Third-party punishment promotes fairness in children

Justin W. Martin, Sophia Martin, Katherine McAuliffe

Department of Psychology, Boston College, Chestnut Hill, MA

Word count: 8,492

Correspondence: justinwmartin@gmail.com (J.W. Martin)

Address: 275 Beacon St, Chestnut Hill, MA 02467

Author note: We are grateful to Richard Ahl, Hannah Bolotin, Jonathan Schulz and Gordon Kraft-Todd for feedback on earlier versions of this manuscript. This manuscript was supported by funding from the John Templeton Foundation (KM). Data and analysis scripts available at <https://osf.io/k4q8c/>.

Abstract

Third-party punishment can promote fair behavior. However, the mechanisms by which this happens are unclear. Third-party punishment may increase fair behavior by providing *direct feedback*, helping shape the behavior of those punished, or through an influence on *reputation*, by encouraging the transgressor to behave appropriately before a third-party audience. Investigating whether and how third-party punishment leads to fairness in children during middle childhood— a key developmental period in the emergence of fairness—presents an ideal context in which to explore these two mechanisms. Six to nine-year-old children ($N = 121$) allocated resources between themselves and a partner in a forced-choice dictator game. In the Direct Feedback condition, a third party punished if the child chose the less fair option. In the Reputation condition, a third party merely observed the child. In the Baseline condition, no third party was present. We find that the Direct Feedback condition increased fairness relative to the other conditions, especially amongst younger children. These results highlight an important link between third-party punishment and fairness in children and, more broadly, help clarify the mechanisms through which third-party punishment can influence fair behavior and thereby, human cooperation.

Keywords: punishment; fairness; cooperation; reputation

1 Introduction

As a testament to the importance placed on fairness, adults pay to punish when fairness norms are violated (Fehr & Fischbacher, 2003; Fehr & Gächter, 2002; FeldmanHall et al., 2014; Henrich, 2006; Henrich et al., 2010). Such punishment motivates fair behavior: contributions to public goods increase when punishment is possible (Fehr & Gächter, 2002; Gächter et al., 2008). However, *how* punishment motivates fair behavior—that is, by which mechanisms punishment leads to fair behavior—is not well understood.

In line with past work (Fehr & Fischbacher, 2004; FeldmanHall et al., 2014; Martin et al., 2019; Raihani & Bshary, 2015), we distinguish between second-party punishment (punishing as a harmed victim) and third-party punishment (punishing as an unharmed observer). Third-party punishment may be particularly informative in shedding light on the mechanisms by which punishment can encourage fair behavior; while both second- and third-party punishment emerge early in development (Bernhard et al., 2020; Lergetporer et al., 2014; McAuliffe et al., 2015) and are found cross-culturally (Henrich et al., 2006, 2010), third-party punishment may be a particularly strong indicator of a commitment to fairness. That is, a victim (i.e. a second party) may punish to prevent you from harming them again (Clutton-Brock & Parker, 1995; Crockett et al., 2014; Trivers, 1971) or to level your payoff relative to theirs (Bone & Raihani, 2015; Raihani & Bshary, 2019), in addition to their concern for fairness. In contrast, only those who care strongly about fairness will pay to punish on behalf of someone else, exposing themselves to potential retaliation in the process (Cinyabuguma et al.,

2006; Herrmann et al., 2008). Thus, being subject to third-party punishment may inform individuals about the importance placed on fair behavior by the one punishing (and those around them) (Barclay, 2006; Fehr & Gächter, 2002; Jordan et al., 2016; Jordan & Rand, 2020).

Limited work has investigated the influence of third-party punishment on fairness. In one study (Charness et al., 2008), participants played a Trust Game and, in some cases, were observed by a third party with the ability to punish (and reward). Both the initial investment from the investor and the amount returned by the trustee were substantially higher when third-party punishment was possible (Charness et al., 2008). While investments in the Trust Game are motivated by considerations beyond just fairness, these results provide important evidence that third-party punishment can lead individuals to engage in more fair behavior. However, because decisions about behavior in the trust game principally involve considerations about trust and trustworthiness, these results tell us little about the mechanisms supporting the relationship between third-party punishment and fair behavior. Our goal was to move beyond these results by examining the mechanisms by which third-party punishment encourages fair behavior, using child development as a critical window into these processes.

1.1 Investigating the influence of third-party punishment as fairness concerns emerge in childhood

Probing the mechanisms by which third-party punishment encourages fair behavior may be most effectively done as concerns for fairness are emerging and are in direct tension with selfishness—the desire to maximize personal rewards. During this period—a period during which an understanding of local norms is actively forming—children may be particularly open to feedback about what kinds of behaviors are appropriate while nevertheless continuing to be strongly influenced by personal benefits. Thus, childhood—and especially the ages at which fairness concerns begin to gain prominence—represents an ideal context in which to test the ways in third-party punishment motivates fair behavior. Before detailing our approach to exploring how third-party punishment motivates fair behavior, we first review how concern for fairness emerges in development.

Fairness concerns first emerge surprisingly early in development, as early as the first three or four years of life, yet do not reach adult-like levels until around 8 years of age (Blake et al., 2015; House et al., 2013; McAuliffe, Blake, et al., 2017; Warneken, 2017). In windfall situations, cases in which an endowment of resources is provided rather than earned, fairness considerations appear in children in the preschool years, manifesting as a relatively egoistic focus on fairness: young children tend to reject only unfair allocations that disadvantage themselves (i.e. disadvantageous inequity aversion) (Blake et al., 2015; McAuliffe, Blake, et al., 2017; Shaw & Olson, 2012). Later in childhood, children transition to a relatively broader and less egoistic—a more impartial—concept of fairness, in which they begin to additionally reject unfair allocations that benefit themselves but disadvantage others (i.e. advantageous inequity

aversion) (Blake et al., 2015; McAuliffe, Blake, et al., 2017). Thus, middle childhood—between 6 and 8 years of age—marks a key transition from a more self-focused form of fairness to a more impartial form of fairness.

Middle childhood represents an ideal time to probe the mechanisms by which third-party punishment motivates fair behavior. That is, as children begin to realize that fairness applies not just to themselves but to others as well, third-party punishment may serve a particularly useful role in providing the necessary feedback to facilitate that transition. This developmental period thus offers the critical window in which to investigate whether and to what extent third-party punishment promotes fair behavior in children and in which to explore the mechanisms supporting this relationship.

Critically, while much work has demonstrated the potential for punishment in promoting behaviors related to fairness in adulthood, especially cooperation (Ambrus & Greiner, 2012; Cinyabuguma et al., 2006; Fehr & Gächter, 2000, 2002; Ones & Putterman, 2007; Pfattheicher et al., 2018; Rockenbach & Milinski, 2006) and some work has explored how punishment increases cooperation in childhood (Lergetporer et al., 2014), no work to our knowledge has demonstrated that third-party punishment promotes fairness in children. Thus, our first goal is to provide, to our knowledge, the first investigation of *whether* third-party punishment can promote fair behavior, resolving this important open question. Next, we ask *how* third-party punishment might accomplish this increased adherence to fairness norms. One possibility is that third-party punishment promotes fairness through the provision of *direct feedback*. This is the role for punishment that is typically assumed. However, another possibility that has

been largely neglected is that third-party punishment operates through an influence on *reputation*.

1.2 The role of third-party punishment through its provision of direct feedback

Punishment might increase adherence to fairness norms through its ability to provide direct feedback. Past work has described punishment as a form of teaching (Caro & Hauser, 2004; Martin & Cushman, 2016; Raihani et al., 2012). That is, by punishing one can send a signal to the one punished that their behavior was not acceptable and should change in the future. Consistent with this, parents typically punish children for particular behaviors and stop punishing after those behaviors are corrected (Davidov et al., 2012; Holden et al., 1995). And, children expect this kind of punishment to work: children between 5 and 8 years of age expect those who are punished for a behavior to not perform that behavior in the future (Bregant et al., 2016). As discussed above, third-party punishment may provide particularly useful feedback about the importance of norms, as third-party punishment is likely to be relatively less selfish than second-party punishment. That is, punishment from a second party (i.e. a victim) may occur because fairness is a strong norm in the present social context, or they may be punishing because they want to prevent you from harming them again in the future (Crockett et al., 2014), or they want to level your payoff relative to theirs (Bone & Raihani, 2015). However, if you are unfair to someone and a third party punishes you, this observer is much less likely to be motivated by personal benefits (there is no harm to them to prevent occurring again, their payoff was not impacted by

your behavior, etc.) and instead is more likely to be motivated by a fairness norm.

Thus, for the one being punished, third-party punishment sends a strong signal that fairness norms are important to adhere to. This role for third-party punishment through direct feedback has not been previously tested, however, nor has it been distinguished from the potential role that third-party punishment may play through an influence on reputation.

1.3 The role of third-party punishment through its influence on reputation

As an alternative to a role for third-party punishment in promoting fair behavior by providing direct feedback, third-party punishment may promote fair behavior through reputational concerns. By definition, third-party punishment involves an observer witnessing a transgression, and it is possible that this observation—in the absence of the enactment of punishment—is sufficient to promote norm adherence among those who would otherwise behave unfairly. This is because one's reputation can have important consequences: those who are perceived as fair (Przepiorka & Liebe, 2016; Raihani & Barclay, 2016) or generous (Barclay & Willer, 2007; Bornhorst et al., 2010; Coricelli et al., 2004; Fehrler & Przepiorka, 2016; Hardy & Van Vugt, 2006; Lammers et al., 2017; Page et al., 2005; Pleasant & Barclay, 2018; Sylwester & Roberts, 2010) get picked for cooperative interactions, and individuals will engage in costly behaviors in order to gain a better reputation (Jordan et al., 2016). This perspective would suggest that simply being observed by a third party is sufficient to promote fair behavior. Indeed, a recent account of the development of fairness concerns highlights a role for

“publicity” in how much children care about fair behavior (Engelmann & Tomasello, 2019). Consistent with this perspective, sensitivity to observation appears to emerge early in life (Engelmann & Rapp, 2018; Warneken, 2017) and becomes more pronounced with age. For instance, 5-year-old children: steal more and help less when a peer is not present to observe (Engelmann et al., 2012); share more with absent individuals when observed by a peer (Engelmann et al., 2013); and are most generous when their behavior is observable by their peers (Leimgruber et al., 2012). This sensitivity to being observed is not fully developed by 5, however; recent work has found that while 8-year-olds engage in costly sharing selectively more when a third party will know about their sharing and will be selecting a partner for a subsequent game, 5-year-olds do not modulate their sharing in this context (Herrmann, Engelmann, & Tomasello, 2019). These results suggest that children begin taking reputational concerns into account by 5 years of age, and this concern becomes more sophisticated over development (Engelmann & Rapp, 2018; Warneken, 2017). That is, between the ages of 6 and 8, children are within a transitional period, shifting from a relatively simple focus on being observed to a more complex, fully-fledged concern about being observed. According to this account, we would expect children as young as 5 or 6 to begin behaving more fairly in third-party contexts due to the mere fact of being observed and that this sensitivity to observation would increase with age.

1.4 Dissociating the influence of direct feedback and reputation in childhood

To explore whether third-party punishment promotes fair behavior through the provision of direct feedback or an influence on reputation, we designed a task in which participants ($N = 121$) play as the dictator in a forced-choice dictator game with an absent peer as recipient. Participants complete 6 trials, choosing between a less fair and more fair option (e.g., between 3 resources for both the participant and the peer, or 4 resources for the participant and 2 for the peer; Figure 1), with the options available varying across trials. In the *Direct Feedback* condition, an additional experimenter is present and engages in costly punishment if the participant does not choose the fairer allocation (this punishment only impacts the participant and not the peer). In the *Reputation* condition, an additional experimenter is present but no punishment is possible (i.e. they are a mere observer), although they use a similar decision-making mechanism as in the Direct Feedback condition (e.g. put tokens into a box following the participant's decision) (Figure 1). We note that reputational concerns are relatively stronger in this condition but they may be present across the other conditions as well. The Reputation condition includes an increased number of individuals who perceive the participant's behavior as well as an observer whose sole task is to watch the participant's choices, thereby posing a greater reputational risk, but there is always at least one experimenter in each condition who perceives the participant's behavior. We also include a *Baseline* condition in which no additional experimenter is present. We analyze the proportion of trials in which the participant chooses the fairer option.

We have three main predictions. First, we predict that third-party punishment will operate through the direct feedback it provides, and so only the Direct Feedback

condition will lead to fairer allocations. In contrast, if the mechanism by which third-party punishment promotes adherence to fairness norms is through reputational concerns, then both those who are observed without the possibility of punishment (children in the Reputation condition) and those who can be punished by a third party (children in the Direct Feedback condition) should enact fairer allocations than those for whom there is no third party present (children in the Baseline condition). Second, we predict that this increase will occur with repeated trials, as participants are more likely to have chosen a less fair allocation and have been punished as the experiment progresses. Finally, we predict the differential influence of condition to be most prominent in younger children. Third-party punishment may be most effective for children between the ages of 6 and 8, encouraging them to adhere strictly to the concept of fairness and choose allocations that are fair even when it costs them to do so. However, by 8 years of age, children will already behave more fairly, and so punishment may have less of an effect.

Figure 1. Schematic of the setup of the Baseline, Direct Feedback, and Reputation conditions. In all conditions, the participant made a choice between two options of how to divide resources between themselves (circled resources) and an absent peer (a forced-choice Dictator Game). In the Reputation condition, after the participant made their choice, an observer (secondary experimenter) randomly placed a token into the red or green hole of a box to enact the participant's decision. In the Direct Feedback condition, after the participant made their choice, the secondary experimenter could either punish, by putting a resource into the red box, in which case the participant would not receive their resources (while the recipient would receive his/hers), or they could accept the decision, by putting a resource into the green box, in which case the participant and recipient would receive their resources.

2. Methods

2.1 Participants

Participants ($N = 132$) were recruited into groups of 6- to 7-year-olds and 8- to 9-year-olds. Participants were tested in a laboratory space ($n = 27$), public parks ($n = 87$), museums ($n = 12$) or a zoo ($n = 6$). Eleven participants were excluded because either they were outside of age range (1), did not like Skittles (3), were not typically developing (1), due to experimenter error (3) or due to interference or previous knowledge of study (3). The final sample included 121 participants (6- to 7-year-olds: $n = 60$, Mean (months) $\pm$ Standard deviation = 84.0 ± 7.1 ; 8-9-year-olds: $n = 61$, 106.1 ± 7.2 ; 60 female, 61 male; Table S1). All methods and procedures were approved by the Boston College Institutional Review Board (protocol number 16.242.01, titled “Cooperation in children”).

Sample sizes were selected to reflect the sample sizes used in similar prior research (McAuliffe et al., 2014, 2015; McAuliffe, Raihani, et al., 2017; Yang et al., 2018). Sample sizes were established *a priori* and data collection was stopped after reaching that point. We conducted a sensitivity analysis to determine the effect sizes our study would have 80% power to detect, given our design and sample size. This analysis focused on the interaction between condition and age. Using the package *simr* (Green & Macleod, 2016) in R, we ran a series of 1000 sample bootstrap tests comparing models including the 2-way interaction between condition and age and models without this interaction. We imputed an interaction effect size between the

range [3 -3], corresponding to an odds ratio range of [20.1 0.05]. We found that our model had ~80% power to detect effects greater than 1.05 (OR = 2.86) and smaller than -0.9 (OR = 0.41), which corresponds to a small to medium effect size¹.

2.2 Procedure

Participants were pseudo-randomly assigned (based upon the availability of second experimenters) to one of three conditions (Baseline, Direct Feedback, Reputation). Across conditions, participants first completed a practice trial of the dictator game (see Supplemental Information for the full experimental protocol for each condition). They were told that tokens could be exchanged for sweets (Skittles) at the end of the game². Participants were given a token bag, and a token bag for the absent recipient was presented. No deception was used: participants were truly making decisions that impacted an absent peer (see Supplemental Material). To practice the role of dictator, participants were allotted three tokens and were told that they could keep as many as they wanted for themselves and give as many as they wanted to their partner. The experimenter presented laminated cards and explained that the side of the card with a dotted circle on it (closer to the participant) contained the participant's tokens, and the blank side of the card (further from the participant) contained the partner's tokens (Figure 1). Two comprehension check questions about the cards were then asked (performance on each comprehension check question was >90% across all

¹ Effect size comparison taken from <http://imaging.mrc-cbu.cam.ac.uk/statswiki/FAQ/effectSize>.

² The exchange rate between tokens and Skittles was intentionally left vague, so as to facilitate payoffs at the end of the game. One participant repeatedly asked the experimenter exactly how many Skittles equated to one token, and was told that they would be traded in at a 1:1 ratio.

participants; Table S2). The experimenter then introduced new forced-choice cards, presenting two different distributions of tokens. The dictator was shown the two distributions of three tokens—a 2/1 and a 1/2 division (side counterbalanced across participants)—and asked how they wanted to divide the three tokens. This decision was then enacted.

Participants in the *Baseline* condition next proceeded directly to the dictator game. For participants in the *Direct Feedback* and *Reputation* conditions, we next explained the roles of the punisher and observer, respectively. The instructions at this stage were directed toward the punisher or observer, but participants were told that it was very important for them to listen to these instructions also. In these conditions, the roles of punisher and observer were played by adult female research assistants.

In the *Direct Feedback* condition, the punisher was endowed with a handful of tokens, which were placed into a green box with a large hole and a small hole (Figure 1). The punisher could exchange any tokens remaining in the green box for Skittles at the end of the game. However, any tokens placed into the red box could not be traded in for Skittles at the end of the game. Each time after the participant made a decision, the punisher could accept the dictator's decision by placing a token into the green box, in which case the participant and the partner received their tokens. Or, the punisher could punish the dictator by placing a token into the red box, in which case the participant's tokens were removed, but the partner's tokens were put into the partner's bag. Thus, punishment was costly and only impacted the participant (not the recipient). This was demonstrated for participants: for practice, a 1/1 distribution was selected by

the experimenter, and the punisher was told by the experimenter to put a token into the green box. This was then enacted—the tokens were put into the participant's and absent peer's bags. Then, another 1/1 distribution was selected by the experimenter, and the punisher was told by the experimenter to put a token into the red box. The distribution was then enacted—the peer's token was put into their bag, while the participant's token was removed.

In the *Reputation* condition, the observer's behavior had no impact on how the tokens were allocated. The observer's endowment of a handful of tokens was placed into a half green and half red box with one large hole in the middle and a small hole on either end (Figure 1). After the participant made his/her decision, the observer put one token into either the red or green hole of the box, at which point the tokens were distributed as chosen by the participant. As in the Direct Feedback condition, this was demonstrated twice for the participant, both times with a 1/1 distribution, once with the observer putting a token into the green hole and once with the observer putting a token into the red hole. During the main trials, the observer used the red and green holes 3 times each, in a pre-randomized order³.

Finally, in the *Baseline* condition, no additional experimenter was present.

In all three conditions, a series of eight comprehension check questions about this game were then asked. Finally, the punisher or observer (dependent upon condition) and participant were told to not speak to each other throughout the game.

³ The decision to pre-randomize the color order was not made until after five participants had completed the experiment in the Reputation condition. Two of these participants saw an uneven amount of tokens placed into the red and green holes.

2.3 Dictator game

The participant then began six forced-choice dictator game trials, in which he/she was asked to make a choice between two different distributions of 6 tokens between themselves and an absent peer. Each of the following distributions was used twice, in the order presented here (the amount for the participant is listed first and the amount for the recipient is listed second): 3/3 vs. 4/2 (trials 1 and 2), 3/3 vs. 5/1 (trials 3 and 4), and 4/2 vs. 5/1 (trials 5 and 6)⁴. Which option was presented first (e.g. 3/3 or 4/2) was counterbalanced. This trial ordering was picked so that the initial trials had a clear fair choice—pitting equality against inequality (3/3 vs. 4/2 and 3/3 vs. 5/1)—while the later trials were slightly more ambiguous—pitting a more selfish option against a less selfish option (4/2 vs. 5/1). This ordering resulted in an asymmetry in the relative benefits across trials, although we find that trial type is unlikely to account for our results (see Results section 3.3). After the participant chose a distribution, the punisher/observer made their decision and the tokens were distributed. After the final trial, participants were asked a series of three to six debrief questions (depending on the condition they were in and the distributions they chose)⁵. Finally, participants were asked whether or not they thought their partner was “real” or “pretend.” Overall, 47.8% of participants answered “real”, which we expect is due to the fact that they either misunderstood pretend as meaning ‘not present’ and/or because they were confused that the

⁴ Due to experimenter error, one participant saw seven trials, in which the fourth trial was repeated before moving on to the fifth trial; no trials were omitted or replaced.

⁵ One participant requested to terminate the session after completing the experimental phases, before the debrief questions. This participant was included in the sample, as the experimental manipulation was complete.

experimenter who introduced the absent peers was now questioning their existence. Importantly, children overwhelmingly answered all other questions about the absent children correctly (performance >92% on all questions; Table S2) and our results are robust to the inclusion of a belief term in our models (Table S5). At the end of the study, children exchanged their tokens for Skittles and the session concluded.

2.4 Data coding and analysis approach

A large majority (91%) of the sessions were videotaped. Children's decisions and debrief responses were live coded on paper by the primary experimenter and later recoded by an independent paper coder as well as an independent video coder. Disagreements between paper and video coding were resolved by re-watching the video (when possible) and looking at the paper data for a second time. The paper and video coder disagreed on whether or not the participant made the more selfish or less selfish decision ~1% of the time. Disagreements were resolved with reference to the video. Debrief responses were coded as belonging to one of 5 categories (non-mutually exclusive) by 2 independent coders. These coders disagreed 2.5% of the time. Coders agreed >90% of the time for each category within each question. Disagreements were resolved by consensus. At the overall level, participants were found to reference 2 categories substantially more than the others (Fairness: 29%; Personal benefit: 26%; All others: <9%). Thus, our use of these responses focuses on these two categories. We also focus on debrief responses that were asked to all participants only (Table S3).

All statistical analyses were conducted using R statistical software (version 3.5.0). Decisions in the dictator game were analyzed using mixed-effects logistic

regression, given the binary response (less selfish = 1, more selfish = 0). Models were run using the package 'lme4' (Bates et al., 2014). The fixed effects of primary interest include condition (3 level factor; Baseline, Reputation or Direct Feedback; Baseline condition was the default reference level, but Reputation was used as the reference level where noted) and participant age (in months, z-scored; the statistical pattern of results was unchanged when age was included as a binary factor, Table S4).

Participant gender was included as an additional fixed effect predictor in all models (the statistical pattern of results is unchanged if this predictor is dropped). Some models additionally included trial number (z-scored) and trial distribution (3 level factor: 3/3 vs. 4/2, 3/3 vs. 5/1, 4/2 vs. 5/1; 3/3 vs. 4/2 was the default reference level) as fixed effect predictors, as noted. All models included a random intercept for participant, except when data from the first trial only was analyzed. For all models, we ran an additional version in which we include self-reported belief in the real vs. pretend nature of dictator game partners as an additional fixed effect predictor of no interest. Results were statistically identical when this predictor was included (Table S5).

In testing the significance of these fixed effects, we take a model comparison approach. First, we compare a null model (only including random effects) to our “full” model that includes all predictors and relevant interactions. This full model is preferred ($\text{LRT } \chi^2(12) = 76.86, p < 0.001$). Next, we compare this model to one without the highest-level interaction term (e.g. the 3-way interaction in a model with 3 predictors). If the model including the interaction provides a better fit to the data, we then investigate this model further. If not, we then compare the reduced model to one that removes the

next highest-level interaction term or main effects terms (if no interaction terms remain).

Predicted effects are plotted in figures, with binomial confidence intervals calculated using the Agresti-Coull method (Agresti & Coull, 1998).

Data and analysis scripts available through the Open Science Framework (<https://osf.io/k4q8c/>; Martin et al., 2021).

2.5 Dictator game recipients

Our experiment did not involve deception. As such, an additional male and female participant were subsequently recruited to serve as recipients in the dictator game. The primary experimenter explained this game to the male and female recipient and then these two individuals were told the exact number of Skittles that each participant in each condition chose to give them; the female was sent home with a total of all of the Skittles that the female participants had allocated to her, and the male was sent home with a total of all of the Skittles that the male participants had allocated to him.

3. Results

3.1 *Third-party punishment leads to more fair behavior selectively in younger children*

We find that the Direct Feedback condition leads to more fair behavior, but only among younger children. In particular, we find a significant interaction between condition and age on participants' decisions (model 1; LRT $\chi^2(2) = 8.97$, $p = 0.01$; Table 1). This interaction was driven principally by a significantly stronger influence of age on behavior in the Direct Feedback condition relative to that of age on behavior in the Reputation condition (OR = 0.40, SE = 0.37, $z = -2.46$, $p = 0.01$) or in the Baseline condition (OR = 0.35, SE = 0.38, $z = -2.78$, $p = 0.005$). Though we include age as a continuous variable in our analyses (see Methods), we further explore this interaction by dividing our sample into our pre-determined age groups: younger (6&7) and older (8&9). Looking first at the younger age group, we find a large difference between the Direct Feedback condition (Figure 2; $M = 0.84$, $SEM = 0.05$, 95% CI = 0.71 – 0.92) and the Reputation ($M = 0.38$, $SEM = 0.08$, 95% CI = 0.23 – 0.55) and Baseline conditions ($M = 0.19$, $SEM = 0.06$, 95% CI = 0.10 – 0.33). This difference is much smaller in the older age group (Direct Feedback $M = 0.77$, $SEM = 0.06$, 95% CI = 0.63 – 0.87; Reputation $M = 0.71$, $SEM = 0.07$, 95% CI = 0.55 – 0.83; Baseline $M = 0.64$, $SEM = 0.08$, 95% CI = 0.48 – 0.78), though the overall pattern is the same. Indeed, whereas the effect of condition remains significant when looking at the youngest group alone (LRT $\chi^2(1) = 31.14$, $p < 0.001$), it is not significant when looking at the older group alone (LRT $\chi^2(1) =$

1.75, $p = 0.42$). And, this significant effect of condition in the younger group is driven by differences between the Direct Feedback condition and the Reputation (OR = 9.16, SE = 0.55, $z = 4.05$, $p < 0.001$) and Baseline conditions (OR = 24.55, SE = 0.59, $z = 5.41$, $p < 0.001$), not by differences between the Baseline and Reputation conditions (OR = 2.67, SE = 0.54, $z = 1.81$, $p = 0.07$).

Probability of choosing the more fair option as a function of condition and age group

Figure 2. Probability of choosing the fairer option extracted (predicted effects) from a model including condition and age group as predictors. Error bars are 95% confidence intervals.

3.2 Third-party punishment promotes fairness through direct feedback and not through reputational concerns

We next turn to the question of the mechanism by which the Direct Feedback condition leads to more fair allocations than either the Reputation or Baseline conditions. One possibility is that punishment promotes fairness through the threat of its imposition as opposed to direct feedback. In such a case, we would expect to find a difference between the Direct Feedback condition and the other two conditions on the first trial, especially because the effect of direct feedback was explicitly demonstrated and comprehension about this aspect of the Direct Feedback condition was assessed before the first trial (see Methods). However, looking only at behavior on the first trial, we find no interaction between age and condition (model 2; LRT $\chi^2(2) = 1.61$, $p = 0.45$; Table 1), nor do we find an overall effect of condition (LRT $\chi^2(2) = 0.60$, $p = 0.74$). We also do not find an influence of condition on the first trial when looking only at the younger group (LRT $\chi^2(2) = 2.90$, $p = 0.23$). This lack of effect suggests that the mere possibility of direct feedback occurring is not sufficient to motivate fair behavior.

Figure 3. Probability of choosing the fairer option extracted (predicted effects) from a model including condition, age group and trial number as predictors. Error bars are 95% confidence intervals.

3.3 Third-party punishment leads to more fair behavior through repeated opportunities for punishment

Rather, we find that direct feedback operates over time to increase fairness (Figure 3). While we do not find a 3-way interaction between age, condition, and trial number (model 3; LRT $\chi^2(2) = 3.26$, $p = 0.20$; Table 1), we do find a significant influence of 2-way interactions (LRT $\chi^2(5) = 27.03$, $p < 0.001$). In particular, we find an interaction between trial number and the Direct Feedback condition, relative to Baseline (OR = 2.45, SE = 0.25, $z = 3.55$, $p < 0.001$) and relative to the Reputation condition (OR = 2.18, SE = 0.25, $z = 3.13$, $p = 0.002$). In contrast, there was no significant influence of trial number for the Reputation condition, relative to Baseline (OR = 1.13, SE = 0.22, $z = 0.53$, $p = 0.60$). In addition, we find that this influence of trial number in the Direct Feedback condition is not due to distributions were available to the participant on a given trial (see Supplemental Results). Thus, the influence of trial number was significantly stronger for those in the Direct Feedback condition than those in the Reputation or Baseline conditions, and there was no influence of trial number in the Reputation condition compared to the Baseline condition.

The nature of our design permits an alternative explanation for the influence of trial number in the Direct Feedback condition. In particular, trial number co-varied with the difficulty of the choice presented to participants. On the first 2 trials, participants chose between an equal split of the stake (3/3) and the participant keeping 4 resources for themselves while giving 2 resources to their partner (4/2). Trials 3 and 4 presented participants with a choice between either a 3/3 or 5/1 split, while the final 2 trials presented

a choice between a 4/2 and 5/1 split. Thus, whereas choosing the fair option may have been easier or less tempting earlier in the experiment (e.g. because earning 3 is not much less than 4), the more tempting unfair options (e.g. earning 5 versus 3) on subsequent trials may have made continuing to be fair difficult. Alternatively, children may have noticed that choosing the more unfair option was more obvious on certain trials than others, for instance when the unfair choice was quite unfair (i.e. a 5/1 split) compared to the other choice (i.e. a 3/3 split). These explanations would predict an influence of the offer at hand especially in the Baseline and Reputation conditions, where there is no threat of direct feedback to deter unfair behavior. However, we find no such influence: in both the Reputation and Baseline conditions, we find no interaction between distributions available and age (Figure S1; Reputation condition [model 4]: LRT $X^2(2) = 1.49$, $p = 0.47$; Baseline condition [model 5]: LRT $X^2(2) = 0.10$, $p = 0.95$) or main effect of distribution (Reputation condition: LRT $X^2(2) = 2.28$, $p = 0.32$; Baseline condition: LRT $X^2(2) = 3.73$, $p = 0.15$). Thus, the influence of trial number that we observe in the Direct Feedback condition cannot be fully explained by an influence of the distributions available.

3.4 Children's free responses support a role for direct feedback in increasing focus on fairness

Additional evidence supporting the unique role of the Direct Feedback condition in promoting adherence to fairness norms in younger participants comes from participants' explanations of their behavior, which were provided at the end of the session (see Methods). When asked why they made the decisions they did, 36% of 6- and 7-year-olds in the Direct Feedback condition referenced fairness, more than those

in the Reputation (23%) or Baseline (18%) conditions. This was not true amongst 8- and 9-year-olds (Direct Feedback: 42%; Reputation: 35%; Baseline: 41%). Additionally, fewer 6- and 7-year-olds in the Direct Feedback condition referenced personal benefit (28%) compared to those in the Reputation (41%) and Baseline (42%) conditions. This was again not seen in the 8- and 9-year-olds (Direct Feedback: 18%; Reputation: 38%; Baseline: 15%), suggesting a unique shift from focusing on personal benefit to a concern for fairness in this group of participants.

4. Discussion

We find that third-party punishment promotes fair behavior, and is particularly influential in shaping the choices of 6- and 7- year-old children. For this age group, those in the Direct Feedback condition chose the fairer allocation a higher proportion of the time, relative to both the Reputation and Baseline conditions. This was not true in the 8- and 9-year-old children: we found no difference in selection of the fairer allocation between our three conditions. The influence of third-party punishment in the younger age group is not due merely to the threat of punishment—we find no difference between conditions on the first trial, before punishment has ever been administered. Rather, punishment operates over trials to motivate fair behavior: we observe significantly higher rates of choosing the fairer option over time in the Direct Feedback condition relative to the Baseline and Reputation conditions, which is a particularly striking pattern when examining the behavior of our younger participants (Figure 3).

This pattern of results is more consistent with the account that third-party punishment motivates fair behavior by providing direct feedback and less consistent

with the account that third-party punishment motivates fair behavior *solely* through influences on reputation. In particular, we find that children are more likely to select the most fair option available to them when in the Direct Feedback condition compared to the Reputation condition. This was especially true for younger children and less so for older children. Furthermore, participants' own justifications for their behavior hint at a greater focus on fairness selectively in the Direct Feedback condition for the younger group only. In addition, only in the Direct Feedback condition did increasing number of trials lead to increasingly fair behavior. Potentially, this influence reflects the increasing likelihood of experiencing punishment in response to selfishness as the number of trials increases. Finally, we find no difference in participants' behavior between the Reputation and Baseline conditions, both overall and when looking at the younger age group only, which is inconsistent with an influence of third-party punishment through reputation.

To our knowledge, our study is the first to demonstrate that third-party punishment promotes fairness in children. This result aligns with complementary work examining punishment and cooperation in childhood (Lergetporer et al., 2014) as well as in adulthood (Ambrus & Greiner, 2012; Cinyabuguma et al., 2006; Fehr & Gächter, 2000, 2002; Ones & Putterman, 2007; Pfattheicher et al., 2018; Rockenbach & Milinski, 2006). While the fact that punishment promotes fairness in childhood is somewhat unsurprising from this perspective, it is nevertheless important to empirically demonstrate. Furthermore, our results are the first to shed light on *how* such

punishment promotes fairness, thereby providing insight into the mechanism underlying the emergence of fair behavior and the role that punishment can play.

These results connect to a number of bodies of research. First, our results speak to work on the development of fairness (Blake et al., 2015; McAuliffe, Blake, et al., 2017). In particular, we find that third-party punishment is most effective at motivating fair behavior in children at 6-7 years of age. At this age, children are within a transitional period, shifting from an egoistic focus on disadvantageous inequity to a more broad concept of fairness (Blake et al., 2015; McAuliffe, Blake, et al., 2017). By 8 years of age, children begin expressing this broader concept of fairness, rejecting offers that benefit themselves and disadvantage others (Blake et al., 2015; McAuliffe, Blake, et al., 2017). At this age, we do not find that third-party punishment motivates fair behavior beyond their elevated baseline. Potentially then, third-party punishment has an impact on younger children by encouraging them to generalize the concept of fairness, which they typically only apply to themselves, to others as well. This would explain why punishment was less effective in older children: at this age, children already believe fairness applies to others, and so third-party punishment is not needed (or effective). Indeed, baseline rates of choosing the fairer allocation are already quite high in the older children (~60%)—much higher than in the younger children (~20%). As a testament to the effectiveness of third-party punishment, we find that such punishment motivates younger children to cooperate by the end of the experiment even more than older children in the Reputation and Baseline conditions (though not the Direct Feedback condition), as seen in the non-overlapping confidence intervals for these

conditions. In sum, our results shed light on the emergence of a more adult-like concept of fairness, suggesting that direct feedback provided by third-party punishment may play a role.

Our results also shed light on the mechanisms by which third-party punishment may shape behavior. We find that third-party punishment leads to more fair behavior over time. This is consistent with the idea that punishment can provide instruction (or teach) (Caro & Hauser, 2004; Martin & Cushman, 2016; Raihani et al., 2012): with repeated opportunities to be punished, children were increasingly fair. Recent work has demonstrated that learning from evaluative feedback, like punishment, is easier if it is treated as communication rather than a form of reward or punishment (Ho et al., 2015, 2016). That is, if you punish me, I can either attempt to minimize that negative feedback or I can think about what you might be communicating. Punishment is a more effective way to change behavior in the latter case. Whether such reasoning underlies the adherence to fairness norms prompted by third-party punishment in younger children or whether a more basic cost/benefit calculation (i.e. how to avoid losing resources from punishment while gaining the most resources from the selection of distribution) is at play is unclear. However, we view this latter possibility as less likely for a few reasons. First, children's explanations of their decisions in our task provide some initial evidence that those in the Direct Feedback condition were more sensitive to fairness norms. Specifically, 6-7-year-olds in the Direct Feedback condition referenced fairness more and personal benefit less than children in the other two conditions; we did not find the same pattern in the 8-9-year-olds. Strong conclusions should not be drawn

from these descriptive data, but they nevertheless are consistent with the idea that those in the Direct Feedback condition were more likely to have fairness norms in mind than those in the other conditions. Second, as just outlined, such a strategy would not be optimal from a learning perspective. Children who simply focus on maximizing payoff (by avoiding punishment) will fail to learn the generalizable norms that may be being taught to them through punishment. Failing to learn in this way would be a more costly strategy in the long term. Of course, that is less true in our task, which only involved 6 trials, but is a large concern throughout life more generally, where failing to learn generalizable norms could be quite costly. Third, the fact that the distribution available varied across trials is inconsistent with a simple version of cost/benefit calculation. In this context, the application of an abstract rule would have been necessary for children to be fair, as they were. Similarly, a simple cost/benefit account would be difficult to implement in our task, as what might maximize one's payoffs given the potential for punishment was dependent on the payoff for oneself, the payoff of the partner as well as the probability of punishment given the selection of each option. Furthermore, children for whom third-party punishment was most effective in our task do not yet have a consistent concept of fairness as applied to others: while children at 6-7 years of age recognize distributions that disadvantage themselves as unfair, they do not yet fully recognize distributions that advantage themselves at the expense of another as unfair (Blake et al., 2015; McAuliffe, Blake, et al., 2017). Rather than implementing such a complex calculation, we view it as more plausible that children simply realized they needed to be fair to avoid punishment. For these reasons, we believe it is more

plausible that learning from evaluative feedback undergirds the behavior exhibited in our task. Of course, our results cannot completely rule such a cost/benefit calculation account out, nor were they designed to; whether third-party punishment promotes fairness through learning from evaluative feedback or a more simple cost/benefit calculation—i.e. the mechanism by which third-party punishment shapes behavior—is an important topic for future work. Our contribution is instead to demonstrate that third-party punishment is more likely to encourage fair behavior by providing direct feedback than an influence on reputation. Moreover, we note that our results do not speak to all forms of punishment, and do not suggest that second-party punishment cannot motivate fair behavior.

Our study points to a few exciting avenues for future work. Foremost among these is the role that reputation plays in motivating fairness. First, we note that our results should not be taken as suggesting that reputation does not influence fairness; it does (Engelmann & Rapp, 2018; Warneken, 2017). Rather, our goal was to specifically determine whether third-party punishment operates through direct feedback or through an influence on reputation. In creating a tight comparison between the Direct Feedback and Reputation conditions, we used a minimal form of reputational influence: participants were observed by a third party who they would never interact with again and who had no impact on any future opportunities the participant might experience. In reality, reputation can have important consequences. A more powerful manipulation might have included the possibility of the participant's reputation having future consequences, such as if the observer had the chance to select the participant for a

valuable opportunity. Recent work on the influence of reputation on competitive altruism has implemented such a design, finding that 8-year-olds share more in such circumstances but 5-year-olds do not (Herrmann et al., 2019). And indeed, while we did not find a difference in fair behavior between the Reputation and Baseline conditions in the younger group, this difference was marginally significant ($p = 0.07$) suggesting reputational concerns may have played a weak role in our task. A sensitivity analysis indicated that we do not have strong power to detect such weak effects, further suggesting we should not strongly reject the idea that reputation may matter in our design. Furthermore, an experimenter observed the child's choice in the Baseline condition (i.e. the decision was not anonymous), which could have contributed to a small difference between these two conditions. Future work on the influence of reputation in the context of third-party punishment should manipulate whether the participant's choice is anonymous. Thus, while a minimal form of reputational influence is sufficient for us to dissociate reputational vs. direct feedback components of third-party punishment, our results should not be taken as suggesting that reputation has no influence on fair behavior. Moreover, these results should also not be taken as suggesting that third-party punishment never operates through an influence on reputation. Rather, our results are intended to demonstrate the efficacy of direct feedback, through third-party punishment, in motivating adherence to fairness norms.

Another avenue for future work is the influence that the *type* of third party has. Here, we used an observer and a punisher who were authorities (an experimenter). For children, authorities tend to be the ones that administer punishment. We also know that

children want to appear fair in front of adults (Shaw et al., 2014). Given that children interact with adults so frequently, we might expect such an authority figure to exert a large influence. Furthermore, participants in our experiment were tested in an experimental setting, where deference to authority may be greater. In addition, there are good reasons to avoid using a peer as the third-party. Principally, given that our goal is to test the influence of punishment, our study design required that the third party punish, and do so in a particular way (only when the less fair option was picked). Using a participant (rather than an adult confederate) as the third party precluded the kind of control necessary to implement this design. Furthermore, while child confederates could in principle be used, doing so would introduce a number of logistical challenges (i.e. finding a gender- and age-matched child for each age and gender tested, finding children who could be sufficiently convincing as confederates, etc.) as well as ethical issues about asking children to be involved in research to this extent. Given that our research question did not require such peer punishers, we avoided this issue. Nevertheless, future work should explore whether similar results are found with other types of third parties, peers in particular. School-aged children are highly sensitive to social peers, and so it is possible that both punishment and concerns about being observed may be different when interacting with those around one's own age.

Finally, our results may have implications beyond the motivation of fair behavior. That is, there is nothing inherent about the relation between the direct feedback provided by third-party punishment in our experiment and fairness. Presumably, such third-party punishment can motivate any behavior which children understand to be

desirable or normative but which is difficult to adhere to. For instance, we might expect direct feedback to motivate children to be generous. Future work should explore this possibility and the ability of third-party punishment to motivate normative behavior in children more generally. In addition, our results may have implications beyond the social context used here. That is, our study does not specifically investigate whether the direct feedback provided by punishment is effective only when administered by a social agent or is similarly effective when provided in an asocial manner. However, there are good reasons to think that direct feedback from social agents may be particularly effective. First, a large body of research suggests that children are particularly sensitive to other social agents (Carpendale & Lewis, 2004; Frith & Frith, 2003; Johnson, Slaughter, & Carey, 1998; Woodward, 1998, 2009; c.f. Gergely & Csibra, 2003). Second, punishment in the real world tends to come from social agents and so it is reasonable to expect children and humans more generally to be particularly attuned to punishment provided by social agents. Third, in our own task, we find some sensitivity to being observed, as noted above. Specifically, we find a trending ($p = 0.07$) influence of the Reputation condition (relative to the Baseline condition) in younger children. This suggests that merely being observed by an additional social agent may lead children to be more fair, though more work and a higher-powered study would be needed to verify this influence.

To conclude, we investigate the influence of third-party punishment on adherence to fairness norms. For the first time, we show that third-party punishment promotes fairness, and does so most strongly for children around 6-7 years of age.

This nicely mirrors a role for punishment in promoting other behaviors important for cooperation (Lergetporer et al., 2014). Our findings provide more support for an influence of third-party punishment through its ability to provide direct feedback than concerns about being observed. This feedback is especially influential as children are transitioning to a relatively broader and less egoistic—a more impartial—concept of fairness. Furthermore, these findings provide insight into the mechanisms by which punishment, and third-party punishment in particular, promote fair behavior.

Punishment provides direct feedback about what kinds of behavior are valued and expected—feedback which can make us more fair, even when fairness comes at personal cost.

Table 1. Estimate and standard error of fixed effects in Generalized Linear Mixed Models predicting children's less selfish behavior. Baselines were set as follows (except where noted): Condition = Baseline, Gender = female, Trial type= 3/3 vs. 4/2.

	Model 1 (Condition X Age)	Model 1 (Condition X Age, Reputation baseline)	Model 2 (Condition X Age, first trial)	Model 2 (Condition X Age, first trial, Reputation baseline)	Model 3 (Condition X Age X Trial)	Model 4 (Distribution X Age, Reputation condition)	Model 5 (Distribution X Age, Baseline condition)
Intercept	-0.28 (0.30)	0.17 (0.29)	0.44 (0.40)	0.23 (0.40)	-0.29 (0.31)	0.36 (0.42)	-0.67 (0.61)
Cond: Direct Feedback	1.76 (0.37)***		0.15 (0.48)		2.06 (0.41)***		
Cond: Reputation	0.45 (0.36)		-0.21 (0.49)		0.46 (0.38)		
Age (months)	0.91 (0.26)***	0.76 (0.25)**	0.90 (0.38)*	0.97 (0.38)**	0.93 (0.28)***	0.77 (0.35)*	1.19 (0.48)*
Gender	-0.05 (0.30)	-0.05 (0.30)	-0.34 (0.40)	-0.34 (0.40)	-0.06 (0.31)	-0.35 (0.52)	0.42 (0.77)
Cond: Direct Feedback X Age (months)	-1.06 (0.38)**		-0.52 (0.52)		-1.30 (0.42)**		
Cond: Rep. X Age (months)	-0.15 (0.36)		0.07 (0.53)		-0.15 (0.37)		
Cond: Baseline (Reputation is ref. level)		-0.45 (0.36)		0.21 (0.49)			
Cond: Direct Feedback (Reputation is ref. level))		1.31 (0.37)***		0.35 (0.48)			
Cond: Baseline (Rep. ref. level) X Age (months)		0.15 (0.36)		-0.07 (0.53)			
Cond: Direct Feedback (Rep. ref. level) X Age (months)		-0.91 (0.37)*		-0.60 (0.52)			
Trial number					0.03 (0.16)		
Cond: Direct Feedback X Trial number					0.99 (0.27)***		
Cond: Rep. X Trial number					0.11 (0.22)		
Age (months) X Trial number					-0.08 (0.16)		
Cond: Pun X Age (months) X Trial number					-0.43 (0.27)		
Cond: Rep. X Age (months) X Trial number					0.01 (0.22)		
Distribution: 4/2 vs. 5/1						0.28 (0.37)	0.27 (0.42)
Distribution: 3/3 vs. 5/1						-0.30 (0.38)	-0.55 (0.43)
Distribution: 4/2 vs. 5/1 X Age (months)						-0.12 (0.38)	-0.08 (0.42)

Distribution: 3/3
vs. 5/1 X Age
(months)

0.34 (0.39) 0.06 (0.43)

AIC	847.64	847.64	163.22	163.22	827.12	304.20	278.47
BIC	884.34	884.34	182.79	182.79	891.34	332.05	306.51
Log Likelihood	-415.82	-415.82	-74.61	-74.61	-399.56	-144.10	-131.24
Number of observations	726	726	121	121	726	240	246
Number of participants	121	121			121	40	41
Variance (participants)	1.52	1.52			1.71	1.61	4.08
Deviance			149.22	149.22			

***p < 0.001, **p < 0.01, *p < 0.05

References

- Agresti, A., & Coull, B. A. (1998). Approximate is better than “Exact” for interval estimation of binomial proportions. *American Statistician*, 52(2), 119–126.
<https://doi.org/10.1080/00031305.1998.10480550>
- Ambrus, A., & Greiner, B. (2012). Imperfect Public Monitoring with Costly Punishment: An Experimental Study. *The American Economic Review*, 102(7), 3317–3332.
<https://doi.org/10.1257/aer.102.7.3317>
- Barclay, P. (2006). Reputational benefits for altruistic punishment. *Evolution and Human Behavior*, 27(5), 325–344. <https://doi.org/10.1016/j.evolhumbehav.2006.01.003>
- Barclay, P., & Willer, R. (2007). Partner choice creates competitive altruism in humans. *Proceedings. Biological Sciences / The Royal Society*, 274(December 2006), 749–753. <https://doi.org/10.1098/rspb.2006.0209>
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). *lme4: Linear mixed-effects models using Eigen and S4. R package version 1.1-7*.
- Bernhard, R., Martin, J. W., & Warneken, F. (2020). Why do children punish? Fair outcomes matter more than intent in children’s second- and third-party punishment. *Journal of Experimental Child Psychology*, 200, 104909.
<https://doi.org/10.31234/osf.io/hdyfv>
- Blake, P. R., McAuliffe, K., Corbit, J., Callaghan, T. C., Barry, O., Bowie, A., Kleutsch, L., Kramer, K. L., Ross, E., Vongsachang, H., Wrangham, R., & Warneken, F. (2015). The ontogeny of fairness in seven societies. *Nature*, 528(7581), 258–261.
<https://doi.org/10.1038/nature15703>

- Bone, J. E., & Raihani, N. J. (2015). Human punishment is motivated by both a desire for revenge and a desire for equality. *Evolution and Human Behavior*.
<https://doi.org/10.1016/j.evolhumbehav.2015.02.002>
- Bornhorst, F., Ichino, A., Kirchkamp, O., Schlag, K. H., & Winter, E. (2010). Similarities and differences when building trust: The role of cultures. *Experimental Economics*, 13(3), 260–283. <https://doi.org/10.1007/s10683-010-9240-x>
- Bregant, J., Shaw, A., & Kinzler, K. D. (2016). Intuitive Jurisprudence: Early Reasoning About the Functions of Punishment. *Journal of Empirical Legal Studies*, 13(4), 693–717. <https://doi.org/10.1111/jels.12130>
- Caro, T. M., & Hauser, M. D. (2004). Is There Teaching in Nonhuman Animals? *The Quarterly Review of Biology*, 67(2), 151–174. <https://doi.org/10.1086/417553>
- Carpendale, J. I. M., & Lewis, C. (2004). Constructing an understanding of mind: The development of children's social understanding within social interaction. *Behavioral and Brain Sciences*, 27(01), 79–151. <https://doi.org/10.1017/s0140525x04000032>
- Charness, G., Cobo-Reyes, R., & Jiménez, N. (2008). An investment game with third-party intervention. *Journal of Economic Behavior and Organization*, 68(1), 18–28.
<https://doi.org/10.1016/j.jebo.2008.02.006>
- Cinyabuguma, M., Page, T., & Putterman, L. (2006). Can second-order punishment deter perverse punishment? *Experimental Economics*, 9(3), 265–279.
<https://doi.org/10.1007/s10683-006-9127-z>
- Clutton-Brock, T., & Parker, G. (1995). Punishment in animal societies. *Nature*, 373, 209–216.

http://www.anth.ucsb.edu/faculty/gurven/anth169/punishment_in_animal_societies.pdf

Coricelli, G., Fehr, D., & Fellner, G. (2004). Partner selection in public goods experiments. *Journal of Conflict Resolution*, 48(3), 356–378.
<https://doi.org/10.1177/0022002704264143>

Crockett, M. J., Özdemir, Y., & Fehr, E. (2014). The Value of Vengeance and the Demand for Deterrence. *Journal of Experimental Psychology: General*.

Davidov, M., Grusec, J. E., & Wolfe, J. L. (2012). Mothers' Knowledge of Their Children's Evaluations of Discipline: The Role of Type of Discipline and Misdeed, and Parenting Practices. *Merrill-Palmer Quarterly*, 58(3), 314–340.
<https://doi.org/10.1353/mpq.2012.0018>

Engelmann, J. M., Herrmann, E., & Tomasello, M. (2012). Five-Year Olds, but Not Chimpanzees, Attempt to Manage Their Reputations. *PLoS ONE*, 7(10).
<https://doi.org/10.1371/journal.pone.0048433>

Engelmann, J. M., Over, H., Herrmann, E., & Tomasello, M. (2013). Young children care more about their reputation with ingroup members and potential reciprocators. *Developmental Science*, 16(6), 952–958. <https://doi.org/10.1111/desc.12086>

Engelmann, J. M., & Rapp, D. J. (2018). The influence of reputational concerns on children's prosociality. *Current Opinion in Psychology*, 20, 92–95.
<https://doi.org/10.1016/j.copsyc.2017.08.024>

Engelmann, J. M., & Tomasello, M. (2019). Children's Sense of Fairness as Equal Respect. *Trends in Cognitive Sciences*, 23(6), 454–463.

<https://doi.org/10.1016/j.tics.2019.03.001>

Fehr, E., & Fischbacher, U. (2003). The nature of human altruism. *Nature*, 425(6960), 785–791. <https://doi.org/10.1038/nature02043>

Fehr, E., & Fischbacher, U. (2004). Social norms and human cooperation. *Trends in Cognitive Sciences*, 8(4). <https://doi.org/10.1016/j.tics.2004.02.007>

Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 151(3712), 867–868. <https://doi.org/10.1126/science.151.3712.867-a>

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415(6868), 137–140. <https://doi.org/10.1038/415137a>

Fehrler, S., & Przepiorka, W. (2016). Choosing a partner for social exchange: Charitable giving as a signal of trustworthiness. *Journal of Economic Behavior and Organization*, 129, 157–171. <https://doi.org/10.1016/j.jebo.2016.06.006>

FeldmanHall, O., Sokol-Hessner, P., Van Bavel, J. J., & Phelps, E. A. (2014). Fairness violations elicit greater punishment on behalf of another than for oneself. *Nature Communications*, 5, 5306. <https://doi.org/10.1038/ncomms6306>

Frith, U., & Frith, C. D. (2003). Development and neurophysiology of mentalizing. *Philosophical Transactions of the Royal Society of London, B*, 358, 459–474.

Gächter, S., Renner, E., & Sefton, M. (2008). The long-run benefits of punishment. *Science*, 322(5907), 1510. <https://doi.org/10.1126/science.1164744>

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

[https://doi.org/10.1016/S1364-6613\(03\)00128-1](https://doi.org/10.1016/S1364-6613(03)00128-1)

Green, P., & Macleod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493–498. <https://doi.org/10.1111/2041-210X.12504>

Hardy, C. L., & Van Vugt, M. (2006). Nice guys finish first: The competitive altruism hypothesis. *Personality and Social Psychology Bulletin*, 32(10), 1402–1413. <https://doi.org/10.1177/0146167206291006>

Henrich, J. (2006). Cooperation, Punishment, and the Evolution of Human Institutions. *Science*, 312(5770), 60–61. <https://doi.org/10.1126/science.1126398>

Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., & Ziker, J. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science (New York, N.Y.)*, 327(5972), 1480–1484. <https://doi.org/10.1126/science.1182238>

Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., Lesorogol, C., Marlowe, F., Tracer, D., & Ziker, J. (2006). Costly punishment across human societies. *Science (New York, N.Y.)*, 312(5781), 1767–1770. <https://doi.org/10.1126/science.1127333>

Herrmann, B., Thoni, C., & Gächter, S. (2008). Antisocial Punishment Across Societies. *Science*, 319(5868), 1362–1367. <https://doi.org/10.1126/science.1153808>

Herrmann, E., Engelmann, J. M., & Tomasello, M. (2019). Children engage in competitive altruism. *Journal of Experimental Child Psychology*, 179, 176–189.

<https://doi.org/10.1016/j.jecp.2018.11.008>

Ho, M. K., Littman, M. L., Cushman, F., & Austerweil, J. L. (2015). Teaching with Rewards and Punishments: Reinforcement or Communication? *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 920–925.

Ho, M. K., MacGlashan, J., Littman, M. L., & Cushman, F. (2016). Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition*. <https://doi.org/10.1016/j.cognition.2017.03.006>

Holden, G. W., Coleman, S. M., & Schmidt, K. L. (1995). Why 3-Year-Old Children Get Spanked: Parent and Child Determinants as Reported by College-Educated Mothers. *Merrill-Palmer Quarterly*, 41(4), 431–452.

<http://www.jstor.org/stable/23087935><http://about.jstor.org/terms>

House, B. R., Silk, J. B., Henrich, J., Barrett, H. C., Scelza, B. A., Boyette, A. H., Hewlett, B. S., McElreath, R., & Laurence, S. (2013). Ontogeny of prosocial behavior across diverse societies. *Proceedings of the National Academy of Sciences*, 110(36), 14586–14591. <https://doi.org/10.1073/pnas.1221217110>

Johnson, S., Slaughter, V., & Carey, S. (1998). Whose gaze will infants follow? The elicitation of gaze following in 12-month-olds. *Developmental Science*, 1(2), 233–238. <https://doi.org/10.1111/1467-7687.00036>

Jordan, J. J., Hoffman, M., Bloom, P., & Rand, D. G. (2016). Third-party punishment as a costly signal of trustworthiness. *Nature*, 530(7591), 473–476. <https://doi.org/10.1038/nature16981>

Jordan, J. J., & Rand, D. G. (2020). Signaling when no one is watching: A reputation

heuristics account of outrage and punishment in one-shot anonymous interactions.

Journal of Personality and Social Psychology, 118(1), 57–88.

<https://doi.org/10.1037/pspi0000186>

Lammers, J., Gast, A., Unkelbach, C., & Galinsky, A. D. (2017). Moral Character

Impression Formation Depends on the Valence Homogeneity of the Context. *Social Psychological and Personality Science*, 194855061771458.

<https://doi.org/10.1177/1948550617714585>

Leimgruber, K. L., Shaw, A., Santos, L. R., & Olson, K. R. (2012). Young Children Are

More Generous when Others Are Aware of Their Actions. *PLoS ONE*, 7(10).

<https://doi.org/10.1371/journal.pone.0048292>

Lergetporer, P., Angerer, S., Glätzle-Rützler, D., & Sutter, M. (2014). Third-party

punishment increases cooperation in children through (misaligned) expectations and conditional cooperation. *Proceedings of the National Academy of Sciences*,

111(19), 6916–6921. <https://doi.org/10.1073/pnas.1320451111>

Martin, J. W., & Cushman, F. (2016). The Adaptive Logic of Moral Luck. *A Companion to Experimental Philosophy*, 190–202.

<https://doi.org/10.1002/9781118661666.ch12>

Martin, J. W., Jordan, J. J., Rand, D. G., & Cushman, F. (2019). When do we punish people who don't? *Cognition*, 193(August), 104040.

<https://doi.org/10.1016/j.cognition.2019.104040>

Martin, J. W., Martin, S., & McAuliffe, K. (2021, March 10). *Third-party punishment promotes fairness in children*. <https://doi.org/DOI.10.17605/OSF.IO/K4Q8C>

- McAuliffe, K., Blake, P. R., Steinbeis, N., & Warneken, F. (2017). The developmental foundations of human fairness. *Nature Human Behaviour*, 1(2), 0042.
<https://doi.org/10.1038/s41562-016-0042>
- McAuliffe, K., Blake, P. R., & Warneken, F. (2014). Children reject inequity out of spite. *Biology Letters*, 10(12), 1–5. <https://doi.org/10.1098/rsbl.2014.0743>
- McAuliffe, K., Jordan, J. J., & Warneken, F. (2015). Costly third-party punishment in young children. *Cognition*, 134, 1–10.
<https://doi.org/10.1016/j.cognition.2014.08.013>
- McAuliffe, K., Raihani, N. J., & Dunham, Y. (2017). Children are sensitive to norms of giving. *Cognition*, 167, 151–159. <https://doi.org/10.1016/j.cognition.2017.01.006>
- Ones, U., & Putterman, L. (2007). The ecology of collective action: A public goods and sanctions experiment with controlled group formation. *Journal of Economic Behavior and Organization*, 62(4), 495–521.
<https://doi.org/10.1016/j.jebo.2005.04.018>
- Page, T., Putterman, L., & Unel, B. (2005). Voluntary association in public goods experiments: Reciprocity, mimicry and efficiency. *Economic Journal*, 115(506), 1032–1053. <https://doi.org/10.1111/j.1468-0297.2005.01031.x>
- Pfattheicher, S., Böhm, R., & Kesberg, R. (2018). The Advantage of Democratic Peer Punishment in Sustaining Cooperation within Groups. *Journal of Behavioral Decision Making*, 31(4), 562–571. <https://doi.org/10.1002/bdm.2050>
- Pleasant, A., & Barclay, P. (2018). Why Hate the Good Guy? Antisocial Punishment of High Cooperators Is Greater When People Compete to Be Chosen. *Psychological*

Science, 095679761775264. <https://doi.org/10.1177/0956797617752642>

Przepiorka, W., & Liebe, U. (2016). Generosity is a sign of trustworthiness-the punishment of selfishness is not. *Evolution and Human Behavior*, 37(4), 255–262. <https://doi.org/10.1016/j.evolhumbehav.2015.12.003>

Raihani, N. J., & Barclay, P. (2016). Exploring the trade-off between quality and fairness in human partner choice. *Royal Society Open Science*, 3(11), 160510. <https://doi.org/10.1098/rsos.160510>

Raihani, N. J., & Bshary, R. (2015). The reputation of punishers. *Trends in Ecology & Evolution*, 1–6. <https://doi.org/10.1016/j.tree.2014.12.003>

Raihani, N. J., & Bshary, R. (2019). Punishment: one tool, many uses. *Evolutionary Human Sciences*, 1, e12. <https://doi.org/10.1017/ehs.2019.12>

Raihani, N. J., Thornton, A., & Bshary, R. (2012). Punishment and cooperation in nature. *Trends in Ecology & Evolution*, 27(5), 288–295. <https://doi.org/10.1016/j.tree.2011.12.004>

Rockenbach, B., & Milinski, M. (2006). The efficient interaction of indirect reciprocity and costly punishment. *Nature*, 444(7120), 718–723. <https://doi.org/10.1038/nature05229>

Shaw, A., Montinari, N., Piovesan, M., Olson, K. R., Gino, F., & Norton, M. I. (2014). Children develop a veil of fairness. *Journal of Experimental Psychology: General*, 143(1), 363–375. <https://doi.org/10.1037/a0031247>

Shaw, A., & Olson, K. R. (2012). Children discard a resource to avoid inequity. *Journal of Experimental Psychology: General*, 141(2), 382–395.

<https://doi.org/10.1037/a0025907>

Sylwester, K., & Roberts, G. (2010). Cooperators benefit through reputation-based partner choice in economic games. *Biology Letters*, 6(5), 659–662.

<https://doi.org/10.1098/rsbl.2010.0209>

Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, 46(1), 35–57. <https://doi.org/10.1086/406755>

Warneken, F. (2017). How Children Solve the Two Challenges of Cooperation. *Annual Review of Psychology*, August 2017, 1–25. <https://doi.org/10.1146/annurev-psych-122216-011813>

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1), 1–34. [https://doi.org/10.1016/S0010-0277\(98\)00058-4](https://doi.org/10.1016/S0010-0277(98)00058-4)

Woodward, A. L. (2009). Infants' Grasp of Others' Intentions. *Current Directions in Psychological Science*, 18(1), 53–57. <https://doi.org/10.1111/j.1467-8721.2009.01605.x>

Yang, F., Choi, Y.-J., Misch, A., Yang, X., & Dunham, Y. (2018). In Defense of the Commons: Young Children Negatively Evaluate and Sanction Free Riders.

Psychological Science, 095679761877906.

<https://doi.org/10.1177/0956797618779061>