



Original Articles

Social evaluation of intentional, truly accidental, and negligently accidental helpers and harmers by 10-month-old infants



Brandon M. Woo, Conor M. Steckler, Doan T. Le, J. Kiley Hamlin*

Department of Psychology, University of British Columbia, Canada

ARTICLE INFO

Article history:

Received 7 July 2016

Revised 28 June 2017

Accepted 28 June 2017

Available online 5 July 2017

Keywords:

Infant development

Social cognition

Intention

Negligence

ABSTRACT

Whereas adults largely base their evaluations of others' actions on others' intentions, a host of research in developmental psychology suggests that younger children privilege outcome over intention, leading them to condemn accidental harm. To date, this question has been examined only with children capable of language production. In the current studies, we utilized a non-linguistic puppet show paradigm to examine the evaluation of intentional and accidental acts of helping or harming in 10-month-old infants. In Experiment 1 ($n = 64$), infants preferred intentional over accidental helpers but accidental over intentional harmers, suggestive that by this age infants incorporate information about others' intentions into their social evaluations. In Experiment 2 ($n = 64$), infants did not distinguish "negligently" accidental from intentional helpers or harmers, suggestive that infants may find negligent accidents somewhat intentional. In Experiment 3 ($n = 64$), we found that infants preferred truly accidental over negligently accidental harmers, but did not reliably distinguish negligently accidental from truly accidental helpers, consistent with past work with adults and children suggestive that humans are particularly sensitive to negligently accidental harm. Together, these results imply that infants engage in intention-based social evaluation of those who help and harm accidentally, so long as those accidents do not stem from negligence.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Who is more blameworthy: Chris, who intentionally poisoned and killed his mother, or Ben, who unknowingly gave his mother poison, accidentally killing her? Beginning with Piaget (1965/1932), a host of research in developmental psychology suggests that how individuals judge intentional versus accidental acts changes with age: Whereas adults assign praise and blame based primarily on mental states like possessing helpful or harmful intent (often referred to as making the "intention-outcome distinction;" Cushman, 2008; Cushman, Young, & Hauser, 2006; Moran et al., 2011; Ohtsubo, 2007), children under age 4 adhere primarily to an outcome rule; for example condemning acts with harmful outcomes even if harm was unintended (Baird & Astington, 2004; Cushman, Sheketoff, Wharton, & Carey, 2013; Margoni & Surian, 2016; Margoni & Surian, 2017; Shultz, Wright, & Schleifer, 1986; Zelazo, Helwig, & Lau, 1996). Although research suggests that reducing processing demands improves children's ability to incorporate mental states into explicit moral judgments (Armsby, 1971;

Farnill, 1974; Yuill & Perner, 1988), even with these reduced demands there remains significant evidence that the ability to incorporate mental states into moral judgments improves with age. This developmental shift has been taken to suggest either that moral judgments undergo fundamental conceptual change over development (Cushman et al., 2013; Kohlberg, 1969; Piaget, 1965/1932), or that changes in non-moral domains like theory of mind and executive function are critical for privileging intention over outcome in moral judgment (Chandler, Sokol, & Hallett, 2001; Grueneich, 1982; Killen, Lynn Mulvey, Richardson, Jampol, & Woodward, 2011; Zelazo et al., 1996; see Margoni & Surian, 2016).

Past research into the development of the intention-outcome distinction has largely utilized explicit verbal measures to explore children's third-party moral judgments (cf. Vaish, Carpenter, & Tomasello, 2009). Critically, these measures constrain research to children who are capable of language production. To address this constraint, researchers have begun to utilize non-linguistic "puppet shows," in which animated puppet agents perform simple helpful and harmful actions toward third-parties while preverbal infants and toddlers look on; relative preference for the helpful versus harmful puppets is subsequently determined based on which puppets infants and toddlers look at, reach toward, or give

* Corresponding author at: University of British Columbia, Department of Psychology, 2136 West Mall, Vancouver, BC V6T 1Z4, Canada.

E-mail address: kiley.hamlin@psych.ubc.ca (J.K. Hamlin).

objects to. Studies like these have demonstrated that from early in the first year infants' social evaluations are consistent with older children's and adults' moral ones: Infants prefer helpful agents who previously facilitated others' goals over harmful agents who blocked them (e.g., Buon et al., 2014; Hamlin, 2015; Hamlin & Wynn, 2011; Hamlin, Wynn, & Bloom, 2007; Scola, Holvoet, Arciszewski, & Picard, 2015; Steckler, Woo, & Hamlin, 2017; Tasimi & Wynn, 2016; cf. Salvadori et al., 2015). By late in the first year infants expect that helped and harmed individuals will prefer their helpers as well (Fawcett & Liszkowski, 2012; Hamlin et al., 2007; Kuhlmeier, Wynn, & Bloom, 2003).

Although such studies suggest that preverbal infants evaluate helpful and harmful actions, it is unclear exactly how these evaluations are generated. Specifically, infants' evaluations could be driven primarily by an analysis of on helpers and harmers' mental states, by their associated outcomes, or both. To disentangle these possibilities, recent studies have begun to explore whether infants incorporate others' mental states into their social evaluations. One common situation in which helpers' and harmers' intentions are distinguishable from the outcomes they are associated with is when they attempt, but fail, to help and harm – in these cases positive intentions are associated with negative outcomes, whereas negative intentions are associated with positive outcomes. In one set of studies examining infants' evaluations of failed attempters to help and harm, 8-month-olds consistently preferred characters who demonstrated positive intent, irrespective of outcomes (in contrast, 5-month-olds failed to distinguish characters in any contrasts involving failed attempts; Hamlin, 2013). Relatedly, a recent study demonstrated that valenced intentions also influence somewhat older infants' expectations for interactions between third-parties: By 16-months infants expect third-parties to approach individuals who previously tried to help them, even if the would-be helpers failed and so outcomes were negative (Lee, Yun, Kim, & Song, 2015). Together, these results suggest that preverbal infants' evaluate attempted help and harm mentalistically.

Like attempted help and harm, accidents involve a conflict between intention and outcome. Critically, however, attempts and accidents can be quite different: Whereas in attempted help and harm intentions are central while outcomes occur peripherally, in accidents outcomes are central while there is an *absence* of intention. To illustrate, in an attempted harm scenario utilized by Cushman et al. (2013), a boy attempts to break a mirror with a ball (clear negative intention), but fails and the ball lands in a bin (neutral or mildly positive outcome). In a comparable accidental harm scenario, a boy accidentally breaks a mirror (clear negative outcome) with a ball intended for a bin (neutral or mildly positive intention). This differential salience of intention and outcome in attempts versus accidents has led to the argument that intentions are less discernible in accidents than in attempts, and so mentalistically evaluating accidents may be particularly difficult¹ (Young & Saxe, 2009). Supporting this possibility, Young, Cushman, Hauser, and Saxe (2007) found that relative to attempted harm, accidental harm activates more brain regions associated with cognitive conflict, suggesting that evaluating accidents takes effort. These results suggest that even if infants are capable of mentalistically evaluating attempts, they may nevertheless fail to mentalistically evaluate accidents. This pattern would support a developmental trajectory wherein mental states play an increasingly privileged role over outcomes in the development of social and moral evaluation.

Despite this possibility, other findings suggest that infants may be able to process accidental acts around the same time they pro-

cess attempts. For example, infants infer intention from visual and prosodic cues as early as 14 months (Behne, Carpenter, & Tomasello, 2005; Sakkalou & Gattis, 2012; see also Carpenter, Akhtar, & Tomasello, 1998), and by 10 months recognize the role of epistemic states such as knowledge versus ignorance in socio-moral action (e.g., Meristo & Surian, 2013; see also Choi & Luo, 2015; Sloane, Baillargeon, & Premack, 2012). These abilities may inform evaluations of agents who lack valenced intention. Furthermore, infants infer how others will treat agents who harm intentionally versus accidentally: 13-month-olds expect third-parties to avoid intentional, but not accidental, harmers (Choi & Luo, 2015). Finally, 10-month-olds may understand the role of knowledge in establishing helpful or harmful intentions: 10-month-olds prefer a puppet who facilitated a goal over one who blocked it, but only if each puppet was previously aware of the goal (Hamlin, Ullman, Tenenbaum, Goodman, & Baker, 2013). Because agents presumably cannot intend to help or harm third-parties without first being aware of the third-parties' goals (Shultz et al., 1986), this result suggests that preverbal infants may be sensitive to the lack of intent characteristic of accidental acts.

Despite this evidence, each of the studies reviewed above suggesting that infants recognize the role of epistemic states in socio-moral action rely on evidence from *null effects*: To date infants have simply failed to distinguish characters or outcomes anytime accidents are involved. Because these null effects could reflect many underlying causes including a failure to process accidents at all, additional research is necessary to determine whether infants can in fact mentalistically evaluate accidental helping and/or harming.

Indeed, if infants have such a strong grasp of social intention and epistemic states, then why has past research found that young children fail to distinguish accidental and intentional outcomes, instead adhering to an outcome rule? One intriguing possibility, recently raised by Nobes, Panagiotaki, and Pawson (2009), is that past research has underestimated children's abilities to mentalistically evaluate accidents by confounding truly accidental harm with negligently accidental harm. Nuñez, Laurent, and Gray (2014) define negligently accidental harm as a situation in which "a person fails to exercise a reasonable degree of care that results in an unintended mishap or injury." (p. 58). They contrast negligently accidental with truly accidental harm, defined as "an unexpected mishap or injury that a person could not have predicted." (p. 58). Under these definitions, the awareness or foreseeability of the possible valenced consequences of one's actions renders negligent harm at least somewhat intentional: A negligently accidental harmer chooses to act carelessly despite knowing a harmful consequence that could result. Indeed, research with both adults and children indicate that foreseen, negligent accidents are viewed as both distinct from and more intentional than similar but unforeseen, true accidents (Monroe & Malle, 2017; Nobes et al., 2009; Nuñez et al., 2014).

The distinction between true and negligent accidents could help to explain cases in the literature in which young children appear to adhere to an outcome rule by condemning accidental acts of harm. For example, in Cushman et al.'s (2013) depiction of accidental harm, a boy "accidentally breaks a mirror when he throws a ball towards the bin where it belongs" (p. 10). Notably, young children are more likely to describe this act as blameworthy, whereas older children see it as relatively more blameless. However, based on the awareness definition of negligence above, the throwing act is potentially blameworthy: The boy is presumably aware that the mirror is near the bin, and so should have known that a negative outcome could result if he did not take special care to avoid it. That he chose not to exercise this care indicates that causing the negative outcome may have been at least a little bit intentional, and therefore blameworthy (indeed, children likely frequently experi-

¹ Of course, the reason for a failed attempt might be something accidental (e.g., tripping on the way to put poison in someone's food); nevertheless, it is the valenced intention to help or to harm that takes center stage in attempted help/harm, which is not the case in accidental help/harm.

ence being admonished for not being more careful in their own lives). Supporting this possibility, Nobes et al. (2009) demonstrated that when vignettes clearly specify that accidental harms were not due to negligence or lack of care (that is, they were truly accidental), even young children tend not to find accidental harmers blameworthy.

To date, children's mentalistic judgments about helping have received much less attention than their mentalistic judgments about harming (but see Leslie, Knobe, & Cohen, 2006; Margoni & Surian, 2017). Noting this asymmetry, Margoni and Surian (2017) created vignettes to explore whether children exhibit the same outcome-to-intent shift when evaluating accidental help as they do when evaluating accidental harm, and found that children under 5 years of age do indeed praise accidentally helpful acts despite their presumed lack of intention, mirroring the developmental shift observed in the harm domain. It seems possible that this outcome-to-intent shift also reflects early negligence judgments of accidentally helpful acts (Nobes et al., 2009); however, to date no studies have systematically explored this possibility. Indeed, the English word "negligent" is typically only used to describe foreseen negative side-effects; we are aware of no suitable alternative word for foreseen positive side-effects. Although this raises the possibility that foreseen positive side-effects are viewed as fundamentally distinct from foreseen negative ones (e.g., Knobe, 2003; Leslie et al., 2006), for the purposes of the current studies we use the phrase "negligently accidental" to refer to both foreseen positive and foreseen negative side-effects, and "truly accidental" to refer to both unforeseen positive and unforeseen negative side-effects. When a side-effect is negative, we refer to negligently or truly accidental harm; when a side-effect is positive, we refer to negligently or truly accidental help.

In three Experiments, we examine whether and to what extent infants differentiate between truly accidental, intentional, and negligently accidental acts of helping and harming. In each Experiment infants viewed puppet shows depicting two "Actors" and a "Protagonist" with an unfulfilled goal. Across conditions, Actors either both facilitated or both blocked a protagonist from achieving its goal (helping and harming conditions; between-subjects); critically, Actors differed in their intention to help or to harm the

protagonist. In Experiment 1 infants chose between intentional and truly accidental helpers or harmers, in Experiment 2 they chose between intentional and negligently accidental helpers or harmers, and in Experiment 3 they chose between truly and negligently accidental helpers or harmers. Across Experiments, we hypothesized that infants would prefer the relatively more intentional helpers, but the relatively less intentional harmers; we anticipated that the magnitude of these preferences might vary across conditions in Experiments 2 and 3, depending on how infants interpret negligently accidental acts. We chose to study 10-month-olds because past research from our lab has demonstrated that 10-month-olds' social evaluations are sensitive to helpers' and harmers' awareness of a third-party's goal (Hamlin et al., 2013), the critical feature that we manipulated to establish whether accidents were true or negligent.

2. Experiment 1: Intentional versus truly accidental help and harm

2.1. Method

2.1.1. Participants

Sixty-four full-term 10-month-olds participated (31 girls; mean age = 10 months, 3 days; range = 9;13–10;20). Subjects were randomly assigned to either the helping condition ($n = 32$; 17 girls; mean age = 10 months, 3 days; range = 9;14–10;20) or the harming condition ($n = 32$; 14 girls; mean age = 10 months, 4 days; range = 9;13–10;17).

An additional 47 babies began the study but were not included in the final sample. Most either fussed out ($n = 15$) or failed to choose between puppets ($n = 13$). This exclusion rate is higher than average for our studies, and may stem from our procedure, which involves a heavy shelf smacking the stage during each event. The sound upset many infants immediately; some infants subsequently refused to choose puppets. The remaining 19 exclusions were procedural errors ($n = 11$), parental interference ($n = 7$), and equipment failure ($n = 1$). Blind experimenters determined exclusions using pre-set criteria.

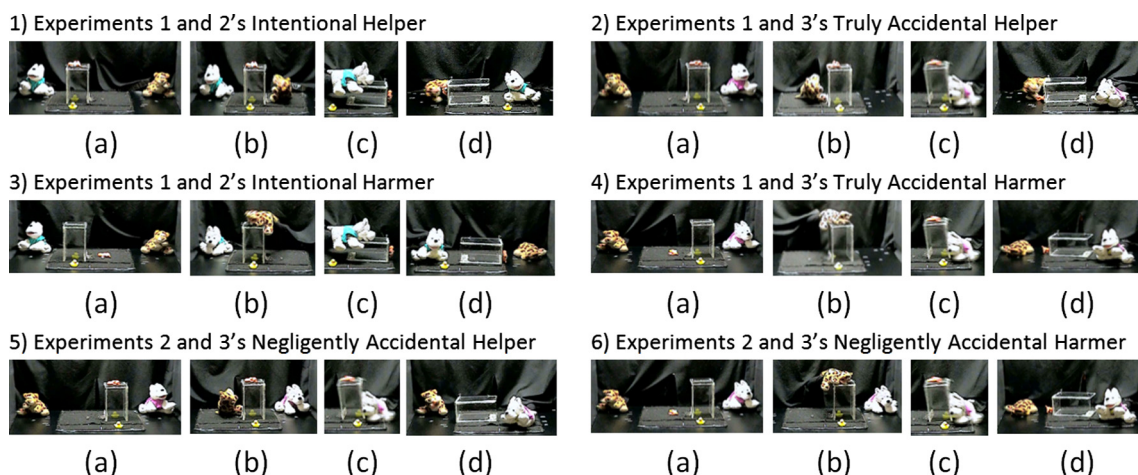



Fig. 1. Stimuli in Experiments 1, 2, and 3. Line 1: Intentional Helping Event, (a) Moose is on shelf; (b) Protagonist tries to reach moose as Intentional Helper is onstage; (c) Intentional Helper pushes shelf over; (d) Protagonist picks up moose before Intentional Helper grasps the duck. Line 2: Truly Accidental Helping Event, (a) Moose is on shelf; (b) Protagonist tries to reach moose as Truly Accidental Helper is off-stage; (c) Truly Accidental Helper knocks shelf over on its way to the duck; (d) Protagonist picks up moose after Truly Accidental Helper grasps the duck. Line 3: Intentional Harming Event, (a) Moose is on the ground; (b) Protagonist tries and eventually succeeds at placing moose on shelf as Intentional Harmer is onstage; (c) Intentional Harmer pushes the shelf over; (d) Protagonist lays down before Intentional Harmer grasps the duck. Line 4: Truly Accidental Harming Event, (a) Moose is on the ground; (b) Protagonist tries and eventually succeeds at placing moose on shelf as Truly Accidental Harmer is off-stage; (c) Truly Accidental Harmer knocks the shelf over on its way to the duck; (d) Protagonist lays down after the Truly Accidental Harmer grasps the duck. Line 5: Negligently Accidental Helping Event; this is the same as the Truly Accidental Helping Event, except that in (b), Negligently Accidental Helper is onstage as Protagonist tries to reach moose. Line 6: Negligently Accidental Harming Event in Experiment 2; this is the same as the Truly Accidental Harming Event, except that in (b), Negligently Accidental Harmer is onstage as Protagonist tries and eventually succeeds at placing moose on shelf.

2.2. Procedure

Infants sat on their parent's lap before a table (W:122 cm) with a curtain (85 cm from infants) that occluded the puppet show when lowered. Parents were instructed to sit quietly and not influence their infants; in Experiments 1 and 2 parents had their eyes open during the puppet shows. Each infant viewed intentional and accidental events in alternation for a total of four events.

Events consisted of a novel puppet scenario involving a large clear plexiglass shelf with a small moose toy on top. For each event in all conditions, the shelf was knocked over, resulting in the moose landing near the ground; events across conditions were therefore highly similar. Critically, whether this outcome was positive or negative for the Protagonist differed across conditions. In the helping condition, this outcome was positive: The Protagonist had previously unsuccessfully attempted to retrieve the moose from the top of the shelf; knocking the shelf over therefore allowed him to achieve his goal. In the harming condition, this outcome was negative: The Protagonist had previously effortfully placed the moose on top of the shelf; knocking the shelf over therefore thwarted his goal. Event sequences are outlined below (see schematic in Fig. 1, and videos in Supplementary Materials).

Supplementary 1. Stimuli Events.



Supplementary 2. Choice.



2.2.1. Familiarization events: Helping condition

All events began identically: A curtain rose, revealing a plexiglass shelf (W:15 cm, L:15 cm; H:28 cm). A small orange stuffed

moose (H:5 cm) sat on top of the shelf and a yellow rubber duck (H:5 cm) rested in front of the shelf. The shelf was positioned near the middle of the stage, 13 cm off-center; the shelf moved 26 cm to the left or right between events so it was on different sides of the stage for intentional and truly accidental events. A white dog puppet (the Intentional or Truly Accidental Actor; 25 cm high, wearing a teal or purple shirt) and a tiger puppet (the Protagonist; 18 cm high) sat at the stage's rear corners; the Actor sat on the side closer to the shelf. At the start of every event, the Protagonist and the Actor turned to look at each other and then turned to face forward. The following describes what happened next for each event type.

2.2.1.1. Helping condition: Intentional events. While the Intentional Actor rested at his rear corner, the Protagonist moved to the side of the shelf nearest to him and strained to "look" at the duck resting in front of the shelf, before turning to face forward. He then made three attempts to reach the moose. During each attempt, the Protagonist turned, "looked" toward the top of the shelf, and jumped twice toward the moose, making grasping motions as though reaching for it, getting closer with each attempt. After his third failed attempt, the Protagonist lay down on the floor as though giving up, and returned to his original position at the rear corner of the stage. Because the Actor was onstage during the Protagonist's failed attempts, he was presumably aware of the Protagonist's goal.

After the Protagonist returned to the corner, the Intentional Actor turned to look at the Protagonist. He ran to the side of the shelf opposite to where the Protagonist had been, turned to face the shelf, and paused. He then raised his arms and pushed the shelf over, so that the moose was near the ground. The Actor then paused next to the fallen shelf. The Protagonist jumped once, ran forward, picked up the moose, and jumped twice as though to celebrate having achieved his goal. Finally, the Actor ran to the front of the shelf, picked up the duck, and ran off-stage with it.

2.2.1.2. Helping condition: Truly Accidental events. Accidental events began identically to intentional events, with the Actor and the Protagonist turning to look at each other. However, rather than wait onstage during the Protagonist's failed attempts as the Intentional Actor had done, the Truly Accidental Actor left the stage, remaining off-stage throughout the Protagonist's attempts. After the Protagonist had given up and returned to his original position, the Truly Accidental Actor returned, presumably ignorant of the Protagonist's goal. The Truly Accidental Actor then turned to acknowledge the Protagonist, as the Intentional Actor had done during intentional events. The Truly Accidental Actor then ran forward and knocked over the shelf, but, critically, his actions suggested that knocking the shelf over was merely a side-effect of his primary goal of retrieving the duck. Specifically, the Actor turned and looked toward the duck resting in front of the shelf, and ran in a straight line toward the duck and immediately grasped it. On the Actor's way past the shelf, he clumsily knocked into its base with his shoulder, knocking the shelf over. He then paused grasping the duck, and the Protagonist jumped once, ran forward, picked up the moose, and jumped twice, as though to celebrate having achieved his goal. Finally, the Truly Accidental Actor sat up and ran off-stage with the duck.

2.2.2. Familiarization events: Harming condition

The harming condition's shelf, moose, duck, dogs, and tiger were the same as those of the helping condition. Harming events began identically to helping events, except for the moose's initial position. Specifically, the moose was not on top of the shelf, but on the floor next to the shelf, in front of the Protagonist. At the start

of each event, the Protagonist and the Actor turned to look at each other, and then turned to face forward. The following describes what happened next for each event type.

2.2.2.1. Harming condition: Intentional events. While the Intentional Actor rested at his rear corner, the Protagonist ran forward, picked up the moose, and strained to “look” at the duck resting in front of the shelf. He then made three attempts to place the moose on the shelf. During each attempt, the Protagonist turned and “looked” toward the top of the shelf and then jumped twice toward it, making stretching motions as though trying to place the moose on the top, getting closer with each attempt. During the third attempt’s second jump, the Protagonist placed the moose on the shelf, achieving his goal. The Protagonist then jumped twice as though to celebrate, and returned to his original position at the rear corner of the stage. Because the Intentional Actor was onstage during the Protagonist’s attempts, he was presumably aware of the Protagonist’s goal.

After the Protagonist returned to the corner, the Intentional Actor turned to look at the Protagonist, and then intentionally knocked over the shelf as in the helping condition. The Protagonist jumped once, lay down on the floor (not grasping the toy) as though unhappy, and paused. Finally, the Intentional Actor ran to the front of the shelf, picked up the duck, and ran off-stage with it.

2.2.2.2. Harming condition: Truly Accidental events. Truly accidental harming events began identically to intentional harming events, with the Actor and the Protagonist turning to look at each other. However, rather than wait onstage while the Protagonist attempted to put the moose on the shelf as the Intentional Actor had done, the Truly Accidental Actor left the stage, and remained off-stage throughout the Protagonist’s attempts and eventual success. After the Protagonist succeeded and returned to his original position, the Truly Accidental Actor returned, presumably ignorant of the Protagonist’s goal.

The Truly Accidental Actor then turned to acknowledge the Protagonist, as the Intentional Actor had done during intentional events. He then ran in a straight line toward the duck and immediately grasped it, knocking over the shelf on his way as a side-effect as in the helping condition. The Protagonist jumped once, lay down on the floor (not grasping the toy) as though unhappy, and paused. Finally, the Truly Accidental Actor sat up and ran off-stage with the duck.

2.2.3. Coding

During all events across conditions, action paused once the Actor was off-stage. Infants’ looking was recorded from this point by an observer using the coding program jHab (Casstevens, 2007), until infants looked away for 2 consecutive seconds or 30 s elapsed. The observer watched infants through a live video feed in another room, could not see the events, and was unaware of condition.

2.2.4. Choice

Parents turned 90 degrees to the right so that they were no longer facing the stage, and closed their eyes. An experimenter blind to the puppets’ identities and condition kneeled in front of infants and held the Intentional and Truly Accidental Actors up, approximately 30 cm apart and initially out of the infants’ reach. Infants were required to look at both puppets and back to the experimenter; the puppets were then moved within reach. The experimenter determined a choice as the first puppet infants touched via a visually guided reach (a touch preceded by a look; see videos in Supplementary Materials).

2.2.5. Counterbalancing

The following were counterbalanced across infants: (1) condition (helping/harming); (2) intentional puppet shirt color (teal/purple); (3) intentional puppet order (first/s); (4) intentional puppet familiarization side (left/right); and (5) intentional puppet choice side (left/right).

2.3. Results

All reported *p*-values are two-tailed. As predicted, patterns of choice for Intentional and Truly Accidental Actors differed significantly across helping and harming conditions ($\chi^2(2, N = 64) = 10.57, p = 0.001$; see Fig. 2). Infants in the helping condition preferred the Intentional Helper over the Truly Accidental Helper (22/32 infants chose the Intentional Helper, binomial $p = 0.050$) whereas infants in the harming condition preferred the Truly Accidental Harmer over the Intentional Harmer (23/32 of infants chose the Truly Accidental Harmer, binomial $p = 0.020$). There was no effect of event order, puppet color, or puppet side on any comparison ($ps > 0.400$).

To determine whether infants attended differently to intentional versus truly accidental helping or harming, we examined each event type’s looking times (see Table 1 for summary). Attention to individual event types did not differ within or across conditions (repeated-measures ANOVA; $ps > 0.900$).

2.4. Discussion

In Experiment 1, infants preferred an Intentional over a Truly Accidental Helper, but a Truly Accidental over an Intentional




Fig. 2. Percentage of infants choosing Intentional Actors and Truly Accidental Actors in Experiment 1 as a function of whether the actors were helping or harming the Protagonist’s goal. Asterisks indicate significant differences ($p \leq 0.05$; $**p \leq 0.01$).

Table 1

Infants mean total looking times (s) during familiarization in Experiments 1 and 2.

	Attention to intentional events	Attention to accidental events
<i>Experiment 1</i>		
Helping condition	15.08 (2.02)	15.17 (1.79)
Harming condition	16.65 (1.85)	16.47 (1.89)
Helping and harming conditions	15.87 (1.37)	15.82 (1.30)
<i>Experiment 2</i>		
Helping condition	19.13 (1.94)	18.39 (1.94)
Harming condition	17.28 (1.56)	16.71 (1.64)
Helping and harming conditions	18.21 (1.25)	17.55 (1.27)

Note. All numbers in parentheses are standard errors.

Harmer. These results suggest that infants' social evaluations are sensitive to the lack of intent characteristic of truly accidental acts, and provide the first evidence for infants' ability to mentalistically evaluate accidents in sociomoral action using an experimental design that does not rely on null effects. Notably, because all events in all conditions in Experiment 1 included an agent knocking over a shelf, resulting in either a positive or a negative outcome, it is unlikely that infants' choices were based on low-level features of our displays. Furthermore, by preferring the Intentional Actor in the helping condition but the Truly Accidental Actor in the harming condition, infants' choices did not reflect simple preferences for actors who performed actions in a particular way (e.g., knocking over a shelf with the hands versus the shoulder), nor for actors they were exposed to for more versus less time (the Intentional Actor was onstage longer in both the helping and harming condition).

If results from Experiment 1 suggest that infants make the intention-outcome distinction in their evaluation of accidental helpers and harmers, why have past studies found that children have such difficulty doing so until after their 4th birthdays? As discussed in the introduction and in line with Nobes et al. (2009), we hypothesize that young children may treat "accidental" acts of help and harm as "negligent;" praising or blaming accidental helpers and harmers to the extent that they could foresee the positive and negative side-effects that would result from their actions (though we note that Nobes et al. (2009) defined negligence as "carelessness," we found this more difficult to operationalize in displays designed for infants and so chose to manipulate foreseeability). We explored this possibility in Experiment 2, by having infants choose between Intentional and Negligently Accidental Actors rather than between Intentional and Truly Accidental Actors using the same sample sizes as in Experiment 1. We hypothesized that infants would not reliably distinguish Intentional from Negligently Accidental Actors of either valence. This pattern of results would suggest that, like young children, infants treat negligently accidental acts as relatively more similar to intentional acts than are truly accidental acts.

To manipulate foreseeability, Negligently Accidental Actors in Experiment 2 remained onstage during the Protagonist's attempts to either take the toy off or put the toy on the shelf, and were therefore aware of the Protagonist's goal. All other aspects of Actors' behaviors were identical to Experiment 1.




Fig. 3. Percentage of infants choosing Intentional Actors and Negligently Accidental Actors in Experiment 1 as a function of whether the actors were helping or harming the Protagonist's goal. Asterisks indicate significant differences ($p \leq 0.05$; $**p \leq 0.01$).

3. Experiment 2: Intentional versus negligently accidental help and harm

3.1. Method

3.1.1. Participants

Sixty-four full-term 10-month-olds participated (27 girls; mean age = 10 months, 0 days; range = 9;17–10;16). Subjects were randomly assigned to either the helping condition ($n = 32$; 13 girls; mean age = 10 months, 3 days; range = 9;17–10;16) or the harming condition ($n = 32$; 14 girls; mean age = 9 months, 28 days; range = 9;17–10;15).

An additional 31 babies began the study but were not included in the final sample due to failure to choose (11 infants), fussiness, (7 infants), procedural error (8 infants), parental interference (3 infants), equipment failure (1 infant), and parent-reported probable colorblindness to teal versus purple (1 infant).

Further, 28 additional babies successfully completed the study but were not included in the final sample because they were run in a significantly different environment after lab construction, which was shown to influence babies' performance across all lab studies conducted in the new space. Notably, the data patterns observed in these 28 babies were the same as the results reported below (see [Supplementary Materials](#) for details).

3.2. Procedure

Experiment 2's procedure was identical to Experiment 1, except that the Negligently Accidental Actors remained onstage throughout the Protagonist's attempts.

3.3. Results

Patterns of choice for the Intentional and Negligently Accidental Actors did not differ significantly across helping and harming conditions ($\chi^2(2, N = 64) = 1.58, p = 0.209$; see [Fig. 3](#)). Infants in the helping condition did not prefer the Intentional Helper over the Negligently Accidental Helper (17/32 chose the Intentional Helper, binomial $p = 0.860$); infants in the harming condition did not prefer the Negligently Accidental Harmer over the Intentional Harmer (20/32 chose the Negligently Accidental Harmer, binomial $p = 0.215$). There was no effect of event order, puppet color, or puppet side on any comparison ($ps > 0.600$).

As in Experiment 1, attention to individual event types did not differ within or across conditions in Experiment 2 (repeated-measures ANOVA; $ps > 0.600$; see [Table 1](#)).

3.4. Discussion

In Experiment 2, infants did not prefer an Intentional over a Negligently Accidental Helper, nor a Negligently Accidental over an Intentional Harmer. Critically, the only change from the procedure of Experiment 1 was that the Accidental Actors caused accidents in a negligent manner, rather than truly accidentally. This pattern of findings rules out the lower-level possibility that infants in Experiment 1 were simply confused by the clumsy actions of the Truly Accidental Actor, and so defaulted to preferring Intentional Helpers and avoiding Intentional Harmers as in previous work ([Hamlin & Wynn, 2011](#)); if confusion over clumsiness had driven the effects in Experiment 1 then infants should also have defaulted to preferring Intentional Helpers and avoiding Intentional Harmers in Experiment 2. Instead, the present pattern of findings suggests that infants noted the lack of intention of Truly Accidental Actors in Experiment 1, and evaluated Intentional versus Truly Accidental

Helpers and Harmers accordingly. Together, results from Experiments 1 and 2 suggest that infants can make the intention–outcome distinction, so long as accidental outcomes do not result from negligence (as in Nobes et al., 2009).

However, these results fail to elucidate exactly how infants understand negligent accidents. Although the pattern of results across Experiments 1 and 2 suggests that infants view Truly Accidental and Intentional Actors to be *more different* than are Negligently Accidental and Intentional Actors, the forced-choice quality of our dependent variables leaves the nature of this difference unclear. One possibility consistent with the results thus far is that infants treat negligently accidental actors as essentially intentional. In this case, infants might view *any* foreseeable positively or negatively valenced outcome of one's actions as a signal of an intention to cause that outcome, whether it was caused directly (in our intentional case) or as a side-effect of some other goal-directed act (in our negligent case). If this possibility were correct, infants in Experiment 2 may not have reliably distinguished Intentional and Negligently Accidental Actors because they viewed them as equally intentional and thus equally positive/negative.

A second possibility is that infants view the valenced side-effects of negligently accidental actors as being caused relatively *more* intentionally than the valenced side-effects of truly accidental actors, but still not quite as intentionally as are the outcomes of intentional actors. Indeed, infants' failure to distinguish Negligently Accidental and Intentional Actors in Experiment 2 might not mean that infants viewed the actors as entirely equivalent, but instead that the “evaluative distance” between Negligently Accidental and Intentional Actors was insufficient for infants to reliably distinguish them, given our sample size and forced-choice design.

Finally, a third possibility is that infants' evaluation of Negligently Accidental Actors differs depending on the valence of the side-effects they cause. For example, infants may view negligently accidental harming as relatively more intentional than they view negligently accidental helping, or vice versa. Consistent with the possibility that infants' assessment of negligent accidents will show a valence asymmetry, a host of work suggests that adults demonstrate a “negative agency bias,” whereby they are more likely to attribute agency to the causes of negative versus positive outcomes (e.g., Morewedge, 2009; Waytz et al., 2010; see Hamlin & Baron, 2014 for evidence with infants). Further and most relevant to the current studies, research suggests that both adults and children view foreseen (negligent) negative side-effects as more intentional than foreseen (negligent) positive side-effects; this asymmetry has been dubbed the “side-effect effect” (e.g., Knobe, 2003; Leslie et al., 2006; see Knobe et al., 2012, for review).

To explore which of these three possibilities best describes how infants evaluate negligently helpful and harmful actors, in Experiment 3 infants chose between Truly Accidental and Negligently Accidental Helpers or Harmers. Critically, in Experiment 3 both Actors caused the same positive or negative outcome for the Protagonist, and both did so when they clumsily knocked over the shelf – a side-effect of their actual goal to obtain the duck. In fact the *only* distinction between the Actors in Experiment 3 was that they were differentially aware that knocking over the shelf would be positive or negative: The Truly Accidental Actor was off-stage during the Protagonist's attempts and so was completely unaware that his goal would result in a valenced side-effect; the Negligently Accidental Actor was onstage during the attempts and so was aware that his goal would result in a valenced side-effect.

We reasoned that if infants view all negligent acts as intentional, they should distinguish the actors in both the helping and harming conditions in Experiment 3 as they did in Experiment 1, preferring the Negligently Accidental to the Truly Accidental

Helper but the Truly Accidental to the Negligently Accidental Harmer. If instead they view negligent acts as somewhat more intentional than true accidents but not as intentional as direct causes, they should not reliably distinguish the Truly Accidental from the Negligently Accidental Actor in either condition, given that we utilize the same sample sizes in Experiment 3 as in the previous Experiments. Finally, if infants' assessment of Negligently Accidental Actors depends on the valence of the outcomes they cause, infants might distinguish Actors in the harming but not the helping condition, or vice versa.

4. Experiment 3: Truly accidental versus negligently accidental help and harm

4.1. Method

4.1.1. Participants

Sixty-four full-term 10-month-olds participated (32 girls; mean age = 10 months, 0 days; range = 9;13–10;19). Subjects were randomly assigned to either the helping condition ($n = 32$; 16 girls; mean age = 10 months, 0 days; range = 9;15–10;19) or the harming condition ($n = 32$; 16 girls; mean age = 10 months, 0 days; range = 9;13–10;17).

An additional 20 babies began the study but were not included in the final sample due to failure to choose (4 infants), fussiness, (5 infants), procedural error (6 infants), parental interference (4 infants), and equipment failure (1 infant).

4.1.2. Procedure

Familiarization events for truly accidental events were identical to that of Experiment 1, and familiarization events for negligently accidental events were identical to that of Experiment 2. Except for the actors involved in familiarization and choice (i.e., a Truly Accidental Actor and a Negligently Accidental Actor, rather than an Intentional Actor and an Accidental Actor), Experiment 3's familiarization and choice were identical to Experiments 1 and 2.

After data had been collected for Experiments 1 and 2, and before we began running Experiment 3, our lab made the across-the-board decision to ask parents to keep their eyes closed during familiarization events as well as during choice procedures. For Experiments 1 and 2, and past studies in the lab, parents had only been instructed to sit quietly and to not influence their infants during familiarization; blind coders made decisions to exclude the data of babies whose parents had not followed these instructions. By having parents' eyes closed in Experiment 3, we hoped to further minimize the possibility of parental interference.

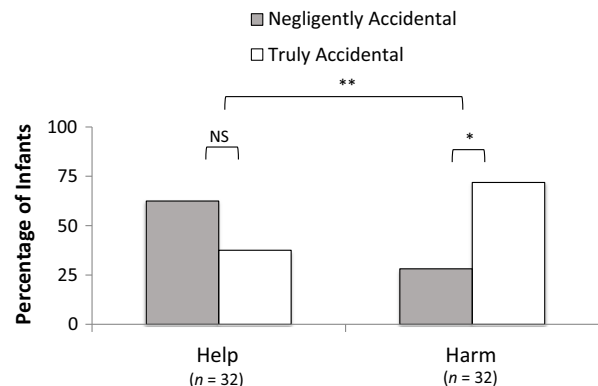


Fig. 4. Percentage of infants choosing Negligently Accidental Actors and Truly Accidental Actors in Experiment 3 as a function of whether the actors were helping or harming the Protagonist's goal. Asterisks indicate significant differences ($* p \leq 0.05$; $** p < 0.01$). NS indicates non-significant findings.

Table 2
Infants mean total looking times (s) during familiarization in Experiment 3.

	Attention to negligently accidental events	Attention to truly accidental events
Helping condition	28.65 (2.62) ^a	26.03 (2.38)
Harming condition	23.86 (2.48)	26.20 (2.34)
Helping and harming conditions	26.26 (1.82)	26.11 (1.66)

Note. All numbers in parentheses are standard errors.

^a Experiment number was significantly associated with log-transformed looking time, $F(2, 189) = 16.04$, $p < 0.001$. Fisher-Hayter's test revealed that Experiment 3's looking times were significantly higher than those of Experiments 1 and 2. Analyses were conducted to determine whether experimenter effects caused these differences in log-transformed looking time, and whether these differences influenced infants' evaluations.

Across experiments, the puppeteer-experimenter (DL, CS, or BW) was associated with looking time, $F(2, 189) = 19.14$, $p < 0.001$. Fisher-Hayter's test revealed that looking times for infants run by BW were longer than those of CS and DL. One explanation could be that the puppeteers' (DL, CS, or BW) show lengths may have differed. DL ran most infants in Experiment 1 (59/64), CS ran most infants in Experiment 2 (43/64), and BW ran all infants in Experiment 3 (64/64). We reasoned that if BW's puppet show lengths were significantly longer (or shorter) than those of DL and CS, this could account for the longer looking following the shows.

To address this possibility, experimenters blind to puppeteer coded puppet show length for forty-two infants, randomly selected across the three puppeteers. These infants included: 14 infants run by DL (7 in Experiment 1, 7 in Experiment 2); 13 infants run by CS (13 infants in Experiment 2); and 15 infants by BW (5 in Experiment 1, 5 in Experiment 2, 5 in Experiment 3). This sample represented about 20 percent of the total sample size across the three experiments, divided by the three puppeteers for the experiments that they were involved in. Considering all trial types together, the puppeteer running each trial was significantly associated with puppet show length, $F(2, 39) = 34.17$, $p < 0.001$. Fisher-Hayter's test revealed that BW's puppet show lengths ($M = 36.23$ s, $SD = 2.86$ s) were significantly longer than those of DL and CS, and that CS's puppet show lengths ($M = 30.33$ s, $SD = 3.07$ s) were significantly longer than those of DL ($M = 26.86$ s, $SD = 3.34$ s). These patterns of differences in puppet show length were similar when considering Intentional, Truly Accidental, and Negligently Accidental trial types individually (all $ps \leq 0.038$).

Critically, although puppeteer was associated with looking time and puppeteers differed in puppet show length, the puppeteer running each experiment did not appear to influence infants' evaluations. Puppet show length did not significantly predict log-transformed looking time for infants whose puppet shows were timed, $r(40) = 0.193$, $p = 0.206$. Additionally, there did not appear to be a relationship between puppeteer and rate of choosing the relatively more intentional helper or the relatively less intentional harmer across the three experimenters: Infants succeeded in distinguishing between Intentional versus Truly Accidental Actors in Experiment 1, which was run mostly by DL (who had the shortest shows), and also between Negligently versus Truly Accidental Harmers in Experiment 3, which was run entirely by BW (who had the longest shows). In addition, in no Experiments was there a difference in looking time between helping and harming conditions.

4.1.3. Counterbalancing

The following were counterbalanced across infants: (a) sub-condition (helping/harming); (b) truly accidental puppet shirt color (teal/purple); (c) truly accidental puppet order (first/s); (d) truly accidental puppet familiarization side (left/right); and (e) truly accidental puppet choice side (left/right).

4.2. Results

Patterns of choice for the Truly Accidental and Negligently Accidental Actors differed significantly across helping and harming conditions ($\chi^2(2, N = 64) = 7.63$, $p = 0.006$; see Fig. 4). Specifically, infants in the harming condition preferred the Truly Accidental Harmer over the Negligently Accidental Harmer (23/32 infants chose the Truly Accidental Harmer, binomial $p = 0.020$) whereas infants in the helping condition did not prefer the Negligently Accidental Helper over the Truly Accidental Helper (20/32 infants chose the Negligently Accidental Helper, binomial $p = 0.215$). There was no effect of event order, puppet color, or puppet side on any comparison ($p = 0.110$), except for puppet side in the helping condition (24/32 infants chose the puppet on their left during choice, binomial $p = 0.007$). Because puppet side was not observed to

influence choice in any other experiment, nor in the harming condition of Experiment 3, we consider this effect spurious and do not consider it further.

As in Experiments 1 and 2, attention to individual event types did not differ within or across conditions in Experiment 3 (repeated-measures ANOVA; $ps > 0.400$; see Table 2).

5. General discussion

In three Experiments we examined 10-month-olds' evaluations of intentional, truly accidental, and negligently accidental helpers and harmers. In Experiment 1, infants differentially evaluated Intentional versus Truly Accidental Actors in both the helping and harming conditions, preferring an agent who helped intentionally over an agent who helped truly accidentally, but preferring an agent who harmed truly accidentally over an agent who harmed intentionally. In contrast, when accidental actors had knowledge of Protagonists' goals in Experiment 2 and were therefore negligent, infants did not distinguish intentional from accidental actors either within or across helping and harming conditions, suggestive that infants differentially evaluate Truly and Negligently Accidental Actors. Experiment 3 explored this possibility specifically, and demonstrated that infants distinguished Truly Accidental and Negligently Accidental Harmers, but not Truly Accidental and Negligently Accidental Helpers, suggestive that there may be a valence asymmetry in infants' evaluation of Negligently Accidental Actors.

The present findings are compelling in several ways. First, although actions were nearly identical both across and within conditions of all Experiments, infants evaluated actors performing the very same accidental actions differentially depending on whether their actions led to positive or negative outcomes and whether they reflected true versus negligent accidents. This feature of our experimental design rules out a host of low-level alternative explanations for our results across studies. Second, whereas past studies suggestive that infants appropriately evaluate accidents have solely relied on null effects, Experiments 1 and 3 provide the first positive evidence for mentalistic accident understanding in infants' social evaluations to date. Notably, this positive evidence adds to previous work suggestive that infants understand epistemic states such as knowledge and ignorance by late in the first year of life (Hamlin et al., 2013; Meristo & Surian, 2013; see also Choi & Luo, 2015; Sloane et al., 2012). Finally, whereas most previous studies have examined the evaluation of harmful accidents alone (e.g., Armsby, 1971; Cushman et al., 2013; Farnill, 1974; Shultz et al., 1986; Yuill, 1984; cf. Leslie et al., 2006; Margoni & Surian, 2017), the current studies examine accidental help and harm. Results from Experiment 1 suggest that infants are capable of mentalistic evaluation of both accidental help and accidental harm.

Despite this, results from Experiment 3 suggest that there may be an asymmetry between how infants evaluate negligently accidental helpful and harmful acts. This pattern of results is consistent with the "side-effect effect," whereby adults and children age 4 and 5 attribute intentionality to foreseen (negligent) harm, but not to foreseen (negligent) help (Knobe, 2003; Leslie et al., 2006): Here infants reliably disliked the Negligently Accidental Harmer (versus the Truly Accidental Harmer) but did not reliably like the Negligently Accidental Helper (versus the Truly Accidental Helper). This pattern is also consistent with effects observed by Hamlin and Baron (2014), in which 6-month-old infants attributed agency to objects that previously caused negative, but not positive, outcomes for an agent with an unfulfilled goal to open a box. Of course, given the "yes/no" nature of comparisons involving forced choices and our relatively limited sample sizes, the nature of this asymmetry is somewhat unclear: Even though the tendency to prefer a negligently accidental versus a truly accidental helper did not reach significance, more babies chose a negligently accidental versus a truly

accidental helper. Thus, it is currently impossible to determine whether infants fail to attribute *any* helpful intention to negligently accidental helpers, or whether they just attribute *less* helpful intention to them, and the design of the current studies was insufficient to reveal significant evaluation differences in the help case. Of course, it is also possible that the asymmetry observed here was a statistical fluke; indeed, the effect sizes in the present research are relatively small. Additionally, the increased time of Experiment 3's puppet shows relative to those of Experiments 1 and 2 may limit the interpretation of infants' lack of preference for the Negligently Accidental versus the Truly Accidental Helper in Experiment 3, if slower puppet show times made events harder to interpret for some reason. Importantly, puppet show length did not predict individual infants' looking times, nor whether infants distinguished between puppets in a given condition (as all puppet shows in Experiment 3 were long). That said, future work should seek to replicate these results while taking care to control for puppet show length across experimenters.

Despite these open questions about infants' perceptions of negligently accidental help versus harm, the current studies provide strong evidence that infants engage in intention-based social evaluation of accidents, so long as those accidents do not result from negligence. How are we to understand these results in light of past studies suggestive that children under 4 fail the intention-outcome distinction? There are at least two non-mutually exclusive explanations. First, as discussed above, past studies have confounded accidents with negligence, and children may view negligent harms as blameworthy (Nobes et al., 2009). The current studies provide similar evidence for 10-month-olds, emphasizing the influence of negligence attributions on social evaluations early in development. Further, this evidence provides support for the possibility that infants' non-linguistic social preferences and children's verbal, sociomoral evaluations may be based on some of the same criteria.

A second explanation is that infant and preschool tasks may tap different types of knowledge (see Margoni & Surian, 2016). Specifically, studies exploring infants' mentalistic evaluations (e.g., Choi & Luo, 2015; Hamlin, 2013; Lee et al., 2015) test for implicit knowledge non-verbally. The spontaneity of these tasks may make them less susceptible to interference by salient states of the world like valenced outcomes (see, e.g., Scott, He, Baillargeon, & Cummins, 2012). By contrast, research with preschoolers mostly uses elicited-responding and tests for explicit knowledge. It seems entirely possible that preschoolers hold the explicit belief that causing negative outcomes is bad even if they were not specifically intended; indeed, it is likely that children are regularly scolded or punished for carelessly causing negative outcomes in their daily lives. If so, in order to successfully judge that children who accidentally cause large negative outcomes are less blameworthy than children who intentionally cause minor negative outcomes, then, children presumably must have sufficient inhibitory control to ignore the salient nature of large negative outcomes. Consistent with this hypothesis, increased cognitive load causes adults to judge accidental harm more harshly (Buon, Jacob, Loissel, & Dupoux, 2012), suggestive that even adults' judgments are sometimes biased by negative outcomes. To investigate this hypothesis for the development of the intention-outcome distinction, future studies should conduct spontaneous-response tasks with preschoolers, while controlling for negligence. To this end, the present studies' puppet show paradigm could be adapted to investigate the evaluation of intentional and accidental help and harm in older infants and young children. Additionally, future studies could also ask adults to evaluate the scenarios in the present studies' puppet-show paradigm, and determine: (a) whether adults view the scenarios as infants do; and (b) whether increased cognitive load influences the harshness of adults' judgments of acciden-

tal harm, and/or the positivity of adults' judgments of accidental help.

In sum, the current studies indicate that infants process helpful and harmful acts mentalistically as early as 10 months of age, adding to a growing body of research demonstrating the impressive evaluative capacities of human infants. These results challenge past work suggestive that children do not possess the intention-outcome distinction until age 4, and call for future work to: (a) explain these apparently discrepant findings; and (b) explore whether and how infants differentially perceive negligent positive versus negative acts.

Author contributions

B.M. Woo performed testing, data collection, and data analysis; created figures; interpreted the findings; and drafted the manuscript. C.M. Steckler performed testing, data collection, and preliminary analyses; and provided critical feedback to the manuscript. D. T. Le contributed to the study's design and concept, performed data collection, and provided critical feedback to the manuscript. J.K. Hamlin contributed to the study's design and concept, oversaw the data collection, provided oversight over data analyses, contributed to interpreting the findings, and provided critical feedback and revisions to the manuscript. Data collection was conducted in J.K. Hamlin's laboratory.

Acknowledgements

This work was supported by the National Sciences and Engineering Research Council of Canada [Grant number 402356-11].

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cognition.2017.06.029>.

References

- Armsby, R. E. (1971). A reexamination of the development of moral judgments in children. *Child Development*, 42(4), 1241. <http://dx.doi.org/10.2307/1127807>.
- Baird, J. A., & Astington, J. W. (2004). The role of mental state understanding in the development of moral cognition and moral action. *New Directions for Child and Adolescent Development*, 2004(103), 37–49. <http://dx.doi.org/10.1002/cd.96>.
- Behne, T., Carpenter, M., & Tomasello, M. (2005). One-year-olds comprehend the communicative intentions behind gestures in a hiding game. *Developmental Science*, 8(6), 492–499.
- Buon, M., Jacob, P., Loissel, E., & Dupoux, E. (2012). A non-mentalistic cause-based heuristic in human social evaluations. *Cognition*, 126(2), 149–155.
- Buon, M., Jacob, P., Margules, S., Brunet, I., Dutat, M., Cabrol, D., & Dupoux, E. (2014). Friend or foe? Early social evaluation of human interactions. *PLoS ONE*, 9(2), e88612. <http://dx.doi.org/10.1371/journal.pone.0088612>.
- Carpenter, M., Akhtar, N., & Tomasello, M. (1998). Fourteen- through 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior and Development*, 21(2), 315–330. [http://dx.doi.org/10.1016/S0163-6383\(98\)90009-1](http://dx.doi.org/10.1016/S0163-6383(98)90009-1).
- Casstevens, R. M. (2007). *jHab: Java habituation software (version 1.0.2)* [computer software]. Chevy Chase, MD.
- Chandler, M. J., Sokol, B. W., & Hallett, D. (2001). Moral responsibility and the interpretive turn: Children's changing conceptions of truth and rightness. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition* (pp. 345–366). Cambridge, MA: MIT Press.
- Choi, Y., & Luo, Y. (2015). 13-month-olds' understanding of social interactions. *Psychological Science*, 26(3), 274–283. <http://dx.doi.org/10.1177/0956797614562452>.
- Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, 108(2), 353–380. <http://dx.doi.org/10.1016/j.cognition.2008.03.006>.
- Cushman, F., Sheketoff, R., Wharton, S., & Carey, S. (2013). The development of intent-based moral judgment. *Cognition*, 127(1), 6.
- Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science*, 17(12), 1082–1089. <http://dx.doi.org/10.1111/j.1467-9280.2006.01834.x>.

- Farnill, D. (1974). The effects of social-judgment set on children's use of intent information. *Journal of Personality*, 42(2), 276.
- Fawcett, C., & Liszkowski, U. (2012). Infants anticipate others' social preferences. *Infant and Child Development*, 21(3), 239–249. <http://dx.doi.org/10.1002/icd.739>.
- Gruenich, R. (1982). The development of children's integration rules for making moral judgments. *Child Development*, 53(4), 887–894. <http://dx.doi.org/10.1111/j.1467-8624.1982.tb01350.x>.
- Hamlin, J. K. (2013). Failed attempts to help and harm: Intention versus outcome in preverbal infants' social evaluations. *Cognition*, 128(3), 451–474. <http://dx.doi.org/10.1016/j.cognition.2013.04.004>.
- Hamlin, J. K. (2015). The case for social evaluation in preverbal infants: Gazing toward one's goal drives infants' preferences for Helpers over Hinderers in the hill paradigm. *Frontiers in Psychology*, 5, 1563. <http://dx.doi.org/10.3389/fpsyg.2014.01563>.
- Hamlin, J. K., & Baron, A. S. (2014). Agency attribution in infancy: Evidence for a negativity bias. *PLoS ONE*, 9(5), e96112.
- 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. <http://dx.doi.org/10.1111/desc.12017>.
- Hamlin, J. K., & Wynn, K. (2011). Young infants prefer prosocial to antisocial others. *Cognitive Development*, 26(1), 30–39. <http://dx.doi.org/10.1016/j.cogdev.2010.09.001>.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, 450(7169), 557–559. <http://dx.doi.org/10.1038/nature06288>.
- Killen, M., Lynn Mulvey, K., Richardson, C., Jampol, N., & Woodward, A. (2011). The accidental transgressor: Morally-relevant theory of mind. *Cognition*, 119(2), 197–215. <http://dx.doi.org/10.1016/j.cognition.2011.01.006>.
- Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis*, 63(279), 190–194.
- Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., & Sommers, T. (2012). Experimental philosophy. *Annual Review of Psychology*, 63, 81–99.
- Kohlberg, L. (1969). Stage and sequence: The cognitive-developmental approach to socialization. In D. A. Goslin (Ed.), *Handbook of socialization theory and research* (pp. 347–480). Chicago: Rand McNally.
- Kuhlmeier, V., Wynn, K., & Bloom, P. (2003). Attribution of dispositional states by 12-month-olds. *Psychological Science*, 14(5), 402–408. <http://dx.doi.org/10.1111/1467-9280.01454>.
- Lee, Y., Yun, J. E., Kim, E. Y., & Song, H. (2015). The development of infants' sensitivity to behavioral intentions when inferring others' social preferences. *PLoS ONE*, 10(9), e0135588. <http://dx.doi.org/10.1371/journal.pone.0135588>.
- 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.
- Margoni, F., & Surian, L. (2016). Explaining the U-shaped development of intent-based moral judgments. *Frontiers in Psychology*, 7, 219. <http://dx.doi.org/10.3389/fpsyg.2016.00219>.
- Margoni, F., & Surian, L. (2017). Children's intention-based moral judgments of helping agents. *Cognitive Development*, 41, 46–64.
- Meristo, M., & Surian, L. (2013). Do infants detect indirect reciprocity? *Cognition*, 129(1), 102–113. <http://dx.doi.org/10.1016/j.cognition.2013.06.006>.
- Monroe, A. E., & Malle, B. F. (2017). Two paths to blame: Intentionality directs moral information processing along two distinct tracks. *Journal of Experimental Psychology: General*, 146(1), 123.
- Moran, J. M., Young, L. L., Saxe, R., Lee, S. M., O'Young, D., Mavros, P. L., & Gabrieli, J. D. (2011). Impaired theory of mind for moral judgment in high-functioning autism. *Proceedings of the National Academy of Sciences of the United States of America*, 108(7), 2688–2692. <http://dx.doi.org/10.1073/pnas.1011734108>.
- Morewedge, C. K. (2009). Negativity bias in attribution of external agency. *Journal of Experimental Psychology: General*, 138(4), 535.
- Nobes, G., Panagiotaki, G., & Pawson, C. (2009). The influence of negligence, intention, and outcome on children's moral judgments. *Journal of Experimental Child Psychology*, 104(4), 382–397. <http://dx.doi.org/10.1016/j.jecp.2009.08.001>.
- Núñez, N., Laurent, S., & Gray, J. M. (2014). Is negligence a first cousin to intentionality? Lay conceptions of negligence and its relationship to intentionality. *Applied Cognitive Psychology*, 28(1), 55–65. <http://dx.doi.org/10.1002/acp.2957>.
- Ohtsubo, Y. (2007). Perceived intentionality intensifies blameworthiness of negative behaviors: Blame-praise asymmetry in intensification effect. *Japanese Psychological Research*, 49(2), 100–110. <http://dx.doi.org/10.1111/j.1468-5884.2007.00337.x>.
- Piaget, J. (1965). *The moral judgment of the child*. New York: Free Press (Original work published 1932).
- Sakkalou, E., & Gattis, M. (2012). Infants infer intentions from prosody. *Cognitive Development*, 27(1), 1–16.
- Salvadori, E., Blazsekova, T., Volein, A., Karap, Z., Tatone, D., Mascaro, O., & Csibra, G. (2015). Probing the strength of infants' preference for helpers over hinderers: Two replication attempts of Hamlin and Wynn (2011). *PLoS ONE*, 10(11), e0140570. <http://dx.doi.org/10.1371/journal.pone.0140570>.
- Scola, C., Holvoet, C., Arciszewski, T., & Picard, D. (2015). Further evidence for infants' preference for prosocial over antisocial behaviors. *Infancy*, 20(6), 684–692. <http://dx.doi.org/10.1111/infa.12095>.
- Scott, R. M., He, Z., Baillargeon, R., & Cummins, D. (2012). False-belief understanding in 2.5-year-olds: Evidence from two novel verbal spontaneous-response tasks: Toddlers pass verbal spontaneous-response false-belief tasks. *Developmental Science*, 15(2), 181–193. <http://dx.doi.org/10.1111/j.1467-7687.2011.01103.x>.
- Shultz, T. R., Wright, K., & Schleifer, M. (1986). Assignment of moral responsibility and punishment. *Child Development*, 57(1), 177. <http://dx.doi.org/10.2307/1130649>.
- Sloane, S., Baillargeon, R., & Premack, D. (2012). Do infants have a sense of fairness? *Psychological Science*, 23(2), 196–204. <http://dx.doi.org/10.1177/0956797611422072>.
- Steckler, C. M., Woo, B. M., & Hamlin, J. K. (2017). The limits of early social evaluation: 9-month-olds fail to generate social evaluations of individuals who behave inconsistently. *Cognition*, 167, 255–265. <http://dx.doi.org/10.1016/j.cognition.2017.03.018>.
- Tasimi, A., & Wynn, K. (2016). Costly rejection of wrongdoers by infants and children. *Cognition*, 151, 76–79. <http://dx.doi.org/10.1016/j.cognition.2016.03.004>.
- Vaish, A., Carpenter, M., & Tomasello, M. (2009). Sympathy through affective perspective taking and its relation to prosocial behavior in toddlers. *Developmental Psychology*, 45(2), 534–543. <http://dx.doi.org/10.1037/a0014322>.
- Waytz, A., Morewedge, C. K., Epley, N., Monteleone, G., Gao, J. H., & Cacioppo, J. T. (2010). Making sense by making sentient: Effort motivation increases anthropomorphism. *Journal of Personality and Social Psychology*, 99(3), 410.
- Young, L., Cushman, F., Hauser, M., & Saxe, R. (2007). The neural basis of the interaction between theory of mind and moral judgment. *Proceedings of the National Academy of Sciences of the United States of America*, 104(20), 8235–8240. <http://dx.doi.org/10.1073/pnas.0701408104>.
- Young, L., & Saxe, R. (2009). An fMRI investigation of spontaneous mental state inference for moral judgment. *Journal of Cognitive Neuroscience*, 21(7), 1396–1405. <http://dx.doi.org/10.1162/jocn.2009.21137>.
- Yuill, N. (1984). Young children's coordination of motive and outcome in judgments of satisfaction and morality. *British Journal of Developmental Psychology*, 2(1), 73–81.
- Yuill, N., & Perner, J. (1988). Intentionality and knowledge in children's judgments of actor's responsibility and recipient's emotional reaction. *Developmental Psychology*, 24(3), 358–365. <http://dx.doi.org/10.1037/0012-1649.24.3.358>.
- Zelazo, P. D., Helwig, C. C., & Lau, A. (1996). Intention, act, and outcome in behavioral prediction and moral judgment. *Child Development*, 67(5), 2478–2492. <http://dx.doi.org/10.2307/1131635>.