



Published in final edited form as:

Psychol Sci. 2017 August ; 28(8): 1160–1170. doi:10.1177/0956797617706394.

Associative learning of social value in dynamic groups

Oriel FeldmanHall^a, Joseph E. Dunsmoor^{b,*}, Marijn C.W. Kroes^{b,*}, Sandra Lackovic^b, and Elizabeth A. Phelps^{b,c,d}

^aDepartment of Cognitive, Linguistic & Psychological Sciences, Brown University, Providence, RI

^bDepartment of Psychology, New York University, New York NY 10003

^cCenter for Neural Science, New York University, New York NY 10003

^dNathan Kline Institute, Orangeburg, NY 10962

Abstract

Despite humans living in societies that regularly demand engaging with multiple people simultaneously, we know little about social learning in group settings. In two experiments, we combine a Pavlovian learning framework with dyadic economic games to test whether blocking mechanisms support value-based social learning in the gain (altruistic dictators) and loss (greedy robbers) domains. Participants first learn about individual dictators. In a second task, dictators make splits collectively with a partner. Results reveal that since the presence of the dictator already predicts the outcome, participants do not learn to associate value with the partner. This social blocking effect was not observed in the loss domain: robbers' partners who could steal the participant's money but refrained from doing so acquired highly positive value—biasing subsequent behavior. These findings reveal how Pavlovian mechanisms support efficient social learning, while also elucidating that violations of social expectations can attenuate how readily this mechanism is recruited.

Keywords

social learning; Pavlovian blocking; social value; classical conditioning

INTRODUCTION

Successfully navigating through our complex and large social world requires constant assessments of whether social interactions produce rewarding outcomes. Individuals must routinely learn whether a person can be trusted, is dependable, or should be cooperated with—oftentimes while engaging with multiple people at once. Research within the non-social domain illustrates that both humans and animals are highly adept at learning from reward

Corresponding Author: Elizabeth A. Phelps, 6 Washington Place, New York University, New York NY 10003, liz.phelps@nyu.edu, 212-998-8317.

*denotes equal contribution

Author Contributions

OFH, JED, & MCWK developed and designed the experiment. OFH & SLA conducted data collection. OFH, JED, MCWK, & EAP performed data analysis and wrote the manuscript.

and punishment contingencies, and regularly exhibit value-based decision-making (Schultz et al. 1997, O'Doherty et al. 2001, O'doherty 2004, Rangel et al. 2008, Glimcher 2009). This work has resulted in a well-characterized account of how associative learning mechanisms underpin value-based decision-making (Pavlov 1927, Rescorla and Wagner 1972, Sutton and Barto 1998). However, much less is known about how value-based learning occurs within the social domain and during group settings (Ruff and Fehr 2014)—despite evidence that optimal social decision-making is fundamentally dependent on the actions of others (Rilling et al. 2008). Here we ask the critical questions of how humans learn social value about others in dynamic group environments, and whether these value representations influence subsequent social choice.

Imagine encountering an individual who repeatedly demonstrates she is trustworthy. The knowledge gained through simply associating this individual with trustworthy outcomes promotes continued trust over ensuing encounters (Mckelvey and Palfrey 1992). Classic associative learning accounts (Vurbic and Bouton 2014) can be used to explain how direct and repetitive experiences influence behavior (Klucharev et al. 2009), including choosing to trust someone who has proven to be highly trustworthy (King-Casas et al. 2005, Phan et al. 2010). Continuous reinforcement helps explain how social behaviors can be learned over time.

However, in our large and ever-changing social world, the conditions under which one gleans information is rarely limited to isolated, repeated interactions with the same individual. Rather, one often learns about individuals in the company of others, requiring simultaneous social evaluations of each individual. Equally likely, one might initially meet an individual alone, and later again amongst friends, requiring evaluations to be updated if outcomes change depending on the context. For example, if we encounter an individual who treats us kindly, and then reencountering the same individual in the company of a stranger who together exhibit kindness, have we learned the stranger is kind? By leveraging an associative learning framework to examine these questions, we probe whether people bind social value to others in a group setting, update these values when conditions change, and apply learned value associations across social domains to make adaptive choices.

Following in the tradition of human causal judgment research (Lovibond et al. 2003), in our experimental structure, participants play a series of dyadic Dictator games where they can learn through interactions whether dictators are characteristically altruistic or selfish—which can be considered a form of associative conditioning. Participants engage with many dictators, some of which first make offers alone (e.g. altruistic dictator) and then later alongside another new dictator (e.g. altruistic dictator's partner). In an ensuing Trust game, participants decide how much of their own money to entrust in each dictator, as well as novel, never before seen strangers.

By giving participants the opportunity to entrust their own money in each dictator, we can specifically test whether 1) participants have learned to associate social value to specific dictators, even those who never made offers by themselves and only as the dictator's partner, and 2) whether value acquired in one social domain—such as altruism—influences subsequent behavior across other domains—such as trust. Such a framework allows us to

test multiple competing hypotheses for how social learning occurs in dynamic groups. Assuming the dictator's behavior is consistent between making decisions alone and later in compound with a partner (for instance, altruistic splits are always offered to the participant), it is possible that the dictator's partner acquires the same value as the dictator, since both are yoked to the same positive outcome. This social category learning account (Kahneman and Miller 1986, Linville et al. 1989, Kashima et al. 2000) posits that the dictator's partner obtains value because he is associated with a dictator previously known to be altruistic and together they continue to be altruistic. In this case, the value learned about the dictator is transferred to his partner since both activate similar exemplar representations (Smith and Zarate 1992).

An alternative hypothesis drawn from cognitive psychology proposes that since the monetary split continues to be altruistic, participants will not learn anything about the dictator's partner, since the partner provides no new information. Such a phenomenon is known as blocking (Kamin 1969, Rescorla and Wagner 1972)—a basic, albeit fundamental Pavlovian learning mechanism that is robustly observed across species within non-social environments. By this account, the altruistic dictator would block his partner from acquiring social value, and participants should consequently trust the altruistic dictator's partner as they would trust a stranger, despite having direct experience of the partner's altruism.

Although blocking has been used to explain efficient error-driven learning (Gluck and Bower 1988), it is not always the case that people ignore additional cues when learning about novel stimuli (Dickinson et al. 1984). For example, blocking is not observed in some category learning tasks whereby people learn more about a stimulus than is necessary to perform a classification (Bott et al. 2007). In fact, prior experiences with a stimulus or its outcome can modify or even reverse the effectiveness of blocking (Dickinson et al. 1976, Le Pelley 2004). This modulation of blocking may also extend to social learning, such that discrete learning dynamics critically determines how much—or little—one learns in the social domain.

MATERIALS AND METHODS

Task Procedures

To test these competing learning accounts, we employed a series of social economic games to interrogate learning within the gain domain. Inspired by classical conditioning paradigms which assess compound pair cues and their associated outcomes (Lovibond et al. 2003), in Experiment 1 participants completed four tasks in the following order (Run 1, Fig 1A): (i) Dictator Run 1 in which all subjects learn about an altruistic dictator (stimulus A) and a selfish dictator (stimulus I); (ii) Dictator Run 2, in which subjects are again exposed to these dictators but this time in compound with partners (A is paired with test stimuli B and I is paired with test stimuli J); (iii) a Trust Game to test whether the dictators partners' social value was learned; and finally, (iv) a surprise memory test in which episodic memory for all dictators and partners is explicitly probed.

Dictator Run 1—In the first Dictator Game (initial conditioning), subjects learn through repeated interactions about dictators who consistently make either altruistic splits [stimulus

A: approximately \$4 out of \$10, where participants are explicitly told that splits above \$5 do not occur] or selfish splits [Stimulus I: approximately \$.18 out of \$10]. In Run 1 altruistic and selfish dictators make splits alone. Once a split is made, participants rate how they feel on a 5-point visual analogue scale (VAS).

Dictator Run 2—In a second Dictator Game, participants re-encounter the same dictators (a within-subjects design), only this time some of the previous dictators are paired with a never before seen partner to form a compound cue. For instance, the altruistic dictator—who makes altruistic splits by himself in Run 1—also makes altruistic splits when collectively deciding how to split the money alongside the altruistic dictator’s partner in Run 2 (Stimuli A/B: participants are told that dictator pairs jointly and equally contribute to deciding how to split the money, see SI). Accordingly, the altruistic dictator’s partner is paired with a positive outcome, and thus has the potential to acquire positive social value.

Critically, participants play with various other dictators and dictator pairs across both runs who either serve as distractors so that learning is non-trivial (i.e. not too easy), or as comparisons to evaluate the magnitude of how well social value is learned (Lovibond, Been et al. 2003). For instance, distractor dictators always make splits alone and are present in both Runs 1 and 2 (stimuli E and F), and dictator pairs who are only presented in Run 2 serve as a test for whether learning can occur from just the second run (stimuli pairs C/D and K/L: see Fig 1E for a full description of all dictators in both runs). Effectively, such a task structure enables us to temporally manipulate information regarding the social value of others in order to examine how value is learned in complex and dynamic group environments.

Trust Game—In the subsequent test phase, participants play a Trust Game and are given the opportunity to entrust their own money with each dictator (dictators are always presented separately during the Trust Game; Fig 1C). Participants are endowed with \$10 and can choose how much money to entrust, knowing that whatever is transferred will be multiplied four times. Participants are aware that second movers in the Trust Game—dictators and partners from Runs 1 and 2—can either keep the transferred money, leaving the participant with nothing, or reciprocate by sharing back half the increased sum (this feedback is not revealed until after the game, when one trial is randomly selected to be paid out). To get a baseline measurement of willingness to trust, strangers with no prior positive or negative associations from Runs 1 and 2 are also presented as second movers in the Trust Game (stimuli M and N). In this way, we can examine whether previous learning from Runs 1–2 biases how participants treat each individual in the Trust Game, such that dictators who exhibited altruistic behavior should be trusted with greater sums of money than selfish dictators, or even strangers who have no previous associations.

Memory Test—Finally, to test whether participants explicitly remembered the social value associated with each dictator (including all partners), we probe episodic memory for dictators and their associated splits in a surprise memory test (measured with a \$1–\$5 VAS with \$.01 increments; see SI).

Participants

In Experiments 1 and 2, 45 participants (sample size based on extant research using the same Dictator Task (Murty et al. 2016) as well as classic human blocking research (Beckers, De Houwer, Pineno, & Miller, 2005)) were recruited for each study from New York University and the surrounding New York City community (Experiment 1: 27 females, mean age 21.5, $SD \pm 2.9$; Experiment 2: 22 females, mean age 22.1, $SD \pm 3.1$). Participants were paid an initial \$15 and received additional compensation based on the result of one randomly selected trial from the Dictator Game and one randomly selected trial from the Trust Game (up to \$25). Informed consent was obtained from each participant in a manner approved by the University Committee on Activities Involving Human Subjects.

Given that players in both experiments were actually computer algorithms yoked to predetermined reinforcement rates (e.g. highly altruistic or selfish), participants underwent a deception manipulation to create a realistic socially dyadic environment. This included participants being photographed in front of a white wall and told that their picture, along with their responses to “how much of \$10 would you split with a future player?” would be used for the next experiment with future participants. The same was done for the Trust Game. This was explained as the most efficient way to feed forward participants’ responses so that multiple people do not need to come in for an experimental session at once. Participants were told that in the event that their decisions as the dictator (or second mover in the Trust Game) were used in subsequent experimental sessions, they would be mailed a check based on that specific decision. Extensive post task debriefing procedures revealed that participants believed the social manipulation (see supplement for further details).

RESULTS

Experiment 1

Behavioral Results—Results from Experiment 1 favor the associative learning account where the altruistic dictator’s partner failed to acquire positive social value in Run 2, presumably because the altruistic dictator previously predicted the same positive outcome when presented alone in Run 1 (rmANOVA for money entrusted in each altruistic dictator, altruistic dictator’s partner, altruistic pair, and Novel: $F(3,132)=4.28$, $p=0.006$, partial $\eta^2=.113$; Fig 2). Participants trusted the altruistic dictator with the most money (stimulus A: \$5.06 $SD \pm 2.7$), and significantly more than what was entrusted to the altruistic dictator’s partner (stimulus B: \$4.06, $SD \pm 2.7$; paired t-test A–B: $t(44)=2.40$, $p=.023$, Fig 2). Instead, the altruistic dictator’s partner was entrusted with the same amount of money as a stranger (Novel: \$3.95 $SD \pm 2.6$; paired t-test B–Novel: $t(44)=-.44$, $p=.667$). Control tests probing whether the failure to associate the altruistic dictator’s partner with positive social value was attributed to a general failure in learning revealed no such effect (see supplement for details of all control tests across experiments).

The same pattern was observed for selfish dictators who kept most of the money (rmANOVA for money entrusted in each selfish dictator, selfish dictator’s partner, selfish pair, and Novel: $F(3,132)=3.10$, $p=0.029$, partial $\eta^2=.07$; Fig 2). The selfish dictator, who singularly made selfish splits in Run 1, and then collectively made selfish splits with his

partner in Run 2, was entrusted with the least amount of money (stimulus I: \$3.42, $SD \pm 2.8$), and significantly less than what was entrusted to the selfish dictator's partner (stimulus J: \$3.93, $SD \pm 2.6$; paired t-test I–J: $t(44) = -2.22$, $p = .031$). Rather, the selfish dictator's partner was entrusted with effectively the same amount as a stranger (Novel: \$3.95 $SD \pm 2.6$; paired t-test J–Novel: $t(44) = -.13$, $p = .895$).

At first blush it may appear surprising that participants did not learn the social value of either the altruistic or selfish dictator's partners, despite having directly experienced positive or negative outcomes in their presence. However, according to classic Pavlovian learning theory (Rescorla and Wagner 1972), when a stimulus such as an altruistic dictator is already associated with positive outcomes, later encountering both the altruistic dictator and his partner together results in the partner acquiring no social value—since the positive outcome is already fully predicted by the presence of the initial altruistic dictator. This phenomenon results in the previously learned associations of the dictator interfering with (i.e. blocking) the participant's ability to form an association with the dictator's partner (Kamin 1969). Such a finding illustrates that social value learning relies on a difference between the expectation of the outcome and the actual outcome (i.e. prediction errors). That is, if a single stimulus already predicts a positive outcome, an added stimulus could be considered redundant, thus failing to acquire associative value.

Memory Results—To understand whether this behavioral blocking phenomenon was due to a failure in explicitly associating the altruistic dictator's partner with altruism and the selfish dictator's partner with selfishness, we examined how accurately participants remembered each of the dictators' offers from Runs 1–2 (accuracy computed by taking the absolute difference between remembered split and actual split in the surprise memory test). For altruistic dictators, participants more accurately remembered splits from the altruistic dictator and the dictator pair compared to the altruistic dictator's partner (rmANOVA: $F(2,50) = 7.12$, $p = .002$, $\eta^2 = .22$, post-hoc tests against the altruistic dictator's partner reveal all $Ps < .01$), despite the fact that *all* altruistic dictators gave the same distribution of splits. A similar pattern was observed for the selfish dictators, where participants more accurately remembered the selfish dictator and the selfish pair, although the test failed to reach significance (rmANOVA: $F(2,56) = 1.98$, $p = .148$). Thus, it seems that the behavioral blocking effect observed for the dictators' partners during the Trust game may be attributed to blocked episodic memory of the dictators partners' offers.

Evidence of a Pavlovian blocking mechanism—from both a behavioral and memory perspective—makes a case for efficient learning in social contexts (Seid-Fatemi and Tobler 2015), where people who appear to add no critical information fail to become associated with social value. Indeed, it is well established that blocking occurs when no new information about the reward probability is elicited from the learning episode, indicating that the psychological process of surprise is a critical feature of learning. While typical behavior in a Dictator game is to give on average 30% of the monetary pie, prior research reveals that altruistic behavior is quite variable and there are many demonstrations of dictators behaving selfishly and keeping large portions of the pie (Engel 2011). This variable behavior—which spans altruistic benevolence to selfish enhancement—is observed across tasks and cultures, contradicting traditional economic models that tout dominant rational behavior (e.g. to share

nothing and keep all the money). The psychological explanation for offering any sized split in the Dictator game (since a dictator can keep all the money without consequence) hinges on the notion that societies have norms which govern wealth sharing: people routinely make some effort to trade-off their material benefit to comply with social norms of sharing with others (Bolton et al. 1998, Dreber et al. 2013). This suggests that both selfish (small) and altruistic (large) splits in the Dictator game operate within a framework of socially expected behavior. Accordingly, so long as the dictators' behavior is normatively expected, blocking effects for an associated partner are facilitated, since the environment leaves little room for surprise.

Experiment 2

A traditional Pavlovian account, however, is agnostic to changes in stimuli valence, which may consequently fail to capture all aspects of social learning. For example, if we consider the various ways in which framing has profound effects on choice (Tversky and Kahneman 1981, Kahneman and Tversky 1984) there may be certain contexts (Kahneman and Tversky 1979, Fox et al. 2008) that give rise to surprise, which should in turn reveal unblocking effects (Dickinson et al. 1976). Applied to social phenomena, if learning is thought to be sensitive to prior beliefs and inferences about the world, behavior that deviates from socially normative expectations may enable learning. For example, one of the oldest proverbs in moral philosophy states that if given the opportunity to steal without consequence, a person invariably steals (Plato 1950), and there is growing evidence of this effect across psychological domains (Greenberg 1993, Mazar et al. 2008, Greene and Paxton 2009). Violations of this social expectation (e.g. robbers who steal little) would be surprising, and unblocking should occur for individuals who behave in such unexpected ways. In Experiment 2 we tested the possibility that blocking is sensitive to the framing of social contexts, theorizing that prior expectancies—beliefs and inferences about future social events—will influence learning such that seemingly 'redundant' individuals result in acquiring social value.

Accordingly, Experiment 2 mirrored the structure of the first experiment with two key differences. First, prior to playing any games, participants completed a short math task in which they could earn up to \$5. Second, instead of playing a series of Dictator Games, participants played a series of Robbery Games. These games were structurally identical to the Dictator Games except that robbers could steal up to \$5 from the participant (rather than receive up to \$5 in the Dictator Game). Robbers were the same as those described in Experiment 1, except that in the Robbery Games kind robbers 'stole' very little (approximately \$0.18), while greedy robbers stole most of the participants' money (approximately \$4). Correspondingly, the kind robber, who is initially presented alone in Run 1, continues to steal tiny amounts when presented as a pair with his partner in Run 2. In contrast, the greedy robber single-handedly steals most of the participant's money in Run 1, and continues to steal most of the participant's money when paired with his partner in Run 2. As before, participants played a Trust Game with each of the robbers, as well as with strangers, before completing a surprise memory test. Thus, Experiment 2—which probes whether blocking of social value also occurs in the loss domain with different social expectations—is structurally and monetarily matched to Experiment 1.

Behavioral Results—With the possibility of being robbed of their earned money, participants no longer failed to associate social value to the kind robber’s partner (rmANOVA for money entrusted to kind robber, kind robber’s partner, kind pair, and Novel: $F(3,132)=9.62$, $p<0.001$, partial $\eta^2=.26$). Instead, the kind robber’s partner was trusted with significantly more money (stimulus B: \$5.11, $SD\pm 2.8$) than a stranger (Novel: \$4.15, $SD\pm 2.6$; paired t-test B-Novel: $t(44)=3.02$, $p=.004$, Fig 3), albeit still less than what is trusted to the kind robber (stimulus A: \$5.94, $SD\pm 2.8$; paired t-test A–B: $t(44)=2.87$, $p=.006$). Interestingly, this unblocking effect was only observed for kind robbers who stole very small amounts, and not for greedy Robbers who stole most of the participant’s money (rmANOVA for greedy robber, greedy robber’s partner, greedy pair and Novel: $F(3,132)=2.70$, $p=0.050$, partial $\eta^2=.07$). The greedy robber’s partner was trusted with the same amount of money (stimulus J: \$4.06, $SD\pm 2.9$) as a stranger (Novel: \$4.15, $SD\pm 2.6$; paired t-test J-Novel: $t(44)=-.38$, $p=0.702$), and more than what was entrusted to the greedy robber (stimulus I: \$3.43 $SD\pm 2.9$; paired t-test I–J is trending: $t(44)=1.88$, $p=0.067$), revealing that—similar to the findings of Experiment 1—the greedy robber blocked his partner from acquiring social value.

Memory Results—When we probed episodic memory for how much money each robber stole, we found that it mirrored the asymmetrical blocking effects observed in the behavioral data. Participants were equally accurate at remembering how much money the kind robbers—and their partners—stole (rmANOVA: $F(2,68)=1.39$, $p=.255$), but failed to accurately remember how much the greedy robber’s partner stole compared to the other robbers who behaved in similarly greedy ways (rmANOVA between greedy robber, greedy robber’s partner, greedy pair: $F(2,76)=3.44$, $p=.037$, $\eta^2=.083$).

To understand why participants were able to associate the kind robber’s partner with social value but failed to associate the greedy robber’s partner (or any of the dictators’ partners in Experiment 1) with value, we examined participants reported subjective feelings after interacting with each dictator and robber in both experiments. We theorized that if there were discrepancies in how participants felt about being robbed versus given money it might elucidate that specific—and possibly divergent—expectations are linked to social contexts involving gains versus losses. Critically, in both experiments, participants retain the same monetary payout (e.g., because participants are first endowed with \$5 in the Robbery Game, there is approximately a \$4 payout after a small amount of money is stolen, mirroring the altruistic split offered during the Dictator Game).

To assess whether social gain and loss are differentially experienced, we subtracted participants’ ratings (5-point analogue scale, where 5=very happy and 1=very unhappy) from 3 (the midpoint, a neutral rating), such that positive feelings were indicated by a positive difference score and negative feelings by a negative difference score. We then compared the degree of difference across both experiments, specifically examining how participants reported feeling about altruistic dictators versus kind robbers, and selfish dictators versus greedy robbers (all matched in their associated value). If the amount of money paid out is the central feature of the task, then receiving large amounts of money versus having fairly little money stolen should not result in divergent subjective ratings, as

participants end up with the same amount of money from both altruistic dictators and kind robbers (the same applies for selfish dictators and greedy robbers).

Results reveal that kind robbers, who only stole very small amounts of money engendered significantly higher positive ratings compared to altruistic dictators who gave participants a lot of money (ANOVA: $F(1,89)=4.20$, $p=0.043$; Fig 4). In contrast, participants felt similarly negative about dictators who selfishly did not share the money and robbers who greedily stole most of the participant's money ($F(1,89)=1.88$, $p=0.174$). Thus, despite the fact that monetary outcomes in Experiment 1 and 2 were matched, the context in which the money was earned and lost (Robbery Game), or simply gained (Dictator Game) generated significantly different subjective experiences. In other words, participants' feelings about the outcomes reflected the asymmetric behavioral and memory blocking effects observed in Experiments 1 and 2, intimating that there may be specific beliefs and expectations for normative social behavior in the loss domain compared to gain domain.

DISCUSSION

Our daily lives are filled with encounters that can adaptively guide choice, and yet little is understood about how humans learn about the value of others in group settings. Here we leverage an associative learning framework to investigate whether mechanisms typically observed within the non-social domain concisely describe the complex ways in which social stimuli acquire value to bias social choice. Results reveal that within the gain domain a Pavlovian blocking mechanism explains social decision-making, however in the loss domain this mechanism fails to fully account for the underlying learning processes.

We found a blocking effect consistent with the idea that there is no prediction error in learning when the same outcome occurs in the presence of both the dictator and the dictator's partner. This blocking phenomenon illustrates that when interacting with multiple dictators at once, people do not associate value with individuals who seem to offer no new information. While participants entrusted much money to altruistic dictators and little to selfish dictators, both altruistic and selfish dictators' partners were entrusted with the same amount of money—which was indiscernible from the amount entrusted to a never-encountered stranger. While this blocking mechanism was systematically observed within the gain domain, in the loss domain, an asymmetric blocking effect was observed: people learned about and entrusted their money to seemingly 'redundant' kind robbers' partners who refrained from stealing (i.e. unblocking), but failed to associate value to greedy robbers' partners (i.e. blocking).

We observed a similar asymmetric blocking effect within episodic memory, such that participants exhibited accurate episodic memory for all dictators and robbers, but were unable to accurately remember the partners' offers. Although decision-making research rarely assesses episodic memory (Murty et al. 2016), this is to our knowledge, the first evidence that blocked episodic memory may bias subsequent social choice. In the one instance where unblocking occurred—for kind robbers' partners—accurate episodic memory for partners and their offers was also observed, supporting the idea that episodic memory could be critical for learning the social value of others.

An important feature of these findings illustrates that social learning appears to recruit basic Pavlovian mechanisms observed across species (Pavlov 1927). The revolutionary discovery of blocking revealed that learning was not the result of mere co-occurrence of conditioned stimuli (in this case, the dictators) and unconditioned stimuli (the social outcome of receiving money) (Mackintosh 1975, Pearce and Hall 1980). Rather, learning relies on the ‘surprise’ of receiving an outcome—e.g. prediction errors—and the quality of the conditioned stimuli in predicting the outcome. In our paradigm, since the dictator already predicts the outcome, there is no prediction error when the partner is present, and thus the partner acquires no social value. While this is consistent from an associative learning framework, from a social perspective, it is surprising that the partner does not hold any informative value and is treated like a stranger—given that participants have explicit information about their behavior.

A conventional Rescorla-Wagner model (Rescorla and Wagner 1972) can easily account for the blocking effects observed in the gain domain, as there is no error to drive learning. However, this account predicts broad blocking effects regardless of context, and it would be unable to explain the asymmetric unblocking effects observed in the loss domain. Why would conventional associative learning models perfectly capture the learning effects within the social gain domain, but fail to capture how humans experience and learn about social loss? One possibility captured by an associability account (Pearce and Hall 1980) posits that the link between a stimulus and its outcome critically depends on the attention paid to the stimuli (Mackintosh 1975, Le Pelley 1993). Since losses garner more attention and are perceived as more emotionally salient (Mellers et al. 1999, Breiter et al. 2001), individuals may be more acutely attuned to the possibility of losing hard-earned money. This account, however, struggles to explain why social value is learned only for robbers who are involved in small losses but not large losses.

An alternative, and perhaps more plausible account is that there are different inferences and expectations of social behavior within the loss domain, which subsequently influences how social phenomena are attended to and experienced. If this is the case, the observed asymmetric learning within the social loss domain may be better captured by contemporary Bayesian learning accounts that describe how prior expectations (Courville et al. 2006) govern how individuals statistically reason about the likelihood of events (Dayan and Long 1998). If social expectations dictate that stealing typically occurs when there is opportunity, then the statistical priors of the environment indicate a high likelihood of stealing. Violations of these statistical priors are anomalies that produce learning and enable unblocking. Applied to our paradigm, the initial act of a kind robber failing to steal money in Run 1 elicits a highly rewarding outcome (evidenced by participants’ subjective reports), and a positive association with the kind robber is produced. When the kind robber is later paired with a partner, the assumption is money will now be stolen, since if people canonically steal, the introduction of a new robber should result in a monetary loss. Because the kind robber and partner fail to steal, the expectation is again violated, allowing the partner to acquire positive value. In this context, the surprising outcome is that stealing does not ensue even with the addition of another robber. That blocking occurs when greedy robbers and partners continuously steal large amounts of money can also be explained with the same logic, since the expectation that people advantageously steal remains intact. A Bayesian account allows

for the existence of asymmetric beliefs or expectations between social loss and gain—which, as recent research within the moral domain has shown—may be how moral phenomena are in fact experienced (Chakroff 2016).

That the same basic mechanisms which govern associations between elementary sensory stimuli and appetitive or aversive outcomes also seem to govern—in part—complex social associations, has important implications for understanding the building blocks of social learning. Just as associative learning processes rely on prediction errors, so to do the processes that underpin how people learn about the social value of others in dynamic groups. Together, these results illustrate that domain-general mechanisms are likely a fundamental feature of human learning, regardless of whether the context is social or not. While this work helps identify and characterize possible core learning mechanisms supporting the representation of social value, we readily recognize that there may be certain aspects of social learning that do not recruit domain-general processes. Future work aimed at disentangling the mechanisms that predominate—or those that fail to work—will help further elucidate the cognitive processes underlying complex social learning.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Funding Source: National Institute of Aging

References

- Beckers T, De Houwer J, Pineno O, Miller RR. Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology-Learning Memory and Cognition*. 2005; 31(2):238–249. DOI: 10.1037/0278-7393.31.2.238
- Bolton GE, Katok E, Zwick R. Dictator game giving: Rules of fairness versus acts of kindness. *International Journal of Game Theory*. 1998; 27(2):269–299. DOI: 10.1007/s001820050072
- Bott L, Hoffman AB, Murphy GL. Blocking in category learning. *Journal of Experimental Psychology-General*. 2007; 136(4):685–699. DOI: 10.1037/0096-3445.136.4.685 [PubMed: 17999579]
- Breiter HC, Aharon I, Kahneman D, Dale A, Shizgal P. Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron*. 2001; 30(2):619–639. [PubMed: 11395019]
- Chakroff AR, PS, Piazzac J, Young L. From impure to harmful: Asymmetric expectations about immoral agents. *J Exp Soc Psychol*. 2016
- Courville AC, Daw ND, Touretzky DS. Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences*. 2006; 10(7):294–300. DOI: 10.1016/j.tics.2006.05.004 [PubMed: 16793323]
- Dayan, P., Long, T. *Statistical models of conditioning*. Vol. 10. MIT Press; 1998.
- Dickinson A, Hall G, Mackintosh NJ. Surprise and Attenuation of Blocking. *Journal of Experimental Psychology-Animal Behavior Processes*. 1976; 2(4):313–322. DOI: 10.1037//0097-7403.2.4.313
- Dickinson A, Shanks D, Evenden J. Judgment of Act-Outcome Contingency - the Role of Selective Attribution. *Quarterly Journal of Experimental Psychology Section a-Human Experimental Psychology*. 1984; 36(1):29–50.
- Dreber A, Ellingsen T, Johannesson M, Rand DG. Do people care about social context? Framing effects in dictator games. *Experimental Economics*. 2013; 16(3):349–371. DOI: 10.1007/s10683-012-9341-9

- Engel C. Dictator games: a meta study. *Experimental Economics*. 2011; 14(4):583–610. DOI: 10.1007/s10683-011-9283-7
- Fox C, Tom S, Trepel C, Poldrack R. The Neural Basis of Loss Aversion in Decision-Making Under Risk. *Advances in Consumer Research*. 2008; 35:129–130.
- Glimcher, PW. Neuroeconomics and the Study of Valuation. In: Gazzaniga, MS., editor. *The Cognitive Neurosciences*. 4. Cambridge, MA: The MIT Press; 2009.
- Gluck MA, Bower GH. From Conditioning to Category Learning - an Adaptive Network Model. *Journal of Experimental Psychology-General*. 1988; 117(3):227–247. DOI: 10.1037/0096-3445.117.3.227 [PubMed: 2971760]
- Greenberg J. Stealing in the Name of Justice - Informational and Interpersonal Moderators of Theft Reactions to Underpayment Inequity. *Organizational Behavior and Human Decision Processes*. 1993; 54(1):81–103. DOI: 10.1006/obhd.1993.1004
- Greene JD, Paxton JM. Patterns of neural activity associated with honest and dishonest moral decisions. *Proc Natl Acad Sci U S A*. 2009; 106(30):12506–12511. DOI: 10.1073/pnas.0900152106 [PubMed: 19622733]
- Kahneman D, Miller DT. Norm Theory - Comparing Reality to Its Alternatives. *Psychological Review*. 1986; 93(2):136–153. DOI: 10.1037/0033-295x.93.2.136
- Kahneman D, Tversky A. Prospect Theory - Analysis of Decision under Risk. *Econometrica*. 1979; 47(2):263–291. DOI: 10.2307/1914185
- Kahneman D, Tversky A. Choices, values, and frames. *American Psychologist*. 1984; 39(4):341.
- Kamin LJ. Predictability, surprise, attention, and conditioning. *Punishment and aversive behavior*. 1969:279–296.
- Kashima Y, Woolcock J, Kashima ES. Group impressions as dynamic configurations: The tensor product model of group impression formation and change. *Psychological Review*. 2000; 107(4): 914–942. DOI: 10.1037//0033-295X.107.4.914 [PubMed: 11089411]
- King-Casas B, Tomlin D, Anen C, Camerer CF, Quartz SR, Montague PR. Getting to know you: reputation and trust in a two-person economic exchange. *Science*. 2005; 308(5718):78–83. DOI: 10.1126/science.1108062 [PubMed: 15802598]
- Klucharev V, Hytonen K, Rijpkema M, Smidts A, Fernandez G. Reinforcement Learning Signal Predicts Social Conformity. *Neuron*. 2009; 61(1):140–151. DOI: 10.1016/j.neuron.2008.11.027 [PubMed: 19146819]
- Le Pelley ME. The Role of Data Administration in Information Engineering. *Journal of Computer Information Systems*. 1993; 34(2):87–91.
- Le Pelley ME. The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology*. 2004; 57(3):193–243. DOI: 10.1080/02724990344000141
- Linville PW, Salovey P, Fischer GW. Perceived Distributions of the Characteristics of in-Group and out-Group Members - Empirical-Evidence and a Computer-Simulation. *Journal of Personality and Social Psychology*. 1989; 57(2):165–188. DOI: 10.1037/0022-3514.57.2.165 [PubMed: 2760805]
- Lovibond PF, Been SL, Mitchell CJ, Bouton ME, Frohardt R. Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory & Cognition*. 2003; 31(1): 133–142. [PubMed: 12699149]
- Mackintosh NJ. Theory of attention - variations in associability of stimuli with reinforcement. *Psychological Review*. 1975; 82(4):276–298. DOI: 10.1037/h0076778
- Mackintosh NJ. A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*. 1975; 82
- Mazar N, Amir O, Ariely D. The Dishonesty of Honest People: A Theory of Self-Concept Maintenance. *Journal of Marketing Research*. 2008; 45(6):633–644. DOI: 10.1509/jmkr.45.6.633
- Mckelvey RD, Palfrey TR. An Experimental-Study of the Centipede Game. *Econometrica*. 1992; 60(4):803–836. DOI: 10.2307/2951567
- Mellers B, Schwartz A, Ritov I. Emotion-based choice. *Journal of Experimental Psychology-General*. 1999; 128(3):332–345. DOI: 10.1037//0096-3445.128.3.332

- Murty VP, FeldmanHall O, Hunter LE, Phelps EA, Davachi L. Episodic memories predict adaptive value-based decision-making. *Journal of Experimental Psychology-General*. 2016; 145(5):548–558. DOI: 10.1037/xge0000158 [PubMed: 26999046]
- O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C. Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*. 2001; 4(1):95–102. [PubMed: 11135651]
- O'doherty JP. Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Curr Opin Neurobiol*. 2004; 14(6):769–776. DOI: 10.1016/j.conb.2004.10.016 [PubMed: 15582382]
- Pavlov, IP. *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex*. London: Oxford University Press: Humphrey Milford; 1927.
- Pearce JM, Hall G. A Model for Pavlovian Learning - Variations in the Effectiveness of Conditioned but Not of Unconditioned Stimuli. *Psychological Review*. 1980a; 87(6):532–552. DOI: 10.1037//0033-295x.87.6.532 [PubMed: 7443916]
- Pearce JM, Hall G. A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*. 1980b; 87(6):532–552. [PubMed: 7443916]
- Phan KL, Sripada CS, Angstadt M, McCabe K. Reputation for reciprocity engages the brain reward center. *Proc Natl Acad Sci U S A*. 2010; 107(29):13099–13104. DOI: 10.1073/pnas.1008137107 [PubMed: 20615982]
- Plato. *The Republic*. New York: Dutton; 1950.
- Rangel A, Camerer C, Montague PR. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*. 2008; 9(7):545–556. DOI: 10.1038/Nrn2357 [PubMed: 18545266]
- Rescorla, RA., Wagner, AR. *Classical conditioning II: current research and theory*. New York, NY: 1972. *A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement*.
- Rilling JK, King-Casas B, Sanfey AG. The neurobiology of social decision-making. *Curr Opin Neurobiol*. 2008; 18(2):159–165. DOI: 10.1016/j.conb.2008.06.003 [PubMed: 18639633]
- Ruff CC, Fehr E. The neurobiology of rewards and values in social decision making. *Nature Reviews Neuroscience*. 2014; 15(8):549–562. DOI: 10.1038/nrn3776 [PubMed: 24986556]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275(5306):1593–1599. DOI: 10.1126/science.275.5306.1593 [PubMed: 9054347]
- Seid-Fatemi A, Tobler PN. Efficient learning mechanisms hold in the social domain and are implemented in the medial prefrontal cortex. *Soc Cogn Affect Neurosci*. 2015; 10(5):735–743. DOI: 10.1093/scan/nsu130 [PubMed: 25326037]
- Smith ER, Zarate MA. Exemplar-Based Model of Social Judgment. *Psychological Review*. 1992; 99(1):3–21. DOI: 10.1037/0033-295x.99.1.3
- Sutton, RS., Barto, AG. *Reinforcement Learning: An Introduction*. Cambridge University Press; 1998.
- Tversky A, Kahneman D. The framing of decisions and the psychology of choice. *Science*. 1981; 211(4481):453–458. [PubMed: 7455683]
- Vurbic D, Bouton ME. *A Contemporary Behavioral Perspective on Extinction*. Wiley Blackwell Handbook of Operant and Classical Conditioning. 2014; :53–76. DOI: 10.1002/9781118468135

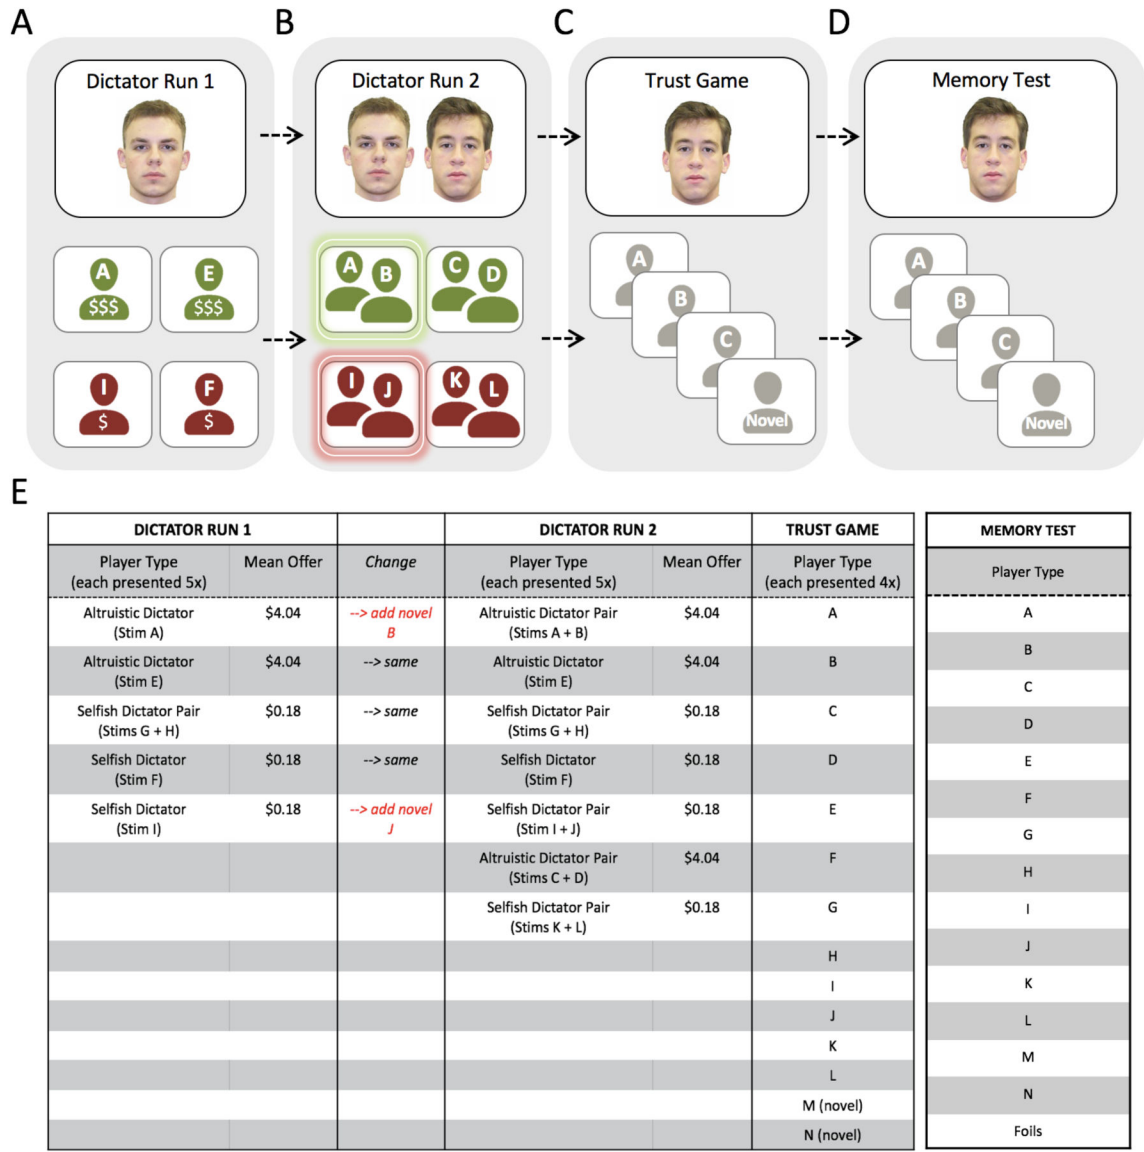


Fig 1. Task Structure of Experiment 1

All participants play two Dictator games before re-encountering the same Dictators in a Trust game. **A)** In Dictator Run 1, participants received various altruistic splits from good dictators (stimuli A and E) and selfish splits from bad dictators (stimuli I and F). After each split, participants rate how they feel. **B)** In Dictator Run 2, participants again received altruistic or selfish splits from dictators, only this time some dictators made decisions as a pair. For example, Dictator I who made selfish splits by himself in Run 1, makes selfish splits collectively with Dictator J in Run 2. **C)** Participants then play a Trust game, and can decide how much of \$10 to entrust in another individual. These individuals are the dictators from Runs 1–2, as well as novel, never before seen people who have no history with the participant (Novel). **D)** Finally, participants complete a surprise memory test in which they are asked to report how much money each dictator offered them in the Dictator games.

Experiment 2 followed a similar task structure. E) The temporal presentation of all stimuli across the entire task.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

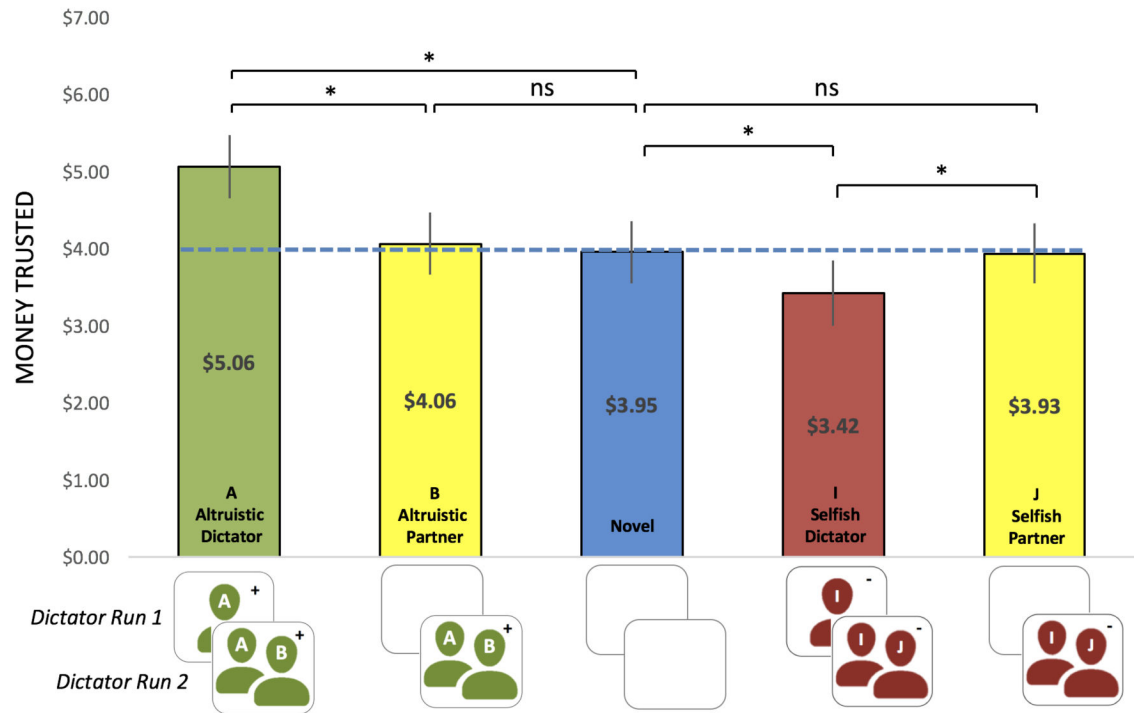


Fig 2. Money Entrusted to Dictators in Experiment 1

During the Trust Game, participants entrusted the most amount of money to the altruistic dictator (A) and the least amount to the selfish dictator (I). The dictators' partners (indicated in yellow: B and J) were trusted with the same amount of money as a novel stranger (indicated by dashed blue line, Novel), effectively revealing that prior conditioning with altruistic and selfish dictators results in a failure to associate their partners with either positive or negative social value. Temporal presentation of dictators and their partners across Runs 1–2 are denoted below each bar. Error bars represent 1 SEM.




Fig 3. Money Entrusted to Robbers in Experiment 2

During the Trust Game, participants entrusted the most amount of money to the kind robber (A) and the least amount to the greedy robber (I). While the kind robber's partner (B) was trusted with more money than a stranger (Novel), the greedy robber's partner (J) was trusted with the same amount as a stranger. Unlike the results from Experiment with dictators, in the loss domain, participants exhibited asymmetric blocking effects, where kind robber's partners acquired positive social value, but greedy robber's partners failed to acquire any social value. Error bars represent 1 SEM.

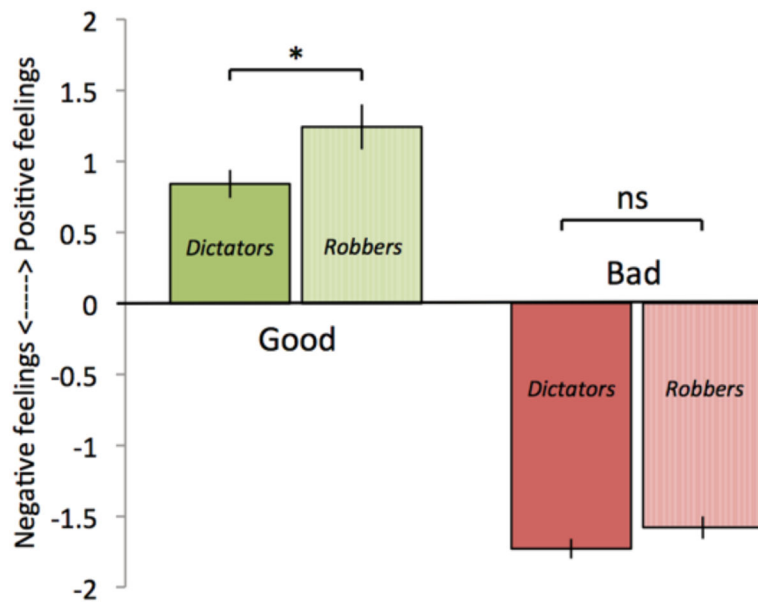


Fig 4. Subjective reported feelings for dictators and robbers across Experiments 1–2

Despite the fact that both Experiments incurred the same monetary payouts between good dictators and robbers, participants reported feeling significantly more positive about outcomes from kind robbers who stole only small amounts of money compared to dictators who made large and altruistic offers. In contrast, participants felt similarly about the same monetary outcomes from selfish dictators and greedy robbers.