaction—that is, through our action and another’s feedback—across contexts. By contrast, models of reward reinforcement in cognitive neuroscience pertain to interactive (i.e., feedback-based) learning but focus on the encoding of reward associations and the kinds of trait inferences that typically characterize social impressions. Here, we integrated these approaches to examine the complementary roles of interaction-based reward and trait learning in support of specific impression formation. We found that feedback-based trait inferences, but not reward inferences, were formed in a context-specific manner, and that adaptive social decisions relied primarily on these context-dependent trait inferences.

Across four experiments examining interaction-based learning of rewards and traits between contexts, participants learned primarily from *context-dependent traits* gleaned from social interactions, secondarily from *global traits*, and least of all from *rewards*. Although prior models of reinforcement learning, which supports feedback-based interactive learning, focused only on reward, participants did not form context-specific reward associations. By contrast, our findings suggest that traits, rather than rewards, provide a cognitive basis for forming stable, context-sensitive impressions from feedback in social interactions.

When confronted with novel contexts, participants generalized previous trait learning to contexts requiring conceptually similar traits. This pattern of selective generalization indicates that trait knowledge can be used flexibly to navigate social decisions. By contrast, prior reward learning had little effect on social decisions, despite that the task afforded equivalent opportunities for learning

about both traits and rewards and that participants would have made the most money during learning by relying entirely on reward feedback (see Hackel et al., 2015). Although instrumental reward feedback has been linked to specific contexts in other tasks and is a key feature distinguishing active and passive learning (Bouton & Todd, 2014; Gershman et al., 2015; Niv et al., 2006; Poldrack et al., 2001), reward learning did not predict the complex social behavior studied here. Together, these results suggest that trait learning during social interaction permits more specific and flexible encoding of associations, in comparison with reward associations, which in turn supports more adaptive social decision-making across contexts.

These findings share features with traditional social cognitive models of trait inference under passive conditions—in particular, models of spontaneous trait inference, in which people form specific trait impressions upon perceiving behaviors (Todorov & Uleman, 2003; Winter & Uleman, 1984). Yet, demonstrating this consistency in interactive learning offers an important advance. Models of reinforcement learning suggest that people learn to repeat actions based on concrete rewards, and past work in single contexts has found that rewards do guide learning in social interaction (Hackel et al., 2015). Moreover, research in cognitive neuroscience suggests that active learning through feedback proceeds through distinct mechanisms relative to more passive learning (Poldrack et al., 2001). Without testing interactive learning, it would therefore remain unclear (a) to what extent trait inferences are spontaneously formed during cross-context interactions (as opposed to passive presentation of information) and (b) to what extent these inferences guide cross-context choices when reward feedback is also present. The present findings demonstrate that people do infer context-specific traits from interacting with others, and that these traits, rather than reward feedback, guide complex decisions across contexts. Thus, they demonstrate the relevance and importance of these social cognitive processes for flexible interactive decision-making.

More broadly, these results support the theoretical integration of reinforcement learning and social cognition to understand how people learn about others through the choices they make and feedback they experience during an interaction. Models of reinforcement learning focus on how people learn from rewarding outcomes (Daw et al., 2011), whereas models of social cognition traditionally focus on how people form impressions of other people’s character traits (Uleman & Kressel, 2013). Recent work indicates that neither approach can explain human social learning alone (Hackel et al., 2015, 2020), consistent with the view that multiple processes of learning and memory contribute to social behavior (Amodio, 2019). The present work highlights the complementary functions of different learning processes in social decision-making, illuminating how social learning through interaction and feedback can support context-dependent social behavior.

Trait Learning Across Social Contexts

A key novel contribution of this research is the finding that trait inferences are particularly informative to cross-context social decision-making: Participants learned traits in a context-specific manner, and they effectively and flexibly generalized these trait impressions to novel contexts as a function of trait similarity. These findings highlight the functionality of impression formation for social decision-making. Past work has emphasized that trait learning

offers a more abstract type of feedback than reward learning, revealing a general disposition that may be expressed in multiple specific ways (Hackel et al., 2020; Rim et al., 2009). For instance, a person who is kind can express their kindness through a heartfelt compliment or a thoughtful gift. Trait-indicating feedback in social interactions therefore lets us estimate the value of interacting with another person across specific different situations.

However, the present work highlights a second function of trait knowledge. People often use cognitive structures not only to reason abstractly but also to make sense of new experiences, notice relevant differences between them, and generalize across settings (Gentner & Markman, 1994; Tenenbaum et al., 2011). In the present work, we demonstrate this functionality of conceptual trait knowledge for interactive social choice. People have fine-grained knowledge about traits and are aware of how different traits correlate with one another (Stolier et al., 2020). Trait knowledge can therefore serve as a conceptual map, allowing one to understand why a person's behavior might vary across contexts and to transcend the specifics of a particular encounter to make social choices in novel situations. Indeed, we found that people generalized from trait feedback not only to other competencies within a similar game (Study 3) but also to preferences for social support (Study 4), even though these preferences were far removed from the initial interactions where learning occurred. Thus, by applying trait-based inferences formed through interaction—rooted in semantic memory systems—people can leverage abstract knowledge structures about traits to choose social partners appropriately.

What cognitive and neural processes underlie this form of instrumental trait learning, and how do these differ from reward learning? One possibility is that trait learning in our studies merely reflects the learning of relative rewards—that is, the proportional value of reward (Holroyd et al., 2004; Palminteri et al., 2015). However, several pieces of evidence suggest that trait-based learning studied here involves contributions of social cognition, beyond reward associations, which rely on semantic representations of trait concepts (Amodio, 2019; Gilbert et al., 2012). First, past work has found that trait feedback during a similar task activates neural regions that have been linked to social impression updating when people read about others' behavior (Hackel et al., 2015; Mende-Siedlecki, Cai, & Todorov et al., 2013). Second, past work has found that participants prioritize learning from proportions over absolute rewards particularly when playing with other humans as opposed to nonsocial slot machines (Hackel et al., 2020). This work demonstrates that people place greater emphasis on proportional feedback specifically when it carries social significance regarding a person's traits. Finally, in the present work, participants applied semantic trait representations (e.g., "intelligent") and generalized trait knowledge based on the conceptual similarity of traits required in old and new settings—processes related to conceptual trait knowledge but not relative reward. Altogether, our data, combined with our prior research, strongly suggest that trait learning involves a unique contribution of social cognition beyond any effect of relative reward encoding.

By relying on preexisting knowledge structures to lend meaning to context-dependent behavior, trait-based learning also differs from traditional studies of statistical learning. Humans generally engage in statistical learning to learn regularities in their environments (Aslin & Newport, 2012; Schapiro & Turk-Browne, 2015), which may help people acquire trait concepts and learn how different traits

are associated with each other (Atzil et al., 2018; Stolier et al., 2020). In turn, that cognitive structure becomes a lens that supports further learning: in our task, participants experienced statistically equivalent signals of traits and rewards, but they relied primarily on traits that allowed them to infer meaning from regularities. Indeed, whereas statistical learning is typically passive, implicit, and relies on neural structures in medial temporal lobe (Schapiro & Turk-Browne, 2015), trait-based reinforcement involves learning from motivationally relevant feedback via a larger set of neural structures linked to impression formation and reward (Hackel et al., 2015), which suggests a more active engagement in social meaning-making.

Finally, it is notable that the patterns of interactive feedback-based learning examined in the present studies are not unique to social interactions. Rather, they reflect domain-general processes that have been shown to occur in response to both human and nonhuman agents (Hackel et al., 2015). Indeed, one can infer both the reward value and trait-like qualities of a slot machine, much like for a human, and use conceptual maps to learn and generalize about any kind of stimulus. However, given the complexity of human behavior, the strong propensity to view humans in terms of trait dispositions, and the multifaceted contexts in which human behavior occurs, these interactive learning mechanisms are particularly relevant to human social cognition (Hackel et al., 2020).

Global Versus Context-Dependent Trait Learning in Social Interaction

Although we observed a predominant role for *context-dependent* trait impressions in interactive social learning, we also found a reliable secondary role for *global* trait impressions. That is, to some degree, participants still chose partners based on math ability for a verbal assignment and vice versa. This influence of global impressions could represent a "halo effect" in which people are drawn to individuals seen in a positive light (Nisbett & Wilson, 1977), even when a person's positive qualities are irrelevant to the task at hand. Alternatively, participants might have expected verbal and math ability to be correlated based on real-world experience, even though verbal and math ability were uncorrelated in the present task. Future work could test how such lay beliefs about traits influence learning and choice on the basis of those traits.

Nonetheless, this finding helps reconcile theories of social cognition that have emphasized either global trait impressions (Fiske et al., 2007; Olivola et al., 2014; Rosenberg et al., 1968) or context-dependent impressions and attitudes (Gawronski & Cesario, 2013; Shoda & Mischel, 1993). In particular, we find that both types of learning play a role in social impression formation, consistent with an interactive memory systems model of social cognition (Amodio, 2019; Amodio & Ratner, 2011). Moreover, our findings suggest that people can construct global trait impressions out of distinct experiences with others: participants reported impressions of overall intelligence that combined information from each individual context (i.e., verbal and math ability). Therefore, even when people do report global trait impressions about others, these impressions may draw on context-dependent knowledge, which plays a stronger role in guiding social choices. These findings begin to illuminate the relationship between context-dependent and global trait impressions in learning, choice, and judgment.

Reward Learning Across Social Contexts

Whereas trait information was learned readily and generalized flexibly across contexts, reward learning was relatively weak and a poor predictor of social behavior. This weak effect of reward learning was observed even though the optimal strategy for winning money during the learning phase would have been to focus on reward feedback (Hackel et al., 2015), and despite that, in Study 2, instructions emphasized that rewards would vary in a context-dependent manner and provided a corresponding narrative to let participants make sense of these shifts.

Why were the effects of rewards so weak, relative to traits, in cross-context learning? If trait knowledge provides a conceptual structure to make sense of differences across contexts, then participants may have found it more intuitive to navigate the task by focusing on trait-relevant feedback, relative to reward feedback. Indeed, people often see traits as a useful guide to predicting another's long-term future behavior (Rim et al., 2009)—especially when others' material resources may fluctuate (Raihani & Barclay, 2016)—and people are well practiced in forming coherent impressions of other humans when faced with inconsistencies (Plaks et al., 2003; Read & Miller, 1993). However, participants may not have applied a similar conceptual structure to make sense of variation in reward. Indeed, in Study 2, participants applied semantic trait labels based on trait feedback (e.g., “intelligent”) but did not apply semantic labels based on reward feedback (e.g., “wealthy,” “lucky”). As a result, even when task instructions emphasized the importance of learning from rewards across contexts, as in Study 2, participants may have found reward feedback to be a less natural and intuitive basis for social impressions and choices in complex cross-context situations.

A second, complementary, reason for this pattern concerns the manner in which reward associations are often formed. Reward learning often involves building incremental associations between rewards and actions through consistent experiences, as has been proposed for habitual (or “model-free”) learning systems (Balleine & Dickinson, 1998; Foerde et al., 2006; Gillan et al., 2015; Hackel et al., 2019; Wood & Runger, 2016). However, the context-dependent nature of our task may have prevented this form of reward learning. In prior research, in which reward and trait learning occurred in a single context, reward learning was robust and significantly influenced choice, albeit to a lesser extent than trait learning, even when it was no longer optimal to choose previously rewarding partners (Hackel et al., 2015, 2019). By contrast, in the present research (Studies 1–3), rapid alternations between contexts appeared to impede reward learning and its potential effects on social choices. Only in Study 4, in which contexts were encountered in separate blocks, which produced more coherent reward contingencies, did rewards have a small influence on later social choices. Hence, reward learning may have a more persistent effect on social decisions in more consistent settings that allow people to build more habitual response tendencies (Wood & Runger, 2016).

It is notable that the distinction between trait and reward learning, observed here, may appear similar to the distinction between model-based and model-free learning in the reinforcement learning literature. Indeed, it is possible that trait learning may overlap conceptually with “model-based” reinforcement learning, wherein agents use an internal representation (or model) of the environment to form plans that maximize expected reward.

However, existing accounts of model-based learning are silent on how trait knowledge may be used or generalized in a context-dependent manner; it is possible that social cognitive knowledge structures serve as an internal model of the world used as input to a model-based choice system. At the same time, both model-based and model-free learning can influence reward-based social attitudes and choices (Hackel et al., 2019), and in principle, both can support context-dependent expression of learning (Niv et al., 2006). The potential interplay of model-based/model-free learning processes and reward- and trait-based impressions remains fertile ground for future work on interaction-based social cognition.

Conclusions

Navigating our exceedingly complex social world depends on our ability to maintain consistent representations of other people across constantly changing situations. Our findings reveal how trait knowledge offers a conceptual structure for learning in direct interaction across contexts, and how these context-dependent social impressions provide a flexible basis for social choices. More broadly, the present results suggest that the study of instrumental social learning—understanding how people learn about others through active experience and feedback—offers a fruitful framework for understanding complex forms of human social cognition and decision-making.

References

- Amodio, D. M. (2019). Social cognition 2.0: An interactive memory systems account. *Trends in Cognitive Sciences*, 23(1), 21–33. <https://doi.org/10.1016/j.tics.2018.10.002>
- Amodio, D. M., & Berg, J. J. (2018). Toward a multiple memory systems model of attitudes and social cognition. *Psychological Inquiry*, 29(1), 14–19. <https://doi.org/10.1080/1047840X.2018.1435620>
- Amodio, D. M., & Ratner, K. G. (2011). A memory systems model of implicit social cognition. *Current Directions in Psychological Science*, 20(3), 143–148. <https://doi.org/10.1177/0963721411408562>
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, 41(3), 258–290. <https://doi.org/10.1037/h0055756>
- Aslin, R. N., & Newport, E. L. (2012). Statistical learning: From acquiring specific items to forming general rules. *Current Directions in Psychological Science*, 21(3), 170–176. <https://doi.org/10.1177/0963721412436806>
- Atzil, S., Gao, W., Fradkin, I., & Barrett, L. F. (2018). Growing a social brain. *Nature Human Behaviour*, 2(9), 624–636. <https://doi.org/10.1038/s41562-018-0384-6>
- Baetens, K. L., Ma, N., & Van Overwalle, F. (2017). The dorsal medial prefrontal cortex is recruited by high construal of non-social stimuli. *Frontiers in Behavioral Neuroscience*, 11, Article 44. <https://doi.org/10.3389/fnbeh.2017.00044>
- Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: Contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37(4–5), 407–419. [https://doi.org/10.1016/S0028-3908\(98\)00033-1](https://doi.org/10.1016/S0028-3908(98)00033-1)
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D., Maechler, 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>

- Bendtsen, K. M., Uekermann, F., & Haerter, J. O. (2016). Expert game experiment predicts emergence of trust in professional communication networks. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(43), 12099–12104. <https://doi.org/10.1073/pnas.1511273113>
- Berry, W. D., & Feldman, S. (1985). *Multiple regression in practice*. SAGE Publications. <https://doi.org/10.4135/9781412985208>
- Bhanji, J. P., & Delgado, M. R. (2014). The social brain and reward: Social information processing in the human striatum. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*(1), 61–73. <https://doi.org/10.1002/wcs.1266>
- Boorman, E. D., O'Doherty, J. P., Adolphs, R., & Rangel, A. (2013). The behavioral and neural mechanisms underlying the tracking of expertise. *Neuron*, *80*(6), 1558–1571. <https://doi.org/10.1016/j.neuron.2013.10.024>
- Bouton, M. E., & Todd, T. P. (2014). A fundamental role for context in instrumental learning and extinction. *Behavioural Processes*, *104*, 13–19. <https://doi.org/10.1016/j.beproc.2014.02.012>
- Cheng, C., Chiu, C. Y., Hong, Y. Y., & Cheung, J. S. (2001). Discriminative facility and its role in the perceived quality of interactional experiences. *Journal of Personality*, *69*(5), 765–786. <https://doi.org/10.1111/1467-6494.695163>
- Chiu, C.-Y., Hong, Y.-Y., Mischel, W., & Shoda, Y. (1995). Discriminative facility in social competence: Conditional versus dispositional encoding and monitoring-blunting of information. *Social Cognition*, *13*(1), 49–70. <https://doi.org/10.1521/soco.1995.13.1.49>
- Cornwell, E. Y., & Cornwell, B. (2008). Access to expertise as a form of social capital: An examination of race-and class-based disparities in network ties to experts. *Sociological Perspectives*, *51*(4), 853–876. <https://doi.org/10.1525/sop.2008.51.4.853>
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models. In M. R. Delgado, E. A. Phelps, & T. W. Robbins (Eds.), *Decision making, affect, and learning: Attention and performance XXIII* (pp 3–38). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199600434.003.0001>
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*(6), 1204–1215. <https://doi.org/10.1016/j.neuron.2011.02.027>
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, *18*(2), 185–196. <https://doi.org/10.1016/j.conb.2008.08.003>
- Doll, B. B., Duncan, K. D., Simon, D. A., Shohamy, D., & Daw, N. D. (2015). Model-based choices involve prospective neural activity. *Nature Neuroscience*, *18*(5), 767–772. <https://doi.org/10.1038/nn.3981>
- FeldmanHall, O., Dunsmoor, J. E., Kroes, M. C. W., Lackovic, S., & Phelps, E. A. (2017). Associative learning of social value in dynamic groups. *Psychological Science*, *28*(8), 1160–1170. <https://doi.org/10.1177/0956797617706394>
- Ferrari, S. L. P., & Cribari-Neto, F. (2004). Beta regression for modelling rates and proportions. *Journal of Applied Statistics*, *31*(7), 799–815. <https://doi.org/10.1080/0266476042000214501>
- Fiske, S. T., Cuddy, A. J., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, *11*(2), 77–83. <https://doi.org/10.1016/j.tics.2006.11.005>
- Fitzmaurice, G. M., Laird, N. M., & Ware, J. H. (2004). *Applied longitudinal analysis*. Wiley.
- Foerde, K., Knowlton, B. J., & Poldrack, R. A. (2006). Modulation of competing memory systems by distraction. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(31), 11778–11783. <https://doi.org/10.1073/pnas.0602659103>
- Friesen, C. A., & Kammrath, L. K. (2011). What it pays to know about a close other: The value of if-then personality knowledge in close relationships. *Psychological Science*, *22*(5), 567–571. <https://doi.org/10.1177/0956797611405676>
- Garrison, J., Erdeniz, B., & Done, J. (2013). Prediction error in reinforcement learning: A meta-analysis of neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, *37*(7), 1297–1310. <https://doi.org/10.1016/j.neubiorev.2013.03.023>
- Gawronski, B., & Cesario, J. (2013). Of mice and men: What animal research can tell us about context effects on automatic responses in humans. *Personality and Social Psychology Review*, *17*(2), 187–215. <https://doi.org/10.1177/1088868313480096>
- Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science*, *5*(3), 152–158. <https://doi.org/10.1111/j.1467-9280.1994.tb00652.x>
- Gershman, S. J. (2017). Context-dependent learning and causal structure. *Psychonomic Bulletin & Review*, *24*(2), 557–565. <https://doi.org/10.3758/s13423-016-1110-x>
- Gershman, S. J., Norman, K. A., & Niv, Y. (2015). Discovering latent causes in reinforcement learning. *Current Opinion in Behavioral Sciences*, *5*, 43–50. <https://doi.org/10.1016/j.cobeha.2015.07.007>
- Gilbert, S. J., Swencionis, J. K., & Amodio, D. M. (2012). Evaluative vs. trait representation in intergroup social judgments: Distinct roles of anterior temporal lobe and prefrontal cortex. *Neuropsychologia*, *50*(14), 3600–3611. <https://doi.org/10.1016/j.neuropsychologia.2012.09.002>
- Gillan, C. M., Otto, A. R., Phelps, E. A., & Daw, N. D. (2015). Model-based learning protects against forming habits. *Cognitive, Affective & Behavioral Neuroscience*, *15*(3), 523–536. <https://doi.org/10.3758/s13415-015-0347-6>
- Hackel, L. M., & Amodio, D. M. (2018). Computational neuroscience approaches to social cognition. *Current Opinion in Psychology*, *24*(8), 92–97. <https://doi.org/10.1016/j.copsyc.2018.09.001>
- Hackel, L. M., Berg, J. J., Lindström, B. R., & Amodio, D. M. (2019). Model-based and model-free social cognition: Investigating the role of habit in social attitude formation and choice. *Frontiers in Psychology*, *10*, Article 2592. <https://doi.org/10.3389/fpsyg.2019.02592>
- 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), 1233–1235. <https://doi.org/10.1038/nn.4080>
- Hackel, L. M., Mende-Siedlecki, P., & Amodio, D. M. (2020). Reinforcement learning in social interaction: The distinguishing role of trait inference. *Journal of Experimental Social Psychology*, *88*, Article 103948. <https://doi.org/10.1016/j.jesp.2019.103948>
- Hackel, L. M., & Zaki, J. (2018). Propagation of economic inequality through reciprocity and reputation. *Psychological Science*, *29*(4), 604–613. <https://doi.org/10.1177/0956797617741720>
- Halekoh, U., Højsgaard, S., & Yan, J. (2006). The R package geePack for generalized estimating equations. *Journal of Statistical Software*, *15*(2), 1–11. <https://doi.org/10.18637/jss.v015.i02>
- Hastie, R. (1980). Memory for behavioral information that confirms or contradicts a personality impression. In R. Hastie, T. M. Ostrom, E. B. Ebbesen, R. S. Wyer, D. L. Hamilton, & D. E. Carlston (Eds.), *Person memory: The cognitive basis of social perception* (pp. 155–177). Lawrence Erlbaum.
- Heider, F. (1958). *The psychology of interpersonal relations*. Wiley. <https://doi.org/10.1037/10628-000>
- Højsgaard, S., & Halekoh, U. (2016). *doBy: Groupwise statistics, LSmeans, linear contrasts, utilities* (R package Version 4.5-15) [Computer software]. <http://CRAN.R-project.org/package=doBy>
- Holroyd, C. B., Larsen, J. T., & Cohen, J. D. (2004). Context dependence of the event-related brain potential associated with reward and punishment. *Psychophysiology*, *41*(2), 245–253. <https://doi.org/10.1111/j.1469-8986.2004.00152.x>
- Honig, W. K., & Urquioli, P. J. (1981). The legacy of Guttman and Kalish (1956): Twenty-five years of research on stimulus generalization. *Journal of the Experimental Analysis of Behavior*, *36*(3), 405–445. <https://doi.org/10.1901/jeab.1981.36-405>
- Jones, E. E., & Nisbett, R. E. (1987). The actor and the observer: Divergent perceptions of the causes of behavior. In E. E. Jones, D. E. Kanouse, H. H.

- Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 79–94). Lawrence Erlbaum.
- Kammrath, L. K., Mendoza-Denton, R., & Mischel, W. (2005). Incorporating if . . . then . . . personality signatures in person perception: Beyond the person-situation dichotomy. *Journal of Personality and Social Psychology*, 88(4), 605–618. <https://doi.org/10.1037/0022-3514.88.4.605>
- Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation*. University of Nebraska Press.
- Kleiner, M., Brainard, D. H., Pelli, D. G., Broussard, C., Wolf, T., & Niehorster, D. (2007). What's new in Psychtoolbox-3? *Perception*, 36(ECVP Abstract Suppl), 14.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). *lmerTest: Tests for random and fixed effects for linear mixed effect models (lme4 objects of lme4 package)* (R package Version 2.0-32) [Computer software]. <https://CRAN.R-project.org/package=lmerTest>
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22(6), 1027–1038. <https://doi.org/10.1016/j.conb.2012.06.001>
- Liang, K.-Y., & Zeger, S. (1986). Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1), 13–22. <https://doi.org/10.1093/biomet/73.1.13>
- Lindström, B., Selbing, I., Molapour, T., & Olsson, A. (2014). Racial bias shapes social reinforcement learning. *Psychological Science*, 25(3), 711–719. <https://doi.org/10.1177/0956797613514093>
- Lüdtke, D., Ben-Shachar, M. S., Patil, I., & Makowski, D. (2020). Extracting, computing and exploring the parameters of statistical models using R. *Journal of Open Source Software*, 5(53), Article 2445. <https://doi.org/10.21105/joss.02445>
- MacKinnon, J. G., & White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3), 305–325. [https://doi.org/10.1016/0304-4076\(85\)90158-7](https://doi.org/10.1016/0304-4076(85)90158-7)
- Maisel, N. C., & Gable, S. L. (2009). The paradox of received social support: The importance of responsiveness. *Psychological Science*, 20(8), 928–932. <https://doi.org/10.1111/j.1467-9280.2009.02388.x>
- Mende-Siedlecki, P., Baron, S. G., & Todorov, A. (2013). Diagnostic value underlies asymmetric updating of impressions in the morality and ability domains. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(50), 19406–19415. <https://doi.org/10.1523/JNEUROSCI.2334-13.2013>
- Mende-Siedlecki, P., Cai, Y., & Todorov, A. (2013). The neural dynamics of updating person impressions. *Social Cognitive and Affective Neuroscience*, 8(6), 623–631. <https://doi.org/10.1093/scan/nss040>
- Mende-Siedlecki, P., & Todorov, A. (2016). Neural dissociations between meaningful and mere inconsistency in impression updating. *Social Cognitive and Affective Neuroscience*, 11(9), 1489–1500. <https://doi.org/10.1093/scan/nsw058>
- Morelli, S. A., Lee, I. A., Armn, M. E., & Zaki, J. (2015). Emotional and instrumental support provision interact to predict well-being. *Emotion*, 15(4), 484–493. <https://doi.org/10.1037/emo0000084>
- Morelli, S. A., Ong, D. C., Makati, R., Jackson, M. O., & Zaki, J. (2017). Empathy and well-being correlate with centrality in different social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(37), 9843–9847. <https://doi.org/10.1073/pnas.1702155114>
- Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, 4(2), 61–64. <https://doi.org/10.20982/tqmp.04.2.p061>
- Nisbett, R. E., & Wilson, T. D. (1977). The halo effect: Evidence for unconscious alteration of judgments. *Journal of Personality and Social Psychology*, 35(4), 250–256. <https://doi.org/10.1037/0022-3514.35.4.250>
- Niv, Y., Joel, D., & Dayan, P. (2006). A normative perspective on motivation. *Trends in Cognitive Sciences*, 10(8), 375–381. <https://doi.org/10.1016/j.tics.2006.06.010>
- Olivola, C. Y., Funk, F., & Todorov, A. (2014). Social attributions from faces bias human choices. *Trends in Cognitive Sciences*, 18(11), 566–570. <https://doi.org/10.1016/j.tics.2014.09.007>
- Otto, A. R., Gershman, S. J., Markman, A. B., & Daw, N. D. (2013). The curse of planning: Dissecting multiple reinforcement-learning systems by taxing the central executive. *Psychological Science*, 24(5), 751–761. <https://doi.org/10.1177/0956797612463080>
- Palminteri, S., Khamassi, M., Joffily, M., & Coricelli, G. (2015). Contextual modulation of value signals in reward and punishment learning. *Nature Communications*, 6, Article 8096. <https://doi.org/10.1038/ncomms9096>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442. <https://doi.org/10.1163/156856897X00366>
- Plaks, J. E., & Higgins, E. T. (2000). Pragmatic use of stereotyping in teamwork: Social loafing and compensation as a function of inferred partner–situation fit. *Journal of Personality and Social Psychology*, 79(6), 962–974. <https://doi.org/10.1037/0022-3514.79.6.962>
- Plaks, J. E., Shafer, J. L., & Shoda, Y. (2003). Perceiving individuals and groups as coherent: How do perceivers make sense of variable behavior? *Social Cognition*, 21(1), 26–60. <https://doi.org/10.1521/soco.21.1.26.21191>
- Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, 414(6863), 546–550. <https://doi.org/10.1038/35107080>
- Raihani, N. J., & Barclay, P. (2016). Exploring the trade-off between quality and fairness in human partner choice. *Royal Society Open Science*, 3(11), Article 160510. <https://doi.org/10.1098/rsos.160510>
- R Development Core Team. (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Read, S. J., & Miller, L. C. (1993). Rapist or “regular guy”: Explanatory coherence in the construction of mental models of others. *Personality and Social Psychology Bulletin*, 19(5), 526–540. <https://doi.org/10.1177/0146167293195005>
- Rim, S., Uleman, J. S., & Trope, Y. (2009). Spontaneous trait inference and construal level theory: Psychological distance increases nonconscious trait thinking. *Journal of Experimental Social Psychology*, 45(5), 1088–1097. <https://doi.org/10.1016/j.jesp.2009.06.015>
- Rosenberg, S., Nelson, C., & Vivekananthan, P. S. (1968). A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology*, 9(4), 283–294. <https://doi.org/10.1037/h0026086>
- Rydell, R. J., Rydell, R. J., Gawronski, B., & Gawronski, B. (2009). I like you, I like you not: Understanding the formation of context-dependent automatic attitudes. *Cognition and Emotion*, 23(6), 1118–1152. <https://doi.org/10.1080/02699930802355255>
- Schapiro, A., & Turk-Browne, N. (2015). Statistical learning. *Brain Mapping*, 3, 501–506. <https://doi.org/10.1016/B978-0-12-397025-1.00276-1>
- Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human orbitofrontal cortex represents a cognitive map of state space. *Neuron*, 91(6), 1402–1412. <https://doi.org/10.1016/j.neuron.2016.08.019>
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237(4820), 1317–1323. <https://doi.org/10.1126/science.3629243>
- Sherer, M., Maddux, J. E., Mercandante, B., Prentice-Dunn, S., Jacobs, B., & Rogers, R. W. (1982). The self-efficacy scale: Construction and validation. *Psychological Reports*, 51(2), 663–671. <https://doi.org/10.2466/pr0.1982.51.2.663>
- Shoda, Y., & Mischel, W. (1993). Cognitive social approach to dispositional inferences: What if the perceiver is a cognitive social theorist? *Personality and Social Psychology Bulletin*, 19(5), 574–585. <https://doi.org/10.1177/0146167293195009>
- Shoda, Y., Mischel, W., & Wright, J. C. (1993). Links between personality judgments and contextualized behavior patterns: Situation-behavior

- profiles of personality prototypes. *Social Cognition*, 11(4), 399–429. <https://doi.org/10.1521/soco.1993.11.4.399>
- Shrout, P. E., Herman, C. M., & Bolger, N. (2006). The costs and benefits of practical and emotional support on adjustment: A daily diary study of couples experiencing acute stress. *Personal Relationships*, 13(1), 115–134. <https://doi.org/10.1111/j.1475-6811.2006.00108.x>
- Soto, F. A., Gershman, S. J., & Niv, Y. (2014). Explaining compound generalization in associative and causal learning through rational principles of dimensional generalization. *Psychological Review*, 121(3), 526–558. <https://doi.org/10.1037/a0037018>
- Spunt, R. P., & Adolphs, R. (2015). Folk explanations of behavior: A specialized use of a domain-general mechanism. *Psychological Science*, 26(6), 724–736. <https://doi.org/10.1177/0956797615569002>
- 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. <https://doi.org/10.1038/s41562-019-0800-6>
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press.
- Tavares, R. M., Mendelsohn, A., Grossman, Y., Williams, C. H., Shapiro, M., Trope, Y., & Schiller, D. (2015). A map for social navigation in the human brain. *Neuron*, 87(1), 231–243. <https://doi.org/10.1016/j.neuron.2015.06.011>
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285. <https://doi.org/10.1126/science.1192788>
- Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. Macmillan. <https://doi.org/10.5962/bhl.title.55072>
- Todorov, A., & Uleman, J. S. (2003). The efficiency of binding spontaneous trait inferences to actors' faces. *Journal of Experimental Social Psychology*, 39(6), 549–562. [https://doi.org/10.1016/S0022-1031\(03\)00059-3](https://doi.org/10.1016/S0022-1031(03)00059-3)
- Trask, S., & Bouton, M. E. (2014). Contextual control of operant behavior: Evidence for hierarchical associations in instrumental learning. *Learning & Behavior*, 42(3), 281–288. <https://doi.org/10.3758/s13420-014-0145-y>
- Uleman, J. S., & Kressel, L. M. (2013). A brief history of theory and research on impression formation. In D. E. Carlston (Ed.), *Oxford handbook of social cognition* (pp. 53–73). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199730018.013.0004>
- Vaux, A., Phillips, J., Holly, L., Thomson, B., Williams, D., & Stewart, D. (1986). The social support appraisals (SS-A) scale: Studies of reliability and validity. *American Journal of Community Psychology*, 14(2), 195–218. <https://doi.org/10.1007/BF00911821>
- Waytz, A., Heafner, J., & Epley, N. (2014). The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of Experimental Social Psychology*, 52, 113–117. <https://doi.org/10.1016/j.jesp.2014.01.005>
- Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of trait inferences. *Journal of Personality and Social Psychology*, 47(2), 237–252. <https://doi.org/10.1037/0022-3514.47.2.237>
- Wojciszke, B. (2005). Morality and competence in person- and self-perception. *European Review of Social Psychology*, 16(1), 155–188. <https://doi.org/10.1080/10463280500229619>
- Wood, W., & Rünger, D. (2016). Psychology of habit. *Annual Review of Psychology*, 67(1), 289–314. <https://doi.org/10.1146/annurev-psych-122414-033417>

Received June 1, 2020

Revision received October 12, 2021

Accepted October 21, 2021 ■