Social Cognition, Vol. 38, No. 3, 2020, pp. 266–286

ATTRIBUTIONS OF SOCIAL INTERACTION DEPEND ON THE INTEGRATION OF THE ACTOR'S SIMPLE GOAL AND THE INFLUENCE ON RECIPIENTS

Xiaoyan He, Yisong Yang, and Jing Lin
Ningbo University

Xiaoying Wu
Zhejiang University

Jun Yin
Ningbo University

Understanding social interactions is an essential aspect of a person's social life. Currently, it is unclear how the object-directedness of an actor's actions affects social interaction attributions and whether/how such a relationship is modulated by his/her action's outcomes. We have found that, in both helping and hindering contexts, when the actor's action had a considerable influence on the recipient (i.e., either decreasing or increasing the recipient's cost to reach his/her goal), the social interaction between agents was attributed as more intense when the actor had a clearer (i.e., strong) object-directed goal than when he/she had an ambiguous (i.e., weak) one. This effect became smaller (or absent) when there was no influence on the recipient. These findings indicate that the object-directedness of an actor's action plays a key role in social interaction attributions, and that this role is based on whether or not the actor is influencing the recipient.

Keywords: social interaction, social goal, object-directed goal, influence, cost

---

Supplemental materials are available from the author.

X. He and J. Yin conceived and designed the experiments. X. He, Y. Yang, J. Lin and X. Wu performed the experiments and analyzed the data. X. He and J. Yin wrote the manuscript.

This work was supported by the MOE (Ministry of Education in China) Project of Humanities and Social Sciences under Grant No. 16YJC190026); and the National Natural Science Foundation of China (Grant no. 31871091); and K. C. Wong Magna Fund of Ningbo University.

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Address correspondence to Jun Yin, Department of Psychology, Ningbo University, No. 616 Fenghua Road, Ningbo, 315211, China. Email: yinjun1@nbu.edu.cn.

Humans can infer the goals of others, even those of agents displayed with moving geometrical shapes (Gao, Newman, & Scholl, 2009; Hamlin, Wynn, & Bloom, 2007; Schachner & Carey, 2013). Upon seeing action patterns within a given context, we can easily and spontaneously infer the goals behind them. Several studies have documented that even infants easily assign a goal to an object-directed action (Feiman, Carey, & Cushman, 2015; Woodward, 1998), such as reaching for an apple. This inference has been suggested to follow the principle of rationality; that is, in a given situation, the most efficient action is selected to achieve its goal (Gergely & Csibra, 2003). Jara-Ettinger and colleagues (2016) advanced this concept—positing that agents choose goals and actions to maximize the rewards they expect to obtain relative to the costs they incur, which is referred to as the "naïve utility calculus." Often, rather than acting alone to influence objects in the environment, individuals have social goals to interact with others; in so doing, they influence others' overt actions or state of mind, which is considered one of the building blocks of society (Hinde, 1976; Knoblich & Sebanz, 2008). In that regard, this study focused on whether and how social interactions are attributed.

No one doubts that adults can recognize social interactions; however, this ability is already blossoming in the early stages of development (Choi & Luo, 2015; Hamlin et al., 2007; Hamlin, Wynn, Bloom, & Mahajan, 2011; Stahl & Feigenson, 2014; Tatone, 2017; Tatone, Geraci, & Csibra, 2015). For instance, Hamlin and colleagues (2007) reported that when evaluating individuals, 6- and 10-month-old infants consider how they interact with others, showing preference for helpful individuals over those who hinder others based on their historical interactions. Further, Choi and Luo (2015) investigated how 13-month-olds use their emergent theory-of-mind understanding to make sense of social interactions; specifically, infants watched three puppets (A, B, and C) interact. They found that after seeing two agents (A and B) interact in a positive manner, infants expected them to continue doing so, even after they saw one hit the third agent (C). Hence, infants understood the social interaction between two agents and expected such an interaction to persist. Moreover, Tatone and colleagues (2015) revealed that 1-year-old children possess an action schema of giving, which is specific to representing social interactions based on material exchange. They designed looking-time studies using a violation-of-expectation paradigm and found that infants could distinguish between kinematically identical giving and taking actions. Despite the surface similarity between these two actions, only giving was represented as an object-mediated social interaction.

Similar with the attributions of simple object-directed goals, researchers proposed that attributions of social interaction goals also accord with the naïve utility calculus (Jara-Ettinger et al., 2016). Specifically, social interactions are treated as second-order interpretations of agents' actions, which differ from the object-directed goal as a first-order goal. In other words, observers utilize the naïve utility calculus to compute everyone's costs and rewards separately; subsequently, they integrate them to infer social interaction goals. In such a computation, the costs of an agent's actions are difficult to interpret from the reward of his/her first-order object-directed goal; however, they are compensated by the second-order

expectation of bringing about the rewarded outcomes incurred by another agent. This idea is widely documented (Hamlin, Ullman, Tenenbaum, Goodman, & Baker, 2013; Tatone et al., 2015; Ullman et al., 2009). For example, Ullman and colleagues (2009) constructed Bayesian models based on inverse planning in multiagent Markov decision problems and found that when the model integrated both individuals' cost and rewards, the performance of models were almost the same as humans.

However, unlike solitarily acting on a physical object, which can be easily attributed as a simple "object-directed" goal (through actions like reaching for the object), action patterns whereby two individuals interact are complex (Malle, 2004). For instance, an individual might have a self-interested object-directed goal, a more complex social interaction goal, or both simultaneously (Hamlin et al., 2013; Ullman et al., 2009). In such cases, the interpretation of actions can be ambiguous. Specifically, observers might attribute the goal of social interaction because of the action–effect association between the two individuals; however, they might not interpret the action as socially oriented (i.e., possessing a social goal) and may instead interpret it as object-directed, with an incidental effect on the other person. For instance, in a situation where an individual opens a door for a stranger, an observer might interpret the opening of the door as helping the stranger or as merely necessary for the individual to enter the building, with the stranger being present in time to benefit from the action. Hence, it seems that attributions about an object-directed goal may influence social interaction attributions. However, whether this effect does occur and how it is exhibited are unclear. Understanding the mechanism would clarify how humans infer social goals from multiple agents' actions.

Intuitively, it is obvious that when the agent's action is compatible with two goals—a social interaction one and an object-directed one—a lesser likelihood of attributing one goal would suggest a greater likelihood of attributing the alternative goal, mainly because the goals behind actions should be consistent, and cannot be two contradicting states (Dennett, 1989; Malle, 2004). For example, in a given moment, a person cannot have both goals: one of interacting with another person and one of directing his/her action to a physical object without any social interaction. As such, when an actor seems less likely to be directing his/her action to an object (i.e., violating the principle of rationality), observers are more likely to attribute more intense social interaction to that action pattern. Similarly, the correspondent inference theory (Jones & Davis, 1965), and its further extension (Jones & McGillis, 1976), has a clear prediction about the perceiver's inferences on what an actor may be trying to achieve by a particular action: it is when people do not do what is socially expected that their behavior is informative (i.e., behavioral expectancy). Namely, when the agent does not perform a behavior that is rationally expected (with respect to the object-related goal), the observer can infer that he/she has a social interaction goal.

Humans' actions, suggestively, are always aimed at obtaining external outcomes, so proper understanding of the reasoning behind these actions relies on one's knowledge of the association between the action and its incurred outcome, and how that outcome justifies the goal behind the action (Dennett, 1989; Marien, Aarts, &

Custers, 2015). For instance, a woman's aim of bringing an apple closer to her (the outcome) by grasping it justifies her goal of fetching the apple. Further, the correspondent inference theory also assumes that the outcomes of a chosen action must be compared with the outcomes of possible alternative actions for it to be appropriately understood (i.e., non-common effects); based on this, the more specific the action, the greater one's ability to infer its underlying intention (Jones & Davis, 1965; Jones & McGillis, 1976). Hence, whether or not the outcomes of actions are specific to a social interaction goal may be another key factor influencing social interaction attributions; nevertheless, this topic and correlating ones (whether people integrate the outcome of an actor's action with its object-directedness and how it influences people's attributions) have yet to be widely explored, and answers in this regard may help us better understand how the sources of behavioral expectancy and non-common effects jointly contribute to one's attribution intention.

To deepen the discussion on one of these topics, one question posits itself: What are the non-common outcomes of social interaction goals? Gillin and Gillin (1942) theorized that a social goal is instantiated in a social interaction, which results in behavioral modification for others. Hinde (1976) defined social interactions as the following basic process: one participant will do X to another, who will then respond with Y. Empirical studies have shown that infants treat an actor causing a change in a recipient as having social goals (Choi & Luo, 2015; Tatone et al., 2015; Van de Vondervoort & Hamlin, 2018). Similarly, adults interpret the actions with a stronger effect on the recipient (i.e., more helpful/harmful) as more interactive than actions with weaker effects (i.e., less helpful/harmful; Wu, Hua, Yang, & Yin, 2018). In summary, the specific outcome of a social goal is to influence another person/agent, and whether an actor influences a recipient would modulate the role of object-directedness of an actor's action on attributing social interactions.

As to the conceptual organization about the intention and outcome, the effect of the outcome can be treated as the evidence justifying the underlying intention of the action and the rejection of the alternatives (Dennett, 1989; Malle, 2014). Hence, if the rationality of the actor's object-directedness is violated when the outcome—influencing a recipient—is observed, the actor's underlying social interaction goal could be supported; accordingly, the attributed social interaction would be more likely or intense as compared to when the rationality of the actor's object-directedness is compatible with the observed effect of the outcome. In contrast, when the influence of actions is absent, there is no evidence to justify the desired outcome of a social interaction; thus, social interaction attributions should not be affected by object-directedness information, or be less affected, relative to when the influence on the recipient is present.

To test the above hypothesis, we used visual animations with simple displays of moving shapes as stimuli. This method has been widely used to explore how to infer goals behind actions, which can be easily and strictly controlled for possible confounding variables (Gao et al., 2009; Hamlin et al., 2007; Schachner & Carey, 2013; Ullman et al., 2009), such as movement patterns. In these stimuli, we operationalized the "object-directedness" of an actor's goal as the likelihood that the actor's action was directed toward an object (i.e., an apple) and manipulated the

desired influence of the action on the recipient by determining whether the action changed the recipient's cost (i.e., traveled path) to the recipient's simple goal.

EXPERIMENT 1: HELPING ACTIONS

In this experiment, we investigated how the object-directedness of an actor's goal and the influence of his/her action on the recipient influences social goal attributions for helping actions. Previous studies have suggested that humans universally obey the principle of rationality for explaining actions; that is, they assume that a to-be-explained action is rational, and that the reason that agents adopt the action is because it is most cost-effective and efficient for achieving their goal/intention and has beneficial outcomes for them (Gergely & Csibra, 2003; Jara-Ettinger et al., 2016). This assumed model of action understanding may be a cognitive constraint to predicting actions, mainly because the observed action aimed at achieving a specific target can be either rational or irrational: for the rational observed action, the likelihood of a person perceiving its directedness to a specific object is very high because it fits well with the characteristics of a rational action; for the irrational observed action, the same likelihood is very low because it does not fit well with these same characteristics (Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017; Schachner & Carey, 2013). For example, in the work of Schachner and Carey (2013), there was a low proportion of participants reporting that the irrational action was directed to an external object, but this proportion was absolutely lower than that of the rational action. Usually, when an agent approaches an object in an irrational way (e.g., travelling a high-cost [e.g., longer] path instead of selecting a low-cost [e.g., shorter] path), observers will interpret his/her directedness to the object as weak; in contrast, when the action is approached rationally, observers will interpret the directedness to the object as strong. Hence, we manipulated the object-directedness of an actor's choice of whether to take a rational (i.e., low-cost) path to approach an apple. Although the helped agent (i.e., recipient) ultimately obtained the interested object, we manipulated the influence of the helping action concerning the recipient's path to the object: the recipient was influenced when they could use a lower cost path than before to obtain an apple with the actor's help, and they were not influenced when the recipient obtained the apple using the same costs as before.

METHODS

Participants. Forty-four paid adult volunteers (19 men and 25 women; mean age = 22 years, age range = 18–27 years) participated in this experiment. The sample size was determined via a power analysis based on a predicted effect using G*power 3, which is described herein: Based on prior results (Wu et al., 2018), we used a more conservative estimate, between a small and medium effect size (medium size: $f = 0.25$, and small size: $f = 0.10$, according to the sizes stipulated by Cohen [1988]); with the alpha set at 0.05 and the power set at 0.80, the suggested sample size was about 44 individuals, which was almost the same as a recent study

that explored social interaction processing (Wu et al., 2018). All participants had normal or corrected-to-normal vision and no history of neurological disorders.

This study was approved by the Research Ethics Board of the Department of Psychology at Ningbo University and was performed in accordance with the relevant guidelines and regulations. Participants were given information sheets before the experiment; after they had learned the purpose and procedure of the experiment, they provided their written informed consent.

Apparatus and Stimuli. All stimuli were presented against a black background on a 19-inch cathode ray tube monitor (resolution = 1024 × 768 pixels; refresh rate = 100 Hz) at a 60-cm viewing distance. The visual stimuli were created using the free, open-source 3D creation suite Blender (https://www.blender.org/) and comprised 10-s computer-animated events displayed within the full area of the screen (22.2°× 16.6°; please see Figure 1 for a screenshot of animations and https://osf.io/xtf2s/?view_only=879f1986a6414afa88ce82bdd626cd2d for video details). The events were presented in full at a 24.2° viewing distance and a 3D perspective view of 57°.

Each event involved two cartoon agents of different colors: blue and green. The blue agent (i.e., the actor) was initially positioned in the upper left corner of the screen. It then approaches an apple positioned at the lower left corner of the screen to pick it up and place it in front of a stone near the actor's starting position. While approaching the apple, the blue agent must pass through a gap in a wall. The wall contains two gaps: one is unobstructed, whereas the other is blocked by a movable obstacle. The path to the apple is shorter through the unobstructed gap than through the blocked gap. However, the width of the unobstructed gap differs according to the condition: in the weak object-directedness condition, the gap is

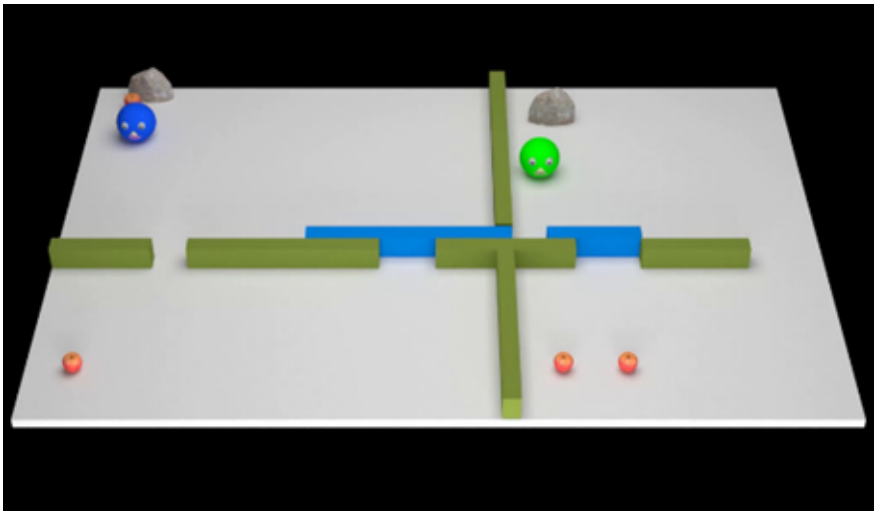

FIGURE 1. Screenshot of animations.

exactly the width of the blue agent, therefore enabling the agent to pass through it easily; however, in the strong object-directedness condition, the gap is narrower than the blue agent, thus forcing the agent to move the obstacle to reach the apple. In both conditions, the blue agent moves the obstacle to take the apple; however, the latter condition is more rational and thereby more object-directed, while the former condition is more irrational and therefore less object-directed in obtaining the apple.

On the other hand, the green agent (i.e., the recipient) remains at the upper right corner of the screen and moves to take two apples positioned at the lower right corner. On his/her way to take the apples, there is a short path that requires him/her to pass through a gap in the wall; however, this is blocked by an obstacle. This forces the green agent to travel a longer distance to bypass the obstacle and obtain the apples. Upon obtaining the first apple, the green agent returns to around its starting position to place the apple in the front of a stone. This action illustrates to observers its intention to take the apple. Afterwards, the green agent moves to fetch the remaining apple. During this period, the obstacle moved by the blue agent can produce two distinct effects on the green agent: (1) forcing the obstacle hindering the green agent away from the gap, allowing the green agent to pick up the apple using the short path; or (2) forcing away the obstacle hindering the green agent and replacing it, thereby forcing the green agent to travel the longer distance. In the first condition, the blue agent's helping action helps the green agent fetch the apple at a lower cost, and thus is considered to have a stronger influence on the green agent than in the second condition. Taken together, there were a total of four conditions arranged in a 2 (actor's goal: strong object-directedness vs. weak object-directedness) by 2 (recipient's change: unchanged cost vs. decreased cost) within-participant design.

Procedure and Design. Before the study, we designed a randomized order for the presentation of the four events. Half of the participants were tested with this predetermined order, called the A version (first event: strong object-directedness with unchanged cost on the recipient; second event: strong object-directedness with decreased cost on the recipient; third event: weak object-directedness with decreased cost on the recipient; fourth event: weak object-directedness with unchanged cost on the recipient), and the other half were tested with a reversed order of the predetermined (and abovementioned) order, called the B version. Each event was replayed until participants pressed the spacebar to proceed to the next event.

For each event, first, participants were asked to answer, "To what extent is the blue agent directed toward obtaining the apple on the lower left corner of the screen?" (rated from 1 = *completely unlikely* to 7 = *completely likely*). This question was designed as a manipulation check for participants' inferences of the blue agent's (i.e., actor's) action and object-directed goal, as the actor's goal should be attributed to the observed physical movements. Regarding the recipient (i.e., green agent), when the changes were directly manifested in its physical movements, there was no manipulation check.

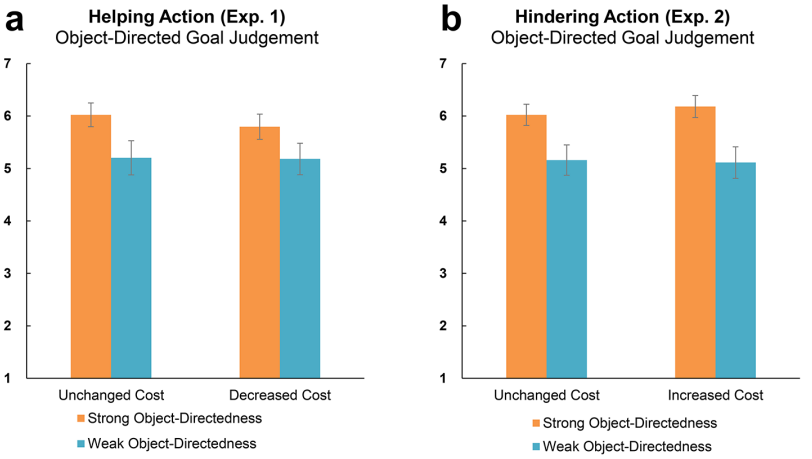

FIGURE 2. Object-directed goal judgement of the actor (i.e., blue agent) to the apple under different acting conditions in Experiment 1 (a) and Experiment 2 (b); 1 means "completely unlikely" and 7 means "completely likely." Error bars indicate the standard error (± SE).

Results for participants' inferences regarding the object-directed goal of the actor's behaviors are presented in Figure 2a and Table 1. A two-way ANOVA with actor's goal and recipient's change as factors revealed a significant main effect of the actor's goal ($F(1,43) = 9.49$, $p = .004$, $\eta^2_p = .18$); however, neither the main effect of the recipient's change ($F(1,43) = 0.58$, $p = .451$, $\eta^2_p = 01$) nor the interaction effect were significant ($F(1,41) = 0.18$, $p = .674$, $\eta^2_p < 01$). Further, the presentation order did not moderate the results (see the supplementary material[1] for detailed results); these results confirmed that participants indeed inferred that an actor's action for obtaining the apple was more object-directed in the strong object-directedness condition than in the weak object-directedness condition.

Besides the actor's goal manipulation check question, participants were required to answer two other questions: "What is the nature of the social interaction between the blue and green agents?" (rated from –5 = *strongly negative* [i.e., strongly hostile] to 5 = *strongly positive* [i.e., strongly friendly]; 0 = *no social interaction between them*); and "To what extent is the blue agent intentionally helping the green agent?" (rated from 1 = *completely unlikely* to 7 = *completely likely*). These two questions were aimed at determining participants' inferences about the social interaction between the two agents from a general and specific understanding of social interaction (Hinde, 1976), respectively. Given that the item measuring the attribution of general social interaction was identical across experiments, this item was asked first. Their answers were used as the dependent variables of interest.

1. Supplementary material can be found online at https://osf.io/xtf2s/?view_only=879f1986a6414afa88ce82bdd626cd2d

TABLE 1. Descriptive Results of Different Rating Dimensions for Helping Action in Experiment 1 (*M*±*SE*)

	Unchanged cost		Decreased cost	
	Strong object-directedness	Weak object-directedness	Strong object-directedness	Weak object-directedness
Actor's object-directed goal	6.02(0.23)	5.20(0.33)	5.80(0.34)	5.18(0.30)
General social interaction attribution	-0.36(0.31)	-0.39(0.35)	1.16(0.30)	2.25(0.30)
Specific social interaction attribution	2.09(0.22)	2.14(0.23)	2.75(0.24)	4.64(0.35)

RESULTS AND DISCUSSION

To examine whether the presentation order influenced participants' evaluations, a three-way analysis of variance (ANOVA) with actor's goal, recipient's change, and presentation order as factors was employed on each variable of interest. We found that the presentation order did not produce any significant interaction effect with other factors (see the supplementary material for detailed results). Hence, in the following analyses, the data with different presentation orders were collapsed.

Figure 3a and Table 1 shows the overall results of the social interaction based on a general understanding of social interaction. The ANOVA showed that both main effects—actor's goal ($F(1,43) = 6.37$, $p = .015$, $\eta^2_p = .13$) and recipient's change ($F(1,43) = 35.43$, $p < .001$, $\eta^2_p = .45$)—were significant, as was the interaction effect ($F(1,43) = 6.27$, $p = .016$, $\eta^2_p = .13$). Post-hoc simple effects analyses revealed that attributed social interactions were more positive for the weak object-directedness condition than for the strong object-directedness condition when the actor's action reduced costs for the recipient ($t(43) = 4.02$, $p < .001$, Cohen's d (hereafter, d) = .61, 95% CI for mean difference = [0.54 1.64]); however, this difference between the weak and strong object-directedness conditions disappeared when the actor's action did not reduce costs for the recipient ($t(43) = 0.07$, $p = .947$, $d = .01$, 95% CI for mean difference = [-0.71 0.66]).

Concerning social goal attributions based on a specific understanding of social interaction (Figure 3b and Table 1), the ANOVA showed that both main effects—actor's goal ($F(1,43) = 18.94$, $p < .001$, $\eta^2_p = .31$) and recipient's change ($F(1,43) = 32.90$, $p < .001$, $\eta^2_p = .43$)—were significant, as was the interaction effect ($F(1,43) = 21.38$, $p < .001$, $\eta^2_p = .33$). Further analyses of the simple effects revealed that social interaction attributions were more positive for the weak object-directedness condition than for the strong object-directedness condition when the actor's action reduced costs for the recipient ($t(43) = 5.03$, $p < .001$, $d = .76$, 95% CI for mean difference = [1.13 2.64]); however, this difference disappeared when the actor's action did not change these costs ($t(43) = 0.24$, $p = .814$, $d = .04$, 95% CI for mean difference = [-0.34 0.43]). Both measures consistently indicated that attributions of social interaction are influenced by the actor's object-directed goal, and that this relationship is modulated by the outcome of influencing the recipient.

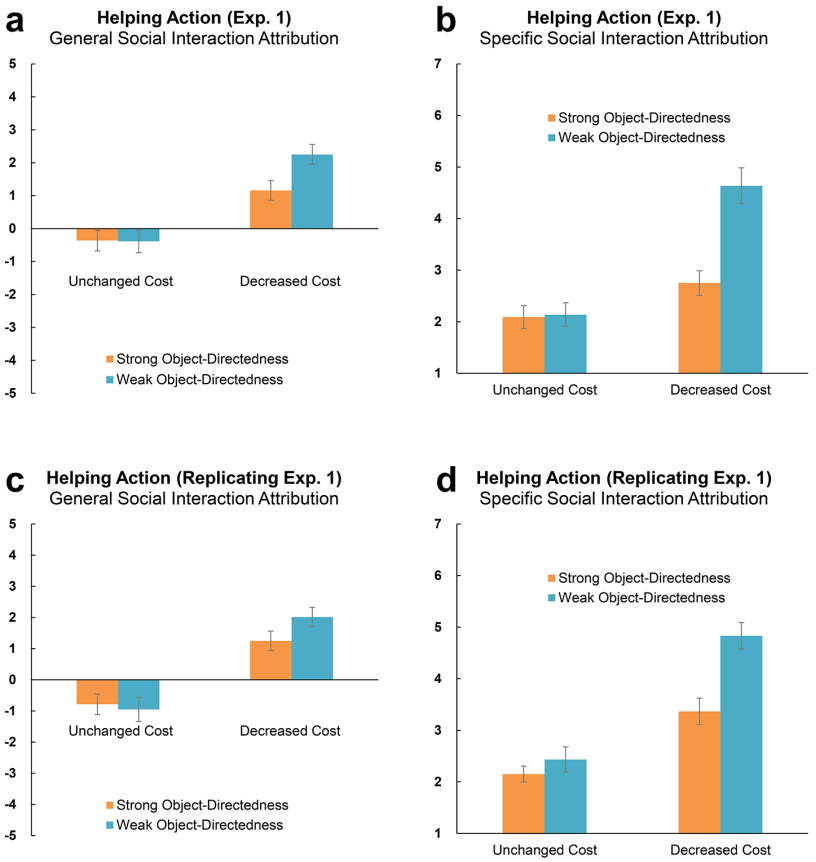

FIGURE 3. Attributions of social interaction goal between the actor (i.e., blue agent) and the recipient (i.e., green agent) under different acting conditions in Experiment 1 and its replication. (a) and (c) represent the averaged rating scores, based on a general understanding of social interaction. A score of -5 means a "strongly negative interaction," 5 means "strongly positive interaction," and 0 means "no social interaction." (b) and (d) represent the averaged rating scores, based on a specific understanding of social interaction. A score of 1 means "completely unlikely," and 7 means "completely likely." Error bars indicate the standard error (± SE).

REPLICATION

In the above experiment, the first question was designed to check the validity of manipulating the actor's goal, and the third was designed to measure the dependent variable of social interaction attribution. This design would allow for participants not to differentiate the roles of these two measures, thus—and importantly—catering to the experimenter's expectation that participants' attribution of the actor's object-directed goal and social interaction would be different.

Hence, we ran another experiment almost identical to Experiment 1, but removed the first question for manipulation check. Further, given that the use of only two different orders may limit the generalization of the conclusions, this experiment used the Latin square design for counterbalancing with the predetermined order

(i.e., first event: weak object-directedness with unchanged cost; second event: weak object-directedness with decreased cost; third event: strong object-directedness with unchanged cost; fourth event: strong object-directedness with decreased cost). A different sample set composed of 60 paid adult volunteers (35 females; mean age = 21 years, ranging from 18 to 25 years) participated in this experiment.

Concerning social goal attributions based on a general understanding of social interaction (Figure 3c and Table 2), the same ANOVA applied in Experiment 1 showed that both the main effect of recipient's change ($F(1,59) = 61.63$, $p < .001$, $\eta^2_p = .51$) and the interaction between actor's goal and recipient's change ($F(1,59) = 5.11$, $p = .028$, $\eta^2_p = .08$) were significant, but the main effect of actor's goal was not significant ($F(1,59) = 1.38$, $p = .245$, $\eta^2_p = .02$). Post-hoc simple effects analyses revealed that social interaction attributions were more positive for the weak than for the strong object-directedness conditions when the actor's action reduced costs for the recipient to obtain the apple ($t(59) = 2.65$, $p = .010$, $d = .34$, 95% CI for mean difference = [0.19 1.35]); however, this difference between the weak and strong object-directedness conditions disappeared when the actor's action did not reduce costs for the recipient ($t(59) = 0.46$, $p = .648$, $d = .06$, 95% CI for mean difference = [-0.89 0.56]).

Concerning social goal attributions based on a specific understanding of social interaction (Figure 3d and Table 2), the ANOVA showed that both main effects—actor's goal ($F(1,59) = 18.52$, $p < .001$, $\eta^2_p = .24$) and recipient's change ($F(1,59) = 57.29$, $p < .001$, $\eta^2_p = .49$)—were significant, as was their interaction effect ($F(1,59) = 11.66$, $p = .001$, $\eta^2_p = .17$). Post-hoc simple effects analyses revealed that social interactions attributions were more positive for the weak than for the strong object-directedness conditions when the actor's action reduced costs for the recipient (t(59) = 4.73, $p < .001$, $d = .61$, 95% CI for mean difference = [0.85 2.09]); however, the difference disappeared when the actor's action did not change these costs ($t(59) = 1.32$, $p = .193$, $d = .17$, 95% CI for mean difference = [-0.15 0.71]). None of the above effects were modulated by the presentation order (ps > 0.210).

Hence, both measures in the two experiments consistently indicated that attributions of social interaction are influenced by the actor's object-directed goal and that this relationship is moderated by the outcome of influencing the recipient.

TABLE 2. Descriptive Results of Different Rating Dimensions for Helping Action in Replication Experiment 1 (*M*±*SE*)

	Unchanged cost		Decreased cost	
	Strong object-directedness	Weak object-directedness	Strong object-directedness	Weak object-directedness
General social interaction attribution	-0.78(0.33)	-0.95(0.39)	1.25(0.31)	2.02(0.31)
Specific social interaction attribution	2.15(0.16)	2.43(0.24)	3.37(0.26)	4.83(0.26)

EXPERIMENT 2: HINDERING ACTION

We then tested the generalizability of the results of Experiment 1 and the robustness of the effects of the actor's goals and the influence of their actions on the recipient on social interaction attributions. Specifically, in Experiment 2, we replicated the effects of Experiment 1 with hindering actions (i.e., increasing the recipient's cost to obtain the apple).

METHODS

Another 44 paid adult volunteers (21 men and 23 women; mean age = 23 years, age range = 18–26 years) participated in this experiment. The stimuli and procedure were the same as those in Experiment 1, with the following exception: the green agent takes the first apple by passing through the short path in the wall. To manipulate the influence of the blue agent's action on the green agent's costs, we used two conditions: (1) the green agent must use the longer path because the blue agent's act of moving the obstacle blocks the short path, thus acting like a harmful action; and (2) the green agent continues to follow the short path because the blue agent's action does not block it.

As in Experiment 1, there were a total of four conditions, which were constructed with a 2 (actor's goal: weak object-directedness vs. strong object-directedness) by 2 (recipient's change: unchanged cost vs. increased cost) within-participant design. Similar to Experiment 1, to control for the possible presentation order effect, two presentation orders were counterbalanced across participants. Half of the participants were tested with the predetermined order, called the A version (first event: strong object-directedness with unchanged cost on the recipient; second event: strong object-directedness with increased cost on the recipient; third event: weak object-directedness with increased cost on the recipient; fourth event: weak object-directedness with unchanged cost on the recipient); and the other half were tested with the reversed order of the predetermined order, called the B version. To fit the current experimental setting, the last question for measuring social interaction attributions was revised to, "To what extent is the blue agent hindering the green agent?" (rated from 1 = *completely unlikely* to 7 = *completely likely*).

For the question of manipulation check, Figure 4a and Table 3 shows the overall results regarding participants' evaluations of the blue agent's object-directedness. The ANOVA revealed that only the main effect of an actor's goal was significant ($F(1,43) = 15.00$, $p < .001$, $\eta^2_p = .26$). As expected, the results confirmed that participants gave higher ratings of object-directedness in the strong object-directedness condition as compared to the weak object-directedness condition. The main effect of recipients' change did not reach significance ($F(1,43) = 0.09$, $p = .768$, $\eta^2_p < .01$), and neither did the interaction effect ($F(1,43) = 0.36$, $p = .553$, $\eta^2_p = .01$). These effects were not modulated by the presentation order of events (see the supplementary materials for detailed results), indicating that the manipulation about the actor's goal was valid.

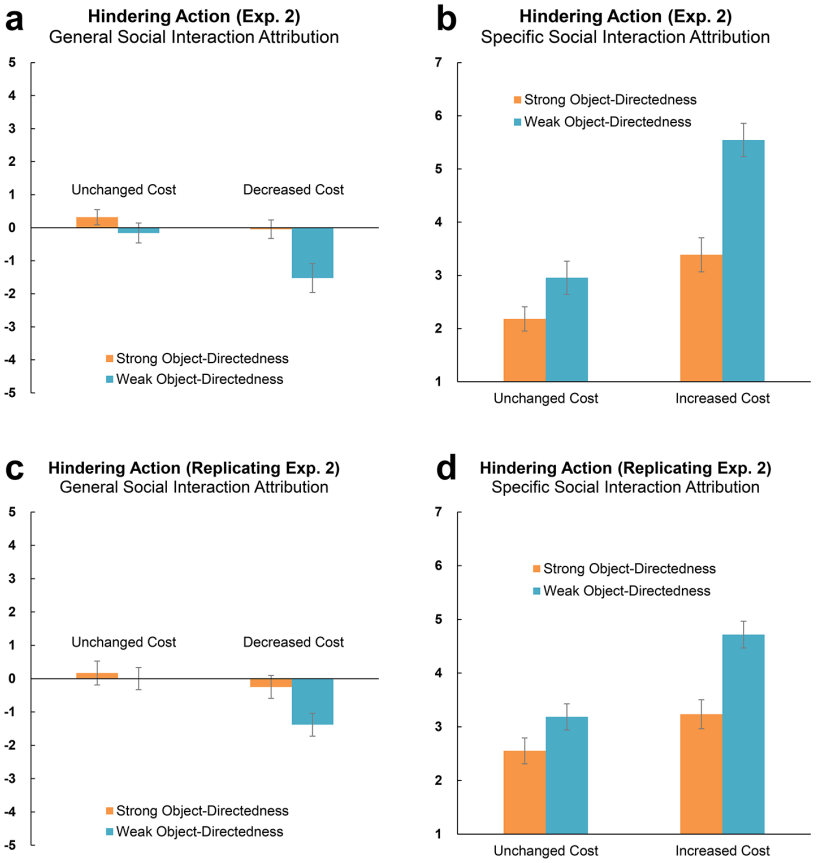

FIGURE 4. Attributions of social interaction goal between the actor (i.e., blue agent) and the recipient (i.e., green agent) under different acting conditions in Experiment 2 and its replication. (a) and (c) represent the averaged rating scores, based on a general understanding of social interaction. A score of -5 means a "strongly negative interaction," 5 means "strongly positive interaction," and 0 means "no social interaction." (b) and (d) represent the averaged rating scores, based on a specific understanding of social interaction. A score of 1 means a "completely unlikely," and 7 means "completely likely." Error bars indicate the standard error (± SE).

RESULTS AND DISCUSSION

Similar to Experiment 1, a three-way ANOVA with actor's goal, recipient's change, and presentation order as factors was employed on each variable of interest to examine the presentation order effect. Similarly, it was found that the factor of presentation order did not produce any significant interaction effect with other factors (see the supplementary materials for detailed results). Hence, in the following analyses, the data with different presentation orders were collapsed.

The overall results for participants' judgments of social interaction between the two agents, based on a general understanding of social interaction, are shown in Figure 4a and Table 3. An ANOVA with actor's goal and recipient's change as factors was then conducted. We found a significant main effect of actor's goal

TABLE 3. Descriptive Results of Different Rating Dimensions for Hindering Action in Experiment 2 (*M*±*SE*)

	Unchanged cost		Decreased cost	
	Strong object-directedness	Weak object-directedness	Strong object-directedness	Weak object-directedness
Actor's object-directed goal	6.02(0.20)	5.16(0.29)	6.18(0.21)	5.11(0.30)
General social interaction attribution	0.32(0.23)	-0.16(0.30)	-0.05(0.28)	-1.52(0.44)
Specific social interaction attribution	2.18(0.23)	2.95(0.31)	3.39(0.32)	5.55(0.31)

($F(1,43) = 10.27$, $p = .003$, $\eta^2_p = .19$), suggesting that when the blue agent had a weak object-directed goal to fetch the apple, participants judged the two agents as having a more negative social interaction than when the blue agent had a strong object-directed goal. The main effect of recipient's change was also significant ($F(1,43) = 12.81$, $p = .001$, $\eta^2_p = .23$), and, importantly, the interaction effect was as well ($F(1,43) = 5.84$, $p = .020$, $\eta^2_p = .13$). Post-hoc simple effects analyses revealed in the increased cost condition, participants evaluated the social interaction as more negative in the weak object-directedness condition than in the strong object-directedness condition ($t(43) = 3.61$, $p = .001$, $d = .54$, 95% CI for mean difference = [-2.30 -0.65]); however, this difference was not found in the unchanged cost condition ($t(43) = 1.48$, $p = .146$, $d = .22$, 95% CI for mean difference = [-1.13 0.17]).

Concerning social interaction judgments based on a specific understanding of social interaction (Figure 4b and Table 3), the ANOVA showed that both main effects—actor's goal ($F(1,43) = 31.43$, $p < .001$, $\eta^2_p = .42$) and recipient's change ($F(1,43) = 35.45$, $p < .001$, $\eta^2_p = .45$)—were significant, as was the interaction effect ($F(1,43) = 9.98$, $p = .003$, $\eta^2_p = .19$). Post-hoc simple effects analyses revealed that, in both the unchanged cost ($t(43) = 3.39$, $p = .002$, $d = .50$, 95% CI for mean difference = [0.47 1.94]) and increased cost ($t(43) = 5.92$, $p < .001$, $d = .89$, 95% CI for mean difference = [1.71 3.47]) conditions, actions that had weak object-directed goals were judged as more indicative of hindering action than were actions that had strong object-directed goals; however, the difference in the former condition was smaller than that in the latter ($t(43) = 3.16$, $p = .003$, $d = .48$).

REPLICATION

Similar to what happened in Experiment 1, we ran a new experiment similar to the previous Experiment 2 with hindering actions, but without asking the first question, and used the Latin square design for counterbalancing with the predetermined order (i.e., first event: strong object-directedness with unchanged cost on the recipient; second event: strong object-directedness with increased cost on the

recipient; third event: weak object-directedness with increased cost on the recipient; fourth event: weak object-directedness with unchanged cost on the recipient). A different sample set of 60 paid adult volunteers (31 females; mean age = 21 years, ranging from 18 to 27 years) participated in this experiment.

Participants' social interaction attributions between the two agents are shown in Figure 4c and Table 4. The ANOVA showed a significant main effect of actor's goal ($F(1,59) = 7.99$, $p = .006$, $\eta^2_p = .12$), suggesting that when the actor had a weak object-directed goal, participants judged the two agents as having a more negative social interaction than when the actor had a strong object-directed goal. The main effect of recipient's change was also significant ($F(1,59) = 10.99$, $p = .002$, $\eta^2_p = .16$), as well as the interaction effect between the variables ($F(1,59) = 6.86$, $p = .011$, $\eta^2_p = .10$). Post-hoc simple effects analyses revealed that, in the increased cost condition, participants evaluated the social interaction as more negative in the weak than in the strong object-directedness condition ($t(59) = 4.38$, $p < .001$, $d = .57$, 95% CI for mean difference = [-1.65 -0.62]); however, this difference was not found in the unchanged cost condition ($t(59) = 0.51$, $p = .612$, $d = 07$, 95% CI for mean difference = [-0.82 0.49]).

Concerning social interaction attributions based on a specific understanding of social interaction (Figure 4d and Table 4), the ANOVA showed that both main effects—actor's goal ($F(1,59) = 27.91$, $p < .001$, $\eta^2_p = .32$) and recipient's change ($F(1,59) = 36.99$, $p < .001$, $\eta^2_p = .39$)—were significant, as was their interaction effect ($F(1,59) = 6.01$, $p = .017$, $\eta^2_p = .09$). Post-hoc simple effects analyses revealed that, in both the unchanged and increased cost conditions, actions that had weak object-directed goals were attributed as more indicative of being hindering actions than those that had strong object-directed goals (unchanged cost: $t(59) = 2.43$, $p = .018$, $d = .09$, 95% CI for mean difference = [0.11 1.15]; increased cost: $t(59) = 5.01$, $p < .001$, $d = .34$, 95% CI for mean difference = [0.94 2.02]); however, the difference between weak and strong object-directed goal conditions in the unchanged cost condition was smaller than that in the increased cost condition ($t(59) = 2.45$, $p = .017$, $d = .32$). None of the above effects were moderated by the presentation order (ps > .145).

These results suggest that when the actor incurred costs on the recipient, the object-directedness of the actor's action had greater influence on attributions of social interaction than when the actor incurred no or rare costs on the recipient. Hence, both measures consistently indicate that the effect of the object-directedness

TABLE 4. Descriptive Results of Different Rating Dimensions for Hindering Action in Replicating Experiment 2 (*M*±*SE*)

	Unchanged cost		Decreased cost	
	Strong object-directedness	Weak object-directedness	Strong object-directedness	Weak object-directedness
General social interaction attribution	0.17(0.36)	0.00(0.33)	-0.25(0.34)	-1.38(0.34)
Specific social interaction attribution	2.55(0.24)	3.18(0.24)	3.23(0.27)	4.72(0.25)

of the actor's goal on social interaction attributions was dependent on whether the actor incurred influence on the recipient.

DISCUSSION

When faced with actions, social goals (i.e., social interactions) may be attributed when observing the desired outcome of an actor influencing a recipient. In addition, each actor might also have a self-interested object-directed goal. We explored how these object-directed goals influenced social interaction attributions. Consistent with our hypothesis, in both helping and hindering contexts, the effect of the object-directedness of the actor's goal on social interaction attributions was dependent on the magnitude of an actor's influence on the recipient. Specifically, when the action was clearly helping/hindering (i.e., it decreased or increased the recipient's cost to reach his/her goal), the agents were judged as having a stronger interaction (either positive or negative) when the blue agent had a strong object-directed goal than when it had a weak goal. This effect became smaller, or even absent, when the actor's action had no influence on the recipient.

These findings accord with previous suggestions that an individual's influence on another person determines observers' social interaction attribution (Choi & Luo, 2015; Tatone et al., 2015; Van de Vondervoort & Hamlin, 2018; Wu et al., 2018). We also extend prior research by showing that this causal relationship is modulated by the directedness of the actor's goal. The less directed an actor is toward achieving an object-related goal, the more directed they appear toward influencing the recipient, at least in the present context. This suggests that observers consider the actor's goal and how their actions influence the recipient jointly, rather than independently, when making social interaction attributions. According to the correspondent inference theory (Jones & Davis, 1965; Jones & McGillis, 1976), how likely the actor is directed to the external object can be treated as behavioral expectancy, as the irrational action to the goal violates the usually assumed rationality; the incurred outcome to the recipient by the actor can be also deemed as the non-common effects to the supposed social goal. As suggested in both previous and current studies, the incompatibility of the action with the usual expectation of the perceiver made the alternative inference more likely to be perceived (Jones & McGillis, 1976), and the uniqueness of the behavioral effects made the corresponding inference more likely to be perceived by observers (Ajzen & Holmes, 1976; Newtson, 1974). More than that, the current study incorporated these two sources of information together, as the results indicated that only when the unique behavioral effect was observed related to the social interaction goal its inference was influenced by the expectancy of rationality with respect to the object-directed goal.

The Knobe effect is a well-known phenomenon that refers to actions with detrimental side effects that are judged as more intentional than actions that produce beneficial side effects (Knobe, 2003; Leslie, Knobe, & Cohen, 2006). In our study, if the actor brought about the influence on the recipient in the strong object-directedness condition, it could be considered a side effect since the strong object-directedness represents that the intention is to approach his/her own object, instead of

to influence others. The comparison of side effects between helping (M = 2.75) and hindering actions (M = 2.95) showed that no difference was attributable to social interaction ($t(86) = 0.52$, $p = .604$, which was based on values of the specific understanding of social interaction); the same was true for the replicated experiments in both Experiments 1 and 2 (helping: M = 3.37; hindering: M = 3.23; $t(118) = 0.72$, $p = .721$). A possible reason for this discrepancy is that most previous research on the Knobe effect adopted textual stimuli to describe actors' intentionality, awareness, and outcome (Knobe, 2003; Leslie et al., 2006; Vonasch & Baumeister, 2017). In contrast, visual stimuli were used in the current study. Thus, participants had to infer all the information from these stimuli, and such inferences might be more difficult to make compared to when text was used. Moreover, literal interpretations may conflict with participants' pragmatic interpretations (Laurent, Clark, & Schweitzer, 2015); hence, the use of textual material might recruit distinct processes as compared to visual materials, thereby leading to diverse results concerning the interpretation of side effects. However, this speculation requires elucidation.

The present findings may be explained by the naïve utility calculus, wherein agents choose goals/intentions and actions to maximize their utilities: the rewards (or beneficial outcomes) they expect to obtain relative to the costs they expect to incur (Jara-Ettinger et al., 2016). Regarding the object-directed goal attribution in the current study, and based on the naïve utility calculus, it is possible to say that observers would deem the utility of an actor's irrational action to be lower when compared with the actor's rational action. Hence, the likelihood of an observer to infer that an irrational action was directed to an external object was lower than that of a rational action. Nevertheless, this situation would be changed for the social interaction goal attribution. Given that social goals are treated as a second-order interpretation of agents' actions, observers utilize this naïve utility calculus to compute everyone's costs and rewards separately; then, they integrate them when evaluating social interactions. Thus, the costs of moving the obstacle and planning a long path to influence others (i.e., weak object-directedness) are compensated by the second-order expectation of bringing about the outcomes of decreasing or increasing the costs of the recipient. Conversely, if these costs are only rational for explaining the actor's first-order expectation (i.e., strong object-directedness), their influence on the recipient is insufficient to support the second-order interpretation of the social interaction. In contrast, when there is no influence between agents, there is no evidence to favor that second-order interpretation. Consequently, the object-directedness of the actor's goal modulates the social interaction attribution only when the influence between agents is present.

In this consideration, where both object-directed and social goals can be attributed to the actor's action, it is necessary to ask which goal attribution comes first, and our results provided a tentative answer for this question. Specifically, regarding the strong object-directedness condition (i.e., rational action), the difference for the social interaction attributions between the two recipients' change conditions was not significant, except when the question about specific social interaction was asked regarding the hindering actions (Experiment 2). In contrast, regarding the weak object-directedness condition (i.e., irrational action), we observed that the

difference for social interaction attributions between the two recipients' change conditions was significant in all experiments and measures. Further, a meta-analysis performed with the JASP software (JASP Team, 2019) confirmed that the effect size of the difference for social interaction attributions between two recipients' change conditions was close to zero in the strong object-directedness condition (95% CI of Cohen's *d* difference = [0.03 0.30]), but its effect size was largely different from zero in the weak object-directedness condition (95%CI of Cohen's *d* difference = [0.47 0.75]). In short, when the object-directed attribution was high, even though the actor produced influence on the recipient, observers were more inclined to attribute to the actor an object-directed goal, not a social interaction goal, mainly owing to its non/rare influence on the recipient—whereas, observers were more inclined to attribute a social interaction goal when the object-directed attribution was low. This dominant role of object-directed attribution was also implied in the study of Schachner and Carey (2013); they found that when the agent's actions (i.e., movements) presented rational actions to achieve a plausible external goal (i.e., putting a star in a box), nearly all participants (i.e., 93% in the third experiment) reasoned that the movements were means to realize the external effect (i.e., external goals); in a contrasting situation (i.e., the agent's actions were irrational to achieve an external object), almost half of the participants concluded that the movement was the agent's goal. In other words, based on their and our results, people's tendency to explicitly infer object-directedness goals from various situations may be stronger than their tendency to infer a situational alternative goal (e.g., the social interaction goal and the movement-based goal of the aforementioned study).

Both the ability of inferring an object-directed goal and a social interaction goal appear in the early stages of development (Choi & Luo, 2015; Feiman et al., 2015; Gergely & Csibra, 2003; Hamlin et al., 2007; Hamlin et al., 2011; Stahl & Feigenson, 2014; Tatone, 2017; Tatone et al., 2015; Woodward, 1998). The current study indicated that, in adults, these two abilities interact when there is a requirement to make explicit attributions, but more studies are warranted to better understand how this develops. Nevertheless, we speculated that the explicit social interaction attribution may follow almost the same developmental trajectories of moral judgment (Cushman, Sheketoff, Wharton, & Carey, 2013), if treating attributed social interactions as the basis for moral judgment (Gray, Schein, & Ward, 2014). Specifically, at the early stage, infants mainly rely on the outcome of the actor's influence on the recipient; with the development of the theory of mind, children and adults are able to explicitly understand the mental states of others, and accordingly shift to intention-based attributions (e.g., the actor may not want to approach the apple, but aim to help others).

One limitation of the present research was that the actions were simulated by cartoon animations, so they were not very realistic. To answer theoretical questions, this method can strictly control for confounding factors, and has been largely documented to induce the perception of animacy, agency, intent, social interaction, and so forth, even at the very beginning of our development (Castelli, Happé, Frith, & Frith, 2000; Hamlin et al., 2007; Joyal, Neveu, Boukhalfi, Jackson,

& Renaud, 2018; Tatone et al., 2015). With the animations utilized in this study, we tried to simulate the opening/closing of a door while helping or hindering others, and they were designed to be able to operationalize two independent variables (i.e., actor's goal and recipient's change). Regarding the real world, the focus on such explicit inferences for a social interaction may not be so obvious, but, sometimes, they are present; for example, when gossiping over a person's disposition when he/she opens or closes a door and how this action changes other's states. Still, the actual social perception (e.g., inferences about social goals and the motivation behind this inference) in the real world is more complex, and the current method draws some limitations regarding broader conclusions in this regard (like how attributions to a real-world social interaction would work). Further research should employ ecologically valid stimuli (e.g., videos depicting real people) and set the contexts for social interaction attribution within a real-life story, thus testing the external validity of our conclusion and allowing for broader generalizations.

CONCLUSION

This study explored how the object-directedness of an actor's actions affects social interaction attributions and whether/how such a relationship is modulated by his/her action's outcomes. Through visual animations, the object-directedness was operationalized as an actor's choice of whether to take a rational (i.e., low-cost) path to approach an external object, and the action outcomes was manipulated as whether the actor's action changed the recipient's cost (i.e., traveled path) to the recipient's simple goal. We demonstrated that the object-directedness of an actor's goal plays an important role in making social interaction attributions; this role is highly dependent on the magnitude of the actor's influence on the recipient.

REFERENCES

Ajzen, I., & Holmes, W. H. (1976). Uniqueness of behavioral effects in causal attribution. *Journal of Personality*, *44*(1), 98–108.

Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour, 1*, 0064.

Castelli, F., Happé, F., Frith, U., & Frith, C. (2000). Movement and mind: A functional imaging study of perception and interpretation of complex intentional movement patterns. *Neuroimage, 12*(3), 314–325.

Choi, Y. J., & Luo, Y. (2015). 13-month-olds' understanding of social interactions. *Psychological Science*, *26*(3), 274–283.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Cushman, F., Sheketoff, R., Wharton, S., & Carey, S. (2013). The development of intent-based moral judgment. *Cognition, 127*(1), 6–21.

Dennett, D. C. (1989). *The intentional stance*. Cambridge, MA: MIT Press.

Feiman, R., Carey, S., & Cushman, F. (2015). Infants' representations of others' goals: Representing approach over avoidance. *Cognition*, *136*, 204–214.

Gao, T., Newman, G. E., & Scholl, B. J. (2009). The psychophysics of chasing: A case study in the perception of animacy. *Cognitive Psychology*, *59*(2), 154–179.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences*, *7*(7), 287–292.

Gillin, J. L., & Gillin, J. P. (1942). *An introduction to sociology*. New York: Macmillan.

Gray, K., Schein, C., & Ward, A. F. (2014). The myth of harmless wrongs in moral cognition: Automatic dyadic completion from sin to suffering. *Journal of Experimental Psychology General*, *143*(4), 1600–1615.

Hamlin, J. K., Ullman, T., Tenenbaum, J., Goodman, N., & Baker, C. (2013). The mentalistic basis of core social cognition: Experiments in preverbal infants and a computational model. *Developmental science, 16(2),* 209–226.

Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, *450*(7169), 557–559.

Hamlin, J. K., Wynn, K., Bloom, P., & Mahajan, N. (2011). How infants and toddlers react to antisocial others. *Proceedings of the National Academy of Sciences*, *108*(50), 19931–19936.

Hinde, R. A. (1976). Interactions, relationships and social structure. *Man*, *11*(1), 1–17.

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, *20*(8), 589–604.

JASP Team. (2019). JASP (Version 0.11.1) [Computer software].

Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. *Advances in Experimental Social Psychology, 2*, 219–266.

Jones, E. E., & McGillis, D. (1976). Correspondent inferences and the attribution cube: A comparative reappraisal. In J. Harvey, W. Ickes, & R. Kidd (Eds.), *New directions in attribution research* (vol. 1; pp. 389–420). Hillsdale, NJ: Erlbaum.

Joyal, C. C., Neveu, S.-M., Boukhalfi, T., Jackson, P. L., & Renaud, P. (2018). Suppression of sensorimotor Alpha power associated with pain expressed by an avatar: A preliminary EEG study. *Frontiers in Human Neuroscience, 12*, 273.

Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis*, *63*(279), 190–194.

Knoblich, G., & Sebanz, N. (2008). Evolving intentions for social interaction: From entrainment to joint action. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *363*(1499), 2021–2031.

Laurent, S. M., Clark, B. A. M., & Schweitzer, K. A. (2015). Why side-effect outcomes do not affect intuitions about intentional actions: Properly shifting the focus from intentional outcomes back to intentional actions. *Journal of Personality and Social Psychology, 108*(1), 18–36.

Leslie, A. M., Knobe, J., & Cohen, A. (2006). Acting intentionally and the side-effect effect: Theory of mind and moral judgment. *Psychological Science*, 17(5), 421–427.

Malle, B. F. (2004). *How the mind explains behavior: Folk explanations, meaning, and social interaction*. Cambridge, MA: MIT Press.

Marien, H., Aarts, H., & Custers, R. (2015). The interactive role of action-outcome learning and positive affective information in motivating human goal-directed behavior. *Motivation Science*, *1*(3), 165.

Newtson, D. (1974). Dispositional inference from effects of actions: Effects chosen and effects forgone. *Journal of Experimental Social Psychology, 10*(5), 489–496.

Schachner, A., & Carey, S. (2013). Reasoning about "irrational" actions: When intentional movements cannot be explained, the movements themselves are seen as the goal. *Cognition*, *129*(2), 309–327.

Stahl, A. E., & Feigenson, L. (2014). Social knowledge facilitates chunking in infancy. *Child Development*, *85*(4), 1477–1490.

Tatone, D. (2017). *The naïve sociology of resource transfer* (doctoral dissertation, Central European University).

Tatone, D., Geraci, A., & Csibra, G. (2015). Giving and taking: Representational building blocks of active resource-transfer events in human infants. *Cognition, 137,* 47–62.

Ullman, T., Baker, C., Macindoe, O., Evans, O., Goodman, N., & Tenenbaum, J. B. (2009). Help or hinder: Bayesian models of social goal inference. In *Advances in neural information processing systems* (pp. 1874–1882).

Van de Vondervoort, J. W., & Hamlin, J. K. (2018). The early emergence of sociomoral evaluation: Infants prefer proso-

cial others. *Current Opinion in Psychology*, *20*, 77–81.

Vonasch, A. J., & Baumeister, R. F. (2017). Unjustified side effects were strongly intended: Taboo tradeoffs and the side-effect effect. *Journal of Experimental Social Psychology*, *68*, 83–92.

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1–34.

Wu, X., Hua, R., Yang, Z., & Yin, J. (2018). The influence of intention and outcome on evaluations of social interaction. *Acta Psychologica*, *182*, 75–81.