Young children use statistical evidence to infer the informativeness of praise

Mika Asaba, **Emily Hembacher**, **Shi Qiu**, **Brett Anderson**, **Michael Frank**, and **Hyowon Gweon**
{masaba, ehembach, shiqiu21, brettand, mcfrank, hyo} @stanford.edu
Department of Psychology, Stanford University

Abstract

Praise is not only rewarding but also informative. It allows us to learn about our skills and competence even when we are uncertain or unable to judge for ourselves. Not all praise is equally meaningful, however: Praise from someone who praises indiscriminately is less informative than from someone who praises selectively. Here we ask whether young children infer the informativeness of others' praise based on the statistical dependence between praise and the quality of work. Exp. 1 shows that adults and 4-5 year-olds were more likely to trust praise from a teacher whose previous praise covaried with the quality of work (i.e., selective praise) than praise from a teacher who indiscriminately praised independent of the quality of work (i.e., overpraise). Exp. 2 addressed the possibility that participants simply prefer a teacher who praises less often. Even for young children, praise is more than something nice. Rather, they can track the informativeness of others' evaluative feedback and use it to learn about the quality of their own work.

Keywords: social cognition, praise, statistical reasoning, selective trust

Evaluative feedback from others is an important source of information for learning about ourselves. In the face of uncertainty about our abilities or performance on a task, feedback from others serves as a useful indicator of quality or success. For instance, a colleague's enthusiastic praise after your big presentation at work (e.g., "that was a great talk!") can be taken as evidence that your talk was indeed good. Not all praise is equally meaningful, however. If your colleague is known to be overgenerous in her evaluation of talks, you might still remain uncertain about how it really went. This is because as adults, we naturally interpret evaluative feedback in the context of the evaluator, considering her expertise, personality, or communicative goals.

Praise is commonly used in parenting and early educational practices across many societies. Although praise is believed to foster intrinsic motivation, it can have negative consequences when it is directed at one's ability or intelligence rather than one's effort (Mueller & Dweck, 1998; Henderlong Corpus & Lepper, 2007). The differential effect of praise depending on *what* is being evaluated suggests that even young children appreciate praise as more than just positive reinforcement (Delin & Baumeister, 1994). Despite abundant work on the effects of praise in early childhood, however, little empirical work has directly investigated how children reason about the informativeness of others' praise and the cognitive capacities that underlie this ability. To begin addressing these questions, here we explore the idea that even the same praise can be differentially informative depending on *who* is giving praise; we ask whether children and adults can use minimal covariation information to infer the informativeness of others' praise and preferentially endorse feedback that is more likely to reflect the true quality of the work.

Researchers have proposed that the perceived sincerity of praise may influence its meaning. When praise is regarded as unjustified or inconsistent with reality, people might infer that the evaluator is merely being "polite" rather than sincere (Brophy, 1981; Delin & Baumeister, 1994). Recent computational work has formalized adults' interpretation of polite speech as inferences about the speaker's communicative goals: an *epistemic* goal to provide accurate information and a *social* goal to make the listener feel good (Yoon, Tessler, Goodman, & Frank, 2016). When the two goals are at odds (e.g., telling the truth can hurt the listener's feelings), adults readily infer that someone who praises poor-quality work is driven by the social goal than the epistemic goal (i.e., wants to be "nice" rather than "honest"). Similarly, given explicit information about the speaker's communicative goal, adults can use praise to infer the actual quality of the work.

In many real-world contexts, however, the meaning of others' feedback can be more ambiguous. When a child seeks feedback on the quality of his drawings, he might be uncertain about his true performance as well as the evaluator's communicative goals. Prior observations of others' praise can help, however. If the learner's past experience suggests that the evaluator's praise has been congruent with the perceived quality of the work being evaluated (e.g., selective praise for high-quality work), he may expect that the evaluator's current feedback is also likely to reflect the true quality of his work. By contrast, if the evaluator's praise has been independent of the quality (e.g., indiscriminate praise for both high- and low-quality work), it suggests that her current praise is also likely nice but uninformative. While opportunities to observe others' praise may be sparse, even minimal covariation information can be sufficient to quickly form a model of the evaluator's praise. Given that even preverbal infants readily draw causal inferences from the statistical dependence between agents and event outcomes (Gweon & Schulz, 2011; Seiver, Gopnik, & Goodman, 2013), such statistical sensitivity might also be recruited for learning about the informativeness of others' evaluative feedback.

It is also possible, however, that young children have difficulty inferring the informativness of others' praise, especially because they might have trouble differentiating informativeness from niceness. Preschool-aged children attribute knowledge to "nice" informants and endorse their testimony even when they lack critical perceptual access (Lane, Wellman, & Gelman, 2013). Children also trust domain-specific information from a nice non-expert over a mean expert (Landrum, Mills, & Johnston, 2013). Even school-aged children are biased to judge positive assessments of others' work (e.g., music or drawings) as more accurate than negative assess-

ments (Boseovski, Marble, & Hughes, 2017). In these studies, children received no information that would allow them to evaluate the quality of others' testimony; this might have led children to rely heavily on information that was explicitly provided in the task, such as informants' niceness or the valence of their feedback. Given clear covariation information about praise and the quality of the work being evaluated, children might be able to distinguish informative from uninformative praise depending on whether it is consistent with the perceived quality of the work.

Further support for this hypothesis comes from prior literature on epistemic trust, which suggests that 3-5 year-old children readily track the informativeness of teachers in pedagogical contexts given clear grounds to evaluate the quality of information provided. Children preferentially learn from teachers who previously provided correct (versus incorrect) labels of familiar objects (e.g., see Sobel & Kushnir, 2013 for a review) and update their evaluations of others' trustworthiness across multiple interactions (Ronfard & Lane, 2017). Beyond tracking inaccuracies, children also recognize more subtle forms of misinformation; when a teacher demonstrates only one of four functions of a toy, children appropriately penalize the teacher for being under-informative (Gweon & Asaba, 2017) and are less likely to trust the teacher in learning about a new toy (Gweon, Pelton, Konopka, & Schulz, 2014). These results suggest that even young children are sensitive to others' informativeness and selectively endorse subsequent information from those who are expected to be informative. Thus, children might demonstrate similar sensitivity to the informativeness of others' praise, especially when there is a ground truth for evaluating the quality of praise.

To test this hypothesis, we investigate adults' and young children's abilities to track the informativeness of others' praise given minimal covariation information. We use an activity with which preschool-aged children are familiar and motivated to improve on: tracing shapes. Given prior work on the effects of person- vs. effort-directed praise (Mueller & Dweck, 1998), here we focus on praise directed at the quality of the tracing (i.e., performance praise). Given a goal to identify which tracing is better, participants' choice to endorse the praise from one of two agents can be used as a proxy for their evaluation of these praisers.

Experiment 1: Selective vs. Overpraise

Adults are sensitive to politeness concerns in communicative contexts (Yoon et al., 2016), but their ability to use minimal statistical information to distinguish informative praise from merely polite comments has not been studied. In Exp.1A, we first verify that adults track others' informativeness and selectively endorse praise from a teacher who had given praise only to the high-quality tracings (Selective Teacher) over a teacher who had given indiscriminate praise (Overpraise Teacher), while choosing the Overpraise Teacher as the one who is trying to be "nice". Exp.1B extends this study to preschool-aged children.

Experiment 1A: Adults

Methods

Participants Eighty-six adults (38 female, M_{Age}(SD) = 36.6(11.9), range: 21-71) were recruited from Amazon's Mechanical Turk. An additional 14 subjects were excluded for failing one or both memory check questions.

Stimuli Images of "good" and "bad" tracings (i.e., a marker tracing that was reasonably aligned or clearly misaligned with the template shape, see Fig.1A) were used in the *Warm-up Phase* and in two videos of teacher-child interactions (Fig.1B). In the video, six tracings were placed in a row on the table, 3 good and 3 bad tracings in alternating order. In both videos, the same child ("Johnny") asked a teacher about the tracings; one video featured "Teacher Jane" who wore a green shirt, and the other featured "Teacher Susan" who wore a red shirt. Different drawings were used for each video.

Procedure In the *Warm-up Phase*, participants were shown two tracings that clearly differed in quality (Fig.1A) and were asked to indicate which one was better; participants who answered incorrectly were excluded from analyses.

In the *Teacher Introduction Phase*, participants watched two videos. In both videos, Johnny first told the teacher that he made the tracings and really wanted to know which of his tracings were good. The teacher then evaluated the tracings one at a time from right to left. In the Overpraise Teacher video, the teacher provided positive, undifferentiated feedback ("Wow, that's great!") for all six tracings and placed a star sticker on each of them (see Fig.1B). In the Selective Teacher video, the teacher provided positive feedback on the good tracings ("Wow, that's great!") and put stickers on them, while giving neutral feedback ("Hm, this one's okay!") for the three bad tracings without giving stickers. Johnny did not receive the stickers himself; they were used to help participants remember which tracings received praise, rather than as an external reward. Both teachers maintained a positive tone for both types of feedback; teacher identity (Jane or Susan), pattern of praise (Selective or Overpraise), and order of presentation were counterbalanced. After each video, participants were asked how many tracings the teacher said were great. Subjects who failed to correctly answer these memory check questions ("3" for the Selective Teacher and "6" for the Overpraise Teacher) were excluded.

Finally, in the *Test Question Phase*, participants were shown a picture of another student (Kristen), and two envelopes, each of which contained one of two tracings she made. Participants were told that the Selective Teacher saw only the tracing in one of the envelopes and praised it (e.g., "Teacher Susan said this tracing is great"), and the Overpraise Teacher saw only the tracing in the other envelope and praised it ("Teacher Jane said this tracing is great"). Participants never saw Kristen's actual tracings, only the envelopes with stickers that indicated which teacher praised the tracing. Participants were asked: "Kristen is going to bring one of her tracings to a contest. Which tracing should she bring?" Addi-

tionally, participants were asked: "One of the teachers wanted to be nice. Who was trying to be nice?"

Results and Discussion

Our primary question was whether adults use teachers' prior patterns of praise to evaluate the informativeness of subsequent praise. As predicted, participants overwhelmingly chose the tracing praised by the Selective Teacher (87.2%, $p < .001$, Binomial Test). Additionally, a majority of participants chose the Overpraise Teacher as the one who was *trying* to be nice (93.0%, $p < 0.001$, Binomial Test). These results suggest that adults readily detect the differences in the informativeness of evaluative feedback from minimal covariation data, and use it to inform their decisions about the quality of hidden products. Furthermore, they inferred the Overpraise Teacher's communicative goal to be "nice".

Experiment 1B: 4-5 year-olds

Next, we investigated whether 4- and 5-year-old children are also sensitive to the informativeness of others' praise. Children were either asked about which of their own tracings was better (Tracing Condition) or which teacher was trying to be nice (Niceness Condition).

Methods

Participants Eighty 4-5 year-olds were recruited from a university preschool. Children were either in the Tracing Condition (N=40, 19 female, M_{Age}(SD) = 4.9(0.4), range = 4.1 - 5.9)[1] or the Niceness Condition (N=40, 21 female, M_{Age}(SD) = 4.9(0.4), range = 4.1 - 5.7). An additional 8 children were tested but excluded due to failure on the warm-up or memory questions.

Stimuli Teacher videos were the same as those from Exp. 1. Additionally, two 8.5"x11" tracing templates (a circle and either an overlapping triangle or rectangle) were used for children to make their own tracings. The tracings in the *Warm-up Phase* were presented on laminated sheets of paper, and videos were presented on a 13" Macbook Pro laptop. We also used printed pictures of the teachers and Johnny, two manila envelopes, and star-shaped red and green stickers.

Procedure Children were tested in a private room in a preschool. In the *Warm-up Phase*, the experimenter first explained what tracing is: "The goal of tracing is to stay as close to the lines as possible" and demonstrated tracing a rectangle. Then, the child traced two templates, and the experimenter put each tracing away into an envelope such that the child could not see them for the remainder of the session. Then, children saw two pairs of tracings (similar in quality to those in the *Warm-up* of Exp. 1A) and were asked to indicate which one was better. Only children who passed both trials were included in the analyses.

In the *Teacher Introduction Phase*, children were shown a picture of a student, Johnny. They were told that Johnny

[1]Planned sample size, exclusion criteria, and analysis plan pre-registered at https://aspredicted.org/4r9dh.pdf.

was working on his tracings earlier and wanted help figuring out which of his tracings were good, because he wanted to show them to his class later. Children then watched the same Selective Teacher and Overpraise Teacher videos as in Exp. 1A. After each video, children saw a still frame of the video (with no stars on the tracings) and were asked which tracings the teacher said were great. If children missed a tracing or incorrectly pointed to a tracing that was not praised by the teacher, they watched the video again and the experimenter asked the same memory check question. Those who failed the memory check even after watching the video again were excluded from analyses.

The only difference between the Tracing and Niceness Conditions was the following *Test Question Phase*. In the Tracing Condition, the experimenter told the child that Teacher Jane and Teacher Susan were nearby and could give feedback on the child's tracings from earlier. The experimenter left the room with the envelopes containing the child's tracings and returned after 15 seconds with stickers attached to the envelopes. The experimenter pointed to the envelope with a green sticker and placed a photo of Teacher Jane next to it, and said: "Teacher Jane looked at this tracing and said that this one is great." She then pointed to the other envelope (with a red sticker and Teacher Susan's photo) and said: "Teacher Susan looked at this tracing and said that this one is great" (teacher order and identity counterbalanced). Finally, with the tracings still in the envelopes, the experimenter said: "Now you can bring back your best tracing to show your teacher! Which one do you think is the best?" Children responded by pointing to one of the envelopes.

In the Niceness Condition, children were not provided feedback on their tracings from the teachers. Instead, they were asked: "Which teacher was trying to be nice?" Children responded by pointing to one of the two teacher photos.

Results and Discussion

Our main question was whether children in the Tracing condition preferentially endorsed the Selective Teacher's praise (i.e., choose the tracing she praised as the best); additionally, we asked whether children in the Niceness condition would choose the Overpraise Teacher as trying to be nice. As predicted, children were significantly more likely to choose the tracing praised by the Selective Teacher (72.5%, $p = 0.006$, Binomial Test, Fig.1D). To investigate whether children's age predicted their choice, we fit a logistic regression model: `Tracing choice ~ Age in Months + (1 | Subject)`. Children's age did not predict their choice of tracing ($B = 0.30$, $z = 0.34$, $p = 0.73$). When asked which teacher was trying to be nice, children overwhelmingly selected the Overpraise Teacher (82.5%, $p < .001$, Binomial Test).[2]

These results suggest that children are more likely to trust

[2]As an exploratory measure, children in the Tracing Condition were also asked this question at the end; they did not show a preference for either teacher (55% chose the Selective Teacher, $p = 0.64$) as trying to be nice, suggesting that these responses may have been influenced by their selection of tracing.
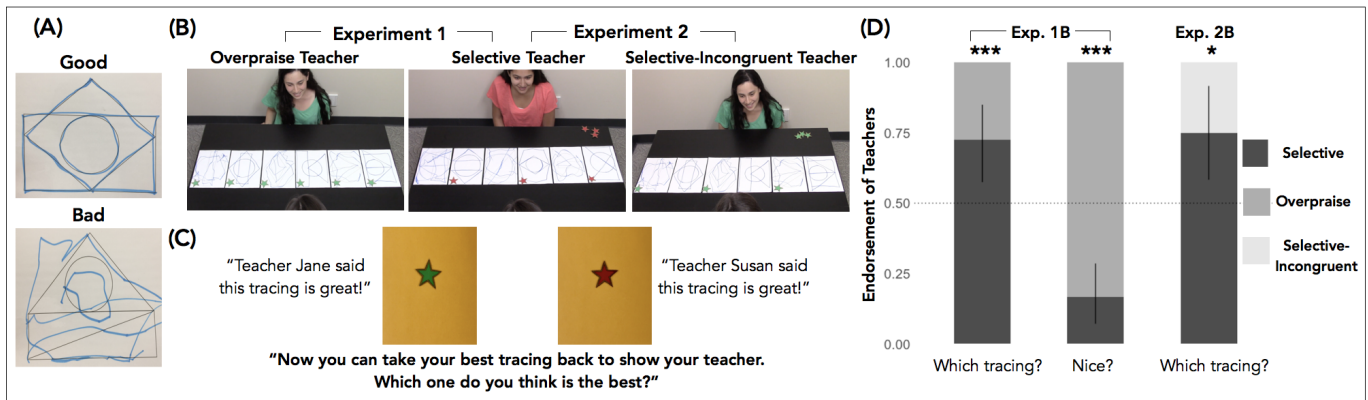
Figure 1: (A) Examples of good and bad tracings for the warm-up questions and teacher videos. (B) Final frames of teacher videos. (C) Set-up and critical question for Exp. 1B & 2B. (D) Results for test question in 4-5 year-olds (* $p < .05$, ** $p < .01$, *** $p < .001$). Error bars are 95% confidence intervals.

the praise given by a teacher who had previously provided selective praise that covaried with the actual quality of the tracings over the praise given by a teacher who had indiscriminately praised all tracings. These results were not driven by a global preference for the Selective Teacher; when asked to choose who was trying to be nice, they appropriately selected the Overpraise Teacher.

The critical difference between the teachers was that the pattern of their feedback appropriately co-varied with the actual quality of the tracings (Selective Teacher) or was indiscriminately positive independent of quality (Overpraise Teacher). Note however that the two teachers' praise also differed in the *frequency* of praise; the Selective Teacher praised only 3 of the 6 tracings while the Overpraise teacher praised all 6. Thus, it is possible that children and adults have a simple heuristic that praise from an infrequent praiser is more informative. We addressed this alternative explanation in Experiment 2 with adults (Exp. 2A) and 4 year-olds (Exp. 2B).

Experiment 2

In Exp. 2, we asked whether participants distinguish between teachers who appropriately provided selective feedback (positive to the good tracings and neutral to the bad tracings; Selective Teacher in Exp.1) from a teacher who provided the reverse (positive to the bad tracings and neutral to the good ones; Selective-Incongruent Teacher). Because the frequency and the valence of their feedback was matched, the critical difference was whether their feedback was congruent or incongruent with the actual quality of the drawings. If participants track the covariance between the content of feedback and the quality of the tracings rather than the frequency of praise, they would trust the Selective Teacher over the Selective-Incongruent Teacher. Given the absence of an age trend in Exp. 1B, we limited our recruitment to 4-year-olds.

Experiment 2A: Adults

Methods

Participants Ninety-one adults (37 female, M_{Age}(SD) = 34.7(10.3), range: 19-65) were recruited from Amazon's Mechanical Turk. An additional 19 subjects were excluded for failing one or both memory check questions, or not completing the task.

Materials Stimuli were identical to Exp.1A except that the Overpraise Teacher video was replaced with the Selective-Incongruent Teacher video. The Selective-Incongruent Teacher was similar to the Selective Teacher but provided praise in the opposite way: She praised the bad tracings and gave a neutral response to the good tracings.

Procedure The procedure was identical to Exp. 1A.

Results and Discussion

As predicted, participants chose the tracing praised by the Selective Teacher (90.1%, $p < .001$, Binomial Test). When asked which teacher was trying to be nice, the majority of participants chose the Selective-Incongruent Teacher (72.5%, $p < .001$, Binomial Test). These results suggest that adults specifically use the congruency of the teacher's praise with the quality of the tracing to evaluate which of two hidden products was likely to be better.

Experiment 2B: 4-5 year-olds

Methods

Participants Twenty-four 4-year-olds (15 female, M_{Age}(SD) = 4.6(0.3), range = 4.0 - 4.9) were recruited from a university preschool. An additional 5 subjects were tested but excluded due to failure on the warm-up or memory check questions.

Stimuli and Design Stimuli were similar to Exp. 1B, except the videos of the Selective Teacher and Selective-Incongruent (as in Exp. 2A) were used.

Procedure The procedure was identical to Exp. 1B.

Results and Discussion

Our main question was whether children would endorse the praise from the teacher who had previously given more informative praise even when the overall valence and the relative frequency of the two teachers' praise was matched. Indeed, children were more likely to choose the tracing praised by the Selective Teacher than the one praised by the Selective-Incongruent Teacher ($75\%, p = 0.02$, Binomial Test). The same logistic regression model as in Exp. 1B (`Tracing choice ~ Age in Months + (1 | Subject)`) revealed an effect of age (B = -6.36, z = -2.11, p = 0.04).

These results suggest that children are not simply responding to the relative frequency of praise; they are sensitive to whether the teacher provides feedback that appropriately co-varies with the actual quality of tracings. Unlike Exp. 1B (and despite the narrower age range), we did find a small but significant effect of age. Given the absence of an age effect in Exp. 1B, further work is needed to better understand how this capacity develops throughout the early childhood years.

General Discussion

Praise is a useful source of information for learning about our skills and abilities; determining whose praise to trust or discount is critical for effectively learning about the self. Across two experiments, we examined whether adults and children infer the informativeness of others' praise and use it to evaluate their own or others' work when they themselves are unable to judge. Exp. 1A verified that adults reliably distinguish a teacher who selectively praised from a teacher who indiscriminately overpraised. Exp. 1B provided support for our main hypothesis that even 4-5 year-old children distinguish between these two teachers and use their relative informativeness to evaluate the quality of their own work. Exp. 2 addressed the alternative explanation that these results were driven by the simple heuristic that people who rarely praise are more informative; when the frequency and the overall valence of praise were matched, adults and 4-year-olds still endorsed the praise from the teacher whose feedback was both selective and congruent with reality.

These results are consistent with a growing body of literature that suggests that children reason about others' informativeness based on the information they provide (Sobel & Kushnir, 2013; Gweon & Asaba, 2017). Going beyond using facts about the physical world (e.g., labels of objects, causal functions of artifacts), children also used information about the quality of work (e.g., quality of tracings) to decide whether or not to trust someone's evaluative feedback. More specifically, children's inferences were based on the statistical dependence between the pattern of praise and the quality of the work being praised. Thus, children's early-emerging

sensitivity to statistical information (Gweon & Schulz, 2011) might also support inferences about others' informativeness.

Just as prior knowledge is critical to evaluate others' testimony about the external world, the ability to independently assess the quality of work was critical for success in our task; if you do not know how good a tracing is, you cannot tell whether a teacher's feedback is appropriate. We used clearly good or bad tracings and used a warm-up task to ensure that all children could accurately assess their quality. Yet, even adults rarely have absolute certainty about the quality of their own work. When do children rely on their own evaluations to infer others' informativeness, and when do they rely on others' evaluative feedback to inform their own evaluations? The current work is a first step towards understanding how children integrate their own certainty with others' informativeness to jointly learn about themselves (i.e., did I do well?) and others (i.e., is this person informative?).

In the current study, children observed repeated instances of praise from two teachers whose praise was clearly aligned or misaligned with the quality of someone else's tracings. In order to ensure well-controlled presentation of two teachers who varied only in their pattern of praise, children watched videos of teachers instead of seeing real teachers praise their own work. Prior work suggests that children tend to be optimistic about their own competence, judging that they have performed well when they have not (Hembacher & Ghetti, 2014), and predicting that they will perform better in the future than they are able to (Schneider, 1998). Thus, an open question is whether this bias would influence their ability to infer the informativeness of others' praise from evaluations directed at their own work (rather than another student). Nevertheless, our initial results do show that such optimism did not completely overpower their ability to use others' feedback to assess their own work; faced with uncertainty about which one of their own tracings is better, children relied on the teacher whose praise is more likely to be informative.

In real life, however, children's observations of evaluative feedback is often noisier, and unfolds in a complex social environment where both the evaluators and the students must navigate multiple competing goals. While our results suggest an early-emerging sensitivity to the informativeness of praise, further work is needed to better understand how such sensitivity might manifest in real-world contexts. Furthermore, children's experience with praise might vary significantly depending on their culture, educational context, and family environment. Our participants were from a university preschool where children come from various cultural backgrounds yet receive ample social support from adults. Thus an important question is how our findings might generalize to a broader population of young children who might experience varying levels of praise, encouragement, and support.

Note that the current study focused on how children infer the informativeness of feedback on one's *performance*. In this context, "informative" feedback meant that it appropriately reflected the quality of the work. However, "infor-

mative" feedback may not always be beneficial. Prior work suggests that praise directed at children's abilities or intelligence (i.e., person-praise) is detrimental when it leads them to take their performance as an indicator of underlying competence and strengthens their belief that competence is static and unchangeable. Thus, it is possible that the negative consequences of person-directed praise are heightened when it comes from a teacher who praises selectively (e.g., a teacher who selectively tells a few students that they are really smart).

Finally, our study emphasized the learner's goal to seek informative feedback: They had to choose a tracing to enter into a contest (Adults) or show their teacher (4-5 year-olds). Yet, just as evaluators may have competing social goals to be informative or nice (Yoon et al., 2016), learners may have diverse goals in approaching others for feedback; they might want honest, informative evaluation of their performance, or warm, encouraging affirmation to *feel better*. An open question is whether young children differentially weight praise based on their own goals, and how this tendency might change with age. It is possible that younger children generally seek more affirmation than evaluation, but their goal might also vary depending on their competence in the domain and the relative difficulty of the task. For instance, learners might prefer encouragement when they are struggling on new or difficult tasks. The kind of feedback children seek might also depend on their relationship to the evaluator, desiring more affirmation from parents and expecting more objective evaluations from teachers. Future work might ask how children actively choose which teacher to approach depending on their goals (e.g., informativeness versus affirmation).

Constructive feedback provides insights into learners' strengths and weaknesses, and guides their future learning to maximize opportunities for growth. Our results suggest that the ability to seek constructive feedback might start early in life. Even for young children, praise is more than something nice; they track the informativeness of others' evaluative feedback and use it to infer the quality of their own work.

Acknowledgements

All data and code for these analyses are available at `https://github.com/masaba/uncertainpraise`

References

Boseovski, J. J., Marble, K. E., & Hughes, C. (2017). Role of expertise, consensus, and informational valence in children's performance judgments. *Social Development*, *26*(3), 445–465.

Brophy, J. (1981). Teacher praise: A functional analysis. *Review of educational research*, *51*(1), 5–32.

Delin, C. R., & Baumeister, R. F. (1994). Praise: More than just social reinforcement. *Journal for the theory of social behaviour*, *24*(3), 219–241.

Gweon, H., & Asaba, M. (2017). Order matters: Children's evaluation of underinformative teachers depends on context. *Child Development*.

Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when teachers are under-informative. *Cognition*, *132*(3), 335–341.

Gweon, H., & Schulz, L. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, *332*(6037), 1524–1524.

Hembacher, E., & Ghetti, S. (2014). Dont look at my answer: Subjective uncertainty underlies preschoolers exclusion of their least accurate memories. *Psychological Science*, *25*(9), 1768–1776.

Henderlong Corpus, J., & Lepper, M. R. (2007). The effects of person versus performance praise on childrens motivation: Gender and age as moderating factors. *Educational psychology*, *27*(4), 487–508.

Landrum, A. R., Mills, C. M., & Johnston, A. M. (2013). When do children trust the expert? benevolence information influences children's trust more than expertise. *Developmental Science*, *16*(4), 622–638.

Lane, J. D., Wellman, H. M., & Gelman, S. A. (2013). Informants' traits weigh heavily in young children's trust in testimony and in their epistemic inferences. *Child Development*, *84*(4), 1253–1268.

Mueller, C. M., & Dweck, C. S. (1998). Praise for intelligence can undermine children's motivation and performance. *Journal of personality and social psychology*, *75*(1), 33.

Ronfard, S., & Lane, J. D. (2017). Preschoolers continually adjust their epistemic trust based on an informant's ongoing accuracy. *Child development*.

Schneider, W. (1998). Performance prediction in young children: Effects of skill, metacognition and wishful thinking. *Developmental Science*, *1*(2), 291–297.

Seiver, E., Gopnik, A., & Goodman, N. D. (2013). Did she jump because she was the big sister or because the trampoline was safe? causal inference and the development of social attribution. *Child development*, *84*(2), 443–454.

Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, *120*(4), 779.

Yoon, E. J., Tessler, M. H., Goodman, N. D., & Frank, M. C. (2016). Talking with tact: Polite language as a balance between kindness and informativity. In *Proceedings of the 38th annual conference of the cognitive science society*.