Children, but not adults, prioritize relational over dispositional interpretations of dominance interactions

Pomiechowska, Barbara; Tatone, Denis; Mészégető, Dorottya; Revencu, Barbu; Csibra, Gergely

Title
Children, but not adults, prioritize relational over dispositional interpretations of dominance interactions

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors
Pomiechowska, Barbara
Tatone, Denis
Meszegeto, Dorottya
et al.

Publication Date
2022

Peer reviewed

Children, but not adults, prioritize relational over dispositional interpretations of dominance interactions

Barbara Pomiechowska*[1] (pomiechowskab@ceu.edu), Denis Tatone*[1] (tatoned@ceu.edu), Dorottya Mészégető[1] (meszegetod@ceu.edu), Barbu Revencu[1] (Revencu_Barbu@phd.ceu.edu), Gergely Csibra[1,2] (csibrag@ceu.edu)

[1] Department of Cognitive Science, Central European University, Quellenstrasse 51, 1100 Wien, Austria
[2] Department of Psychological Science, Birkbeck, University of London, Malet Street, London WC1E 7HX, UK

* denotes joint first authorship

Abstract

Humans routinely monitor social interactions to learn about the relational make-up of their groups and select social partners. It is unclear however whether social interactions primarily invite inferences about the dispositions of the participants involved or about underlying social relations. In the present study we tested which of these two inferences children and adults draw when observing interactions based on dominance. Children expected dominants to prevail over previous subordinates but did not generalize this expectation to interactions with novel agents, whereas adults did. These results suggest that children interpreted dominance as specific to a particular social relation, whereas adults interpreted it as a stable, target-invariant trait. This asymmetry supports the proposal that children may first interpret social interactions through a relational stance, and only later in development apprehend them through the lenses of trait attribution.

Keywords: cognitive development; trait attribution; social dominance; naïve sociology

Introduction

Mapping the social terrain in terms of its constituent relations is a task of utmost adaptive significance for gregarious species, such as humans. A prominent source of information that people exploit to this end is the monitoring of social interactions from a third-party perspective. However, the observation of social interactions (e.g., Malvin offered food to Joe) minimally invites two distinct types of inference: *relational* inferences, which appeal to underlying relations to explain the distribution of observed behaviors between the individuals involved (e.g., Malvin and Joe are friends); and *dispositional* inferences, which appeal to individual traits motivating the observed behavior in the focal individual (e.g., Malvin is generous). Crucially, these two types of inferences support different assumptions about the generalizability of social behaviors: if relational inferences restrict the occurrence of a social behavior to encounters with the same partner, dispositional inferences license its generalization to novel partners (Repacholi et al. 2016).

Which of these inferences do children and adults prioritize when observing third-party interactions? The developmental literature does not provide a clear answer. On one hand, studies on early social evaluation based on the manual-choice task (e.g., Hamlin et al., 2007) produced evidence indirectly compatible with a dispositional interpretation: the infants' preference for prosocial agents may be supported by a "global assessment" of the agents' behavior in terms of sociomoral dispositions generalizable to uninvolved individuals such as the infants themselves (Wynn, 2008). On the other hand, studies that used looking-time tasks to more directly test across-target generalization failed to find evidence of dispositional construals: e.g., after having seen an agent A prevailing over (Mascaro & Csibra, 2012) or giving an object to (Tatone et al., 2015) another agent B infants did not expect A to perform the same action to novel targets (cf. Surian et al., 2018).

Preschoolers also struggle with predicting behavioral consistency from single exemplars (Boseovski & Lee, 2006) as well as with inferring dispositions from trait labels (Liu et al., 2007). It is only around the age of seven that children begin to use traits to consistently explain others' behaviors (Rosati et al., 2001; Kalish, 2002). It should be however noted that previous studies used vignette-based tasks which heavily rely on verbal descriptions. The late-emerging use of trait explanation may thus reflect the correspondingly late emergence of linguistically transmitted "trait-like schemas" (Seiver et al., 2012), and may be masking an earlier preverbal ability to reason about traits in younger children. Furthermore, studies typically contrasted construals based on traits with alternatives based on prevailing norms (Kalish & Shiverick, 2004; Rholes & Ruble, 1984), overlooking the possibility that social interactions may instead prompt inferences about the social relations that these are instance of.

The present study aims to assess the relative influence of relational and dispositional inferences on the interpretation of third-party interactions in children and adults. To overcome the above concerns, we devised a task that involved minimal linguistic scaffolding, and which directly compared relational and dispositional construals of social interactions via a target-generalization measure. Participants were presented with animations without any accompanying verbal narration. The animations depicted interactions over a contested resource resolving in nonviolent priority of access – an event that has been shown to induce the representation of social dominance already in preverbal infants (e.g., Mascaro & Csibra, 2012). Instead of being asked to predict or explain the agent's behavior, participants had to produce postdictive inferences about the resolution of the events observed. Participants were first shown a resource contest featuring two agents (A and B) that always ended with one (A) seizing the resource. At test, they were then shown a new contest event whose outcome was not visible (involving A and B again or one them with a novel agent C), and asked which of the two agents took the

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

resource. A relational interpretation of these events should lead participants to assume power asymmetries to be dyad-specific, and thus select the dominant (A) as the agent likely to prevail only when paired with the previous partner (B). Conversely, a dispositional interpretation should lead participants to generalize the behavior (to prevail, to yield, or both) with novel partners (C) (see Figure 1B).

We compared the performance of adults, who have been reported to infer traits even from single exemplars (e.g., Reeder et al., 2004; Uhlmann et al., 2014), to that of four-year-olds. At this age, children already draw inferences about social power from the outcome of resource-contest events (Gulgöz & Gelman, 2017), but cannot yet consistently predict how an agent will behave towards new partners in vignette-based tasks (Rholes & Ruble, 1984). This is thus an ideal age to assess whether the children's reluctance to infer traits may be due to the language-heavy and predictive nature of previous tasks or whether it instead reflects a more general propensity to interpret interactions as indicative of particular social relations.

Experiment 1

Experiment 1 tested whether children interpreted the resolution of a resource contest between two agents (A and B) in favor of A as indicative of a dominance relation ("A dominates B": relational inference) or of an individual trait ("A is a dominant": dispositional inference). After watching agent A seize a resource coveted by both agents, children were shown either these two agents (A and B: same-partners condition), or one of them with a novel agent (A and C: new-subordinate condition; B and C: new-dominant condition) in an identical resource-contest event with non-visible outcome, and were asked to indicate who took the resource. If children interpret the event in relational terms ("A dominates B"), they should pick agent A when paired with B, while choosing at random when either agent is paired with a novel character. Conversely, if they construe the agents' behavior in dispositional and partner-invariant terms, children should produce one of three possible choice patterns for the new-agent pairs: picking both A and C, if they infer two distinct dispositions to prevail and yield, respectively; and picking only A or C, if they infer a disposition to either prevail or yield.

Participants. The final sample consisted of 60 four-year-olds (range: 4.0 to 4.11 years). An additional 14 children were tested but excluded from the final analysis (no video recording of the session, n = 1; side bias, n = 3; insufficient number of valid trials, n = 5; parental interference, n = 2; incorrect answers on warm-up trials, n = 2; experimenter error: n = 1). Families were recruited from the laboratory's database and via social media. Caregivers provided verbal informed consent for participation and video recording of the session. Children provided verbal assent to participate.

Procedure. Children were tested online during a Zoom call. They were shown animations on their caregiver's computer screen and asked test questions by the experimenter. The session was video-recorded. Children's responses were coded offline.

Stimuli The stimuli were 3D animations made with Blender (https://www.blender.org) and modelled on Mascaro and

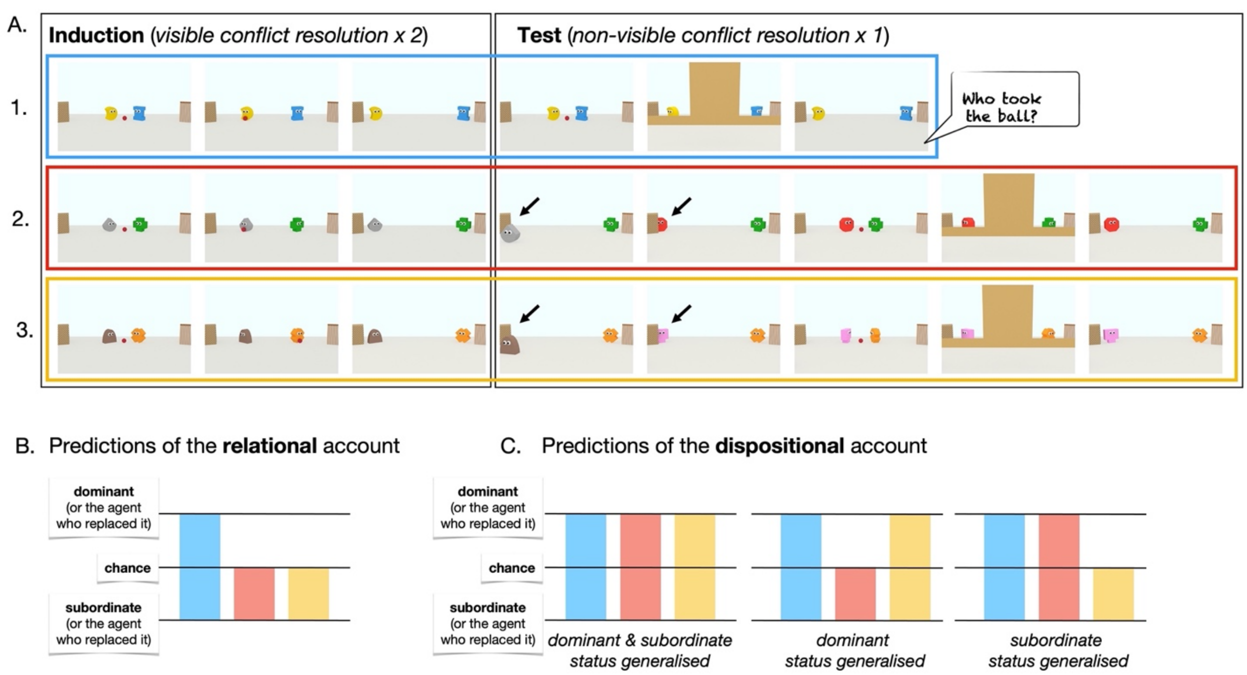

Figure 1: (A) Schematic visualization of a single experimental trial. Numbers 1-3 denote experimental conditions (1: same partners; 2: new dominant; 3: new subordinate). The arrows indicate the departure of one of the familiarized agents and the arrival of the new agent at test. (B-C) Predictions of the two accounts proposed. Values on the Y-axis indicate the probability of choosing the dominant (or the agent replacing it).

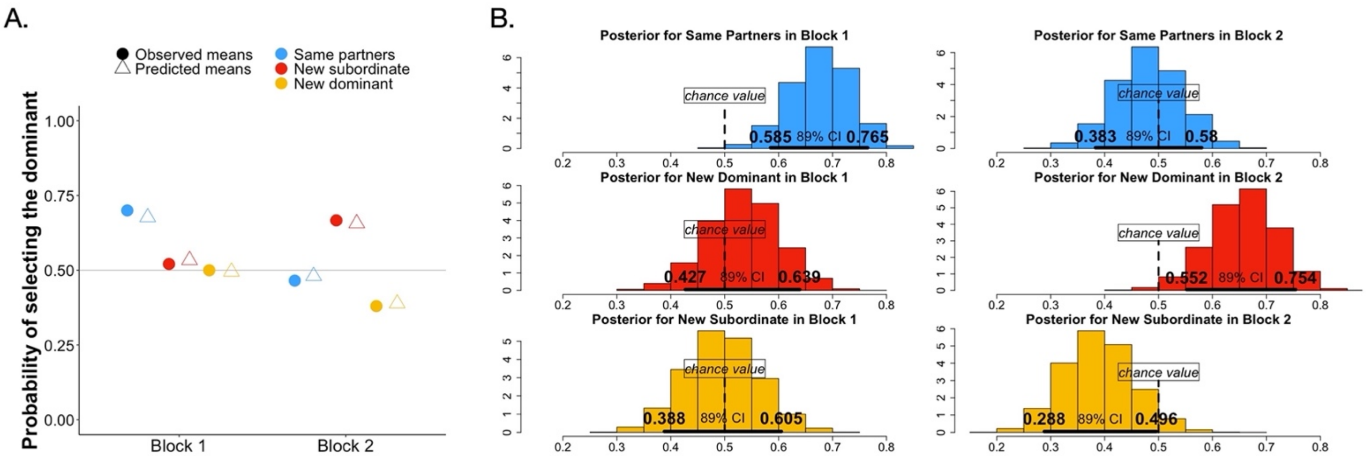

Figure 2 : Results of Experiment 1. (A) The observed probability of selecting the dominant and the model's predications. (B) Posterior distributions for probability of selecting the dominant agent. Thick black horizontal lines give the 89% credible interval around the mean. The chance value (.5) is marked with a vertical, dotted line.

Csibra (2012). Colored geometric shapes with eyes served as agents. There were three types of events. All contained the same display elements: two light brown patterned boxes (one striped, one dotted) located at the opposite sides of a stage; two agents varying in color and shape, each located nearby a different box; a ball (*Figure 1B*).

Conflict events involved two agents set to fulfill the same goal (i.e., collect a ball) and only one of them succeeding. No physical conflict was involved. The agents first observed a ball falling in the middle of the stage, which they then simultaneously approached, until one seized it while the other remained still. At this point, the agents moved behind their respective boxes and re-emerged empty-handed.

Test events unfolded in the same way as conflict events, but when the agents approached the middle of the stage, an occluder covered the scene, concealing the conflict outcome. The occluder moved away only when the agents remerged from behind the boxes.

Depending on the condition, the test events featured either the same two agents shown in the conflict events, or only one. In the latter case, before the test participants would see one agent moving out of the stage and another featurally distinct agent coming from behind the same box to replace her (see Figure 1A).

Different looking agents were presented on each trial. Full list of agents' colors and shapes used (as well as their counterbalancing) is available in our OSF repository.

Design. The experimental session consisted of a warm-up and an experimental task.

Warm-up. First, children participated in a short *color-naming game*. The aim was to assess whether they could name all the colors used in the experimental test trials (n = 6; yellow, green, red, blue, orange, pink), and determine which color labels they would provide. Children were shown two displays, each containing three animals of different colors (display 1: a yellow fish, a red dinosaur, and a blue dog; display 2: an orange turtle, a green snail and a pink owl) and were asked about each animal's color (e.g., *"Can you tell me what color the dinosaur is?"*).

Then, children were introduced to the new creatures called "tegies" through a series of short animations narrated by the experimenter. She began by labeling the agents arriving one by one on the stage ("*Now we will meet some tegies. Look, here is one. And another one, and a few more are arriving.*"); then, she emphasized that the agents have different colors and shapes (*"There are many different tegies, see? They have different colors and shapes."*), but they all have the same favorite toy, a ball (*"Do you know what is the tegies' favorite toy? Balls! Their favorite toys are balls.*"). Four different tegis were then shown one by one collecting balls (*"Look at this lucky tegi! He found a ball, his favorite toy.")*. At the end, the experimenter remarked again about the tegies' preference for balls (*"All tegies' favorite toys are balls. Tegies love balls."*).

After the introduction, children were administered two working-memory trials to assess whether they could track which agent prevailed in a conflict event. These trials had the same structure as the induction events used in the experimental task. After the event ended, children were asked who took the ball ("Do you remember who took the ball?"). If children did not provide any answer, they were asked the test question again, with the agents' colors explicitly named (e.g., "What do you think, which one took the ball, the blue one or the yellow one?"). Participants were given verbal feedback. Each trial featured a different pair of agents.

Experimental task. Each experimental trial involved an induction, consisting of two conflict events with visible resolution, immediately followed by one test event with non-visible resolution. As soon as the test event ended, children were asked *"Who took the ball?"*. If they did not provide any answer, they were asked the test question again, with the agents' colors explicitly named (e.g., "*What do you think, which one took the ball, the blue one or the yellow one?"*). The test responses received no feedback.

There were three within-subject test conditions that differed only with respect to the identity of the agents present at test: (1) in the **same-partners condition**, the same agents were present throughout the trial; (2) in the **new-dominant** condition the agent who previously prevailed (seized the ball) was replaced by a new agent; (3) in the **new-subordinate** condition the agent who previously yielded (did not seize the ball) was replaced by a new agent. In both new-partner conditions the test question was followed by an additional memory question: *"Who is new?"* (followed *"Who arrived later*?" in case the first question elicited no answer). The aim of this question was to determine whether participants detected the presence of a new agent.

Children received 2 blocks of 3 trials (1 per condition), for a total of 6 trials. Side of the dominant agent and order of trial presentation were counterbalanced within subjects, whereas the color-shape mappings between dominant and subordinate agents were counterbalanced across subjects.

Inclusion criteria. Children had to fulfill the following preregistered criteria to be included in the final sample. They had to (1) label all colors in the color-naming game (mislabeling was accepted unless the child confounded two colors co-occurring in the same test); (2) answer correctly to both warm-up memory questions; (3) contribute a minimum of one valid experimental trial per condition; and (4) not display a side bias (i.e., defined as picking the same side in all test questions). A trial was considered valid only if the child answered by naming the color or the shape of an agent, and additionally, in new partners conditions, only if she correctly answered the memory question following the test.

Coding. The children's choices were scored to produce a *domgenscore* (i.e., the probability of choosing the dominant agent), as follows: choosing the agent who prevailed or its substitute was scored as 1; choosing the agent who yielded or its substitute was scored as 0. A *domgenscore* of 1 would thus indicate that participants selected the dominant agent or its substitute.

Results and discussion

Auxiliary analyses. Children had no problems naming colors or indicating who took the ball when the conflict resolution was visible (70/74 tested children answer correctly to both warm-up questions). They also performed well on the memory questions after the new-partner test trials (32/60 children included in the final sample answered correctly to all memory questions) which shows that they were able to track the identity of the agents throughout the events.

Main analysis. Children's choices at test were analyzed by a Bayesian logistic regression family of models in *rstan* 2.21.2, using the *rethinking* package 2.13 (McElreath, 2020). The code is available in our OSF repository. We modeled each observed choice as a Bernoulli random process with parameter p_i, the probability of choosing the dominant agent or its substitute (i.e., *domgenscore)* in each trial i. We then modeled the log-odds of each choice, μ_i, as a function of the condition (same partners v. new dominant v. new subordinate), block (1 v. 2), condition-block interaction, and/or participant generating that choice.

Priors and assumptions. The priors for these parameters were chosen based on prior predictive checks to satisfy three criteria. First, given the novelty of the paradigm, we allowed the models to discover a wide range of possible *domgenscores*, so we used an uninformative prior for the overall baseline. Second, we avoided biasing the models to find an effect in any direction, so we centered the priors on the effects (condition and block) on 0. Third, we allowed the models to explore a wide range of possible effects while imposing some prior skepticism towards extreme effect sizes.

Modeling. First, we obtained the best-fitting model by comparing three different models. The first model (*m1*) included only subject and condition as predictors: $\mu = \alpha + \beta_{ID} + \beta_{condition}$. The second model (*m2*) included an additive block term: $\mu = \alpha + \beta_{ID} + \beta_{condition} + \beta_{block}$. In the third model (*m3*), we added an interactive block term: $\mu = \alpha + \beta_{ID} + \beta_{condition} + \beta_{conditionblock} \cdot (block - 1)$. Using the Widely Applicable Information Criterion (WBIC: Watanabe, 2010; McElreath, 2020), we found that the interaction model (*m3*) obtained 85% of the total weight and thus accounted

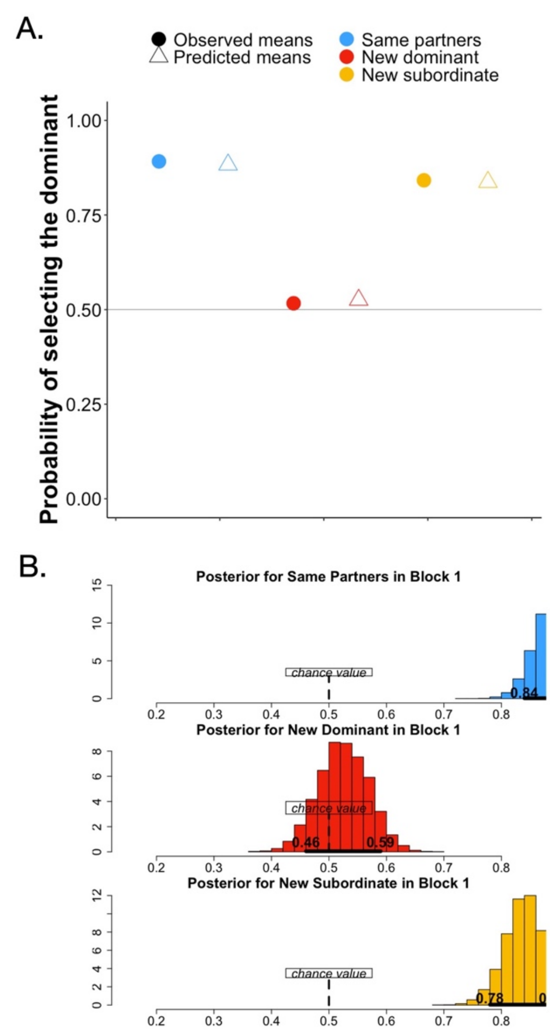

Figure 3: Results of Experiment 2. (A) The observed probability of selecting the dominant and the model's predictions. (B) Posterior distributions for probability of selecting the dominant agent. Thick black horizontal lines give the 89% credible interval around the mean. The chance value (.5) is marked with a vertical, dotted line.

best for the children's data patterns. This indicated that children's choice patterns changed across blocks.

Having found the best-fitting model, we examined the posterior estimates for children's choice probabilities outputted by *m3* across conditions and blocks. Figure 2 shows that in Block 1, in the same partners condition children consistently selected the dominant agent, as indicated by the fact that 89% of distribution of *domgenscore* lied within .59 and .77. In contrast, in both new partners conditions, the 89% credible interval comprised .5, suggesting that children chose randomly between the agents present at test (new dominant: 89% CI = [.43, .64], new subordinate: 89% CI = [.39, .61]). To investigate whether the choices differed between conditions, we computed differences between posteriors across conditions. When the new partners conditions was subtracted from the same partners condition, 89% of the difference distribution lied above the null value (*Same Partners – New Dominant:* 89% CI = [.003, .28]; *Same Partners – New Subordinate:* 89% CI = [.04, .32], thus suggesting that children were more likely to choose the dominant agent only when it faced the same partner as before and more likely to do so than to choose the agent who replaced it. The 89% credible interval for the difference between *New Dominant – New Subordinate* included 0, 89% CI = [-.11, .19], thus indicating that children behaved comparably across both new partners conditions.

In sum, in Block 1, children tracked which agent prevailed in the induction phase and used this knowledge to infer the likely dominant at test. However, when either of the familiar agents were paired with new partners, children chose at random. The lack of consistent choice in the new-partner trials suggests that children did not interpret the resolution of the conflict events as indicative of a partner-invariant trait (to prevail or yield), supporting instead an interpretation of these events as indicative of a dyad-specific dominance relation.

In block 2, however, children did not consistently pick the dominant agent in the same-partners conditions, but chose at random, 89% CI = [.38, .58]. Furthermore, in both new-partner conditions, children consistently chose the new agent (i.e., replacing the previous dominant: 89% CI = [.55, .75]; replacing the previous subordinate: 89% CI = [.29, .50]). Children were more likely to select a novel agent in the new-partner conditions than in the same-partners condition (*Same Partners – New Dominant:* 89% CI = [-.32, -.03]; *Same Partners – New Subordinate:* 89% CI = [-.05, .23]). This pattern of choices does not align to the predictions of either the relational or the dispositional account. Rather, it is consistent with a preference for the novel agent, which developed over the course of the task. We assume that this novelty bias may have been induced by the memory question ("Who is new?") administered immediately after the test question in the new-partner trials. Future experiments should explicitly control for this possibility by manipulating the experimenter's questions.

Experiment 2

Experiment 1 showed that 4-year-olds inferred a dominance relation from the outcome of a contest event and did not generalize the disposition to prevail or yield to interactions with novel agents. These findings suggests that children prioritized a relational over a dispositional construal of the observed interactions. Experiment 2 examined which of these interpretations adults adopt using an adapted version of the task used with children.

Methods

Participants. Participants (N = 64) were recruited from the Testable Minds pool (https://minds.testable.org). The final sample consisted of 60 participants (age: M = 30 years; *R:* 19 to 43 years). All participants provided written informed consent before the testing session.

Stimuli, Design, and Procedure. The experiment was conducted online using Testable (https://www.testable.org). The stimuli and experimental design were the same as in Experiment 1, but we modified the testing procedure such that the task could be administered without supervision and made it appropriate to an adult sample. All instructions and test questions were presented in a written form, and responses were collected through key presses. Further, we removed the color-naming game and introduced the goal object as a food item (rather than a toy) to make the stimuli more ecologically valid for adults.

Inclusion criteria. As in Experiment 1, to be included in the final sample, participants had to (1) answer correctly to both warm-up memory questions; (2) contribute a minimum of one trial per condition. Trials in new-partner conditions were excluded if the memory question after the test was answered incorrectly.

Results and discussion

We used the same analysis protocol as in Experiment 1. The condition-only model (*m1:* $\mu = \alpha + \beta_{ID} + \beta_{condition}$) had 60% of the model weight, indicating that the adults' responses were consistent across blocks. The same priors were used as in Experiment 1.

The results are depicted in Figure 3. The posterior estimates of choice revealed that adults consistently selected the dominant agent in the same partners condition, with the 89% of the predicted *domgenscore* distribution falling between .84 and .91, and the new subordinate condition, with the 89% of the predicted *domgenscore* distribution falling between .78 and .89. Conversely, in the new dominant condition, participants had no preference for either of the agents, the subordinate or the new agent who replaced the dominant, 89% CI = [.46, .60]. Further, they were more likely to choose the dominant agent in the same partners and new subordinate conditions than its substitute in the new dominant condition, as indicated by the differences between conditions not including the null value, i.e., 0 (*Same Partners – New Dominant,* 89% CI = [.27, .44]; *New Subordinate – New Dominant,* 89% CI = [-.39, -.22]).

Unlike children, adults not only selected the previously prevailing agent when paired with the previous partner, but also with a new agent. These findings suggests that adults interpreted the outcome of the interaction as indicative of a partner-invariant disposition to prevail over, but not to yield to, others.

General Discussion

We sought to examine what type of construal (relational vs. dispositional) children and adults adopt when interpreting third-party interactions. Across two experiments, we found that both children and adults inferred dominance from the outcome of a contest event: they consistently chose the dominant agent as more likely to prevail in a contest event with no visible resolution when paired with the previous partner (subordinate). When the event involved a new partner, however, the choice pattern differed across the two age groups: adults selected the previous dominant, but children did not. The children's lack of generalization is consistent with an interpretation of the interaction as indicative of a dominance relation, which restricts inferences about power asymmetries to the observed dyad. Conversely, the tendency to generalize dominance to interactions with novel agents, as observed in adults, suggests an interpretation of the dominant's behavior as manifesting a target-invariant trait (to prevail).

Taken together, these results suggest that dispositional interpretations of social interactions may not emerge until later in development. This conclusion is largely congruent with the literature on trait attribution reviewed earlier and extends the reluctance to deploy dispositional inferences to tasks involving minimal linguistic mediation and postdictive (rather than predictive) questions. This suggests that neither the language-heavy nor the predictive nature of the measures used in earlier studies can adequately account for the children's failure to generalize behaviors across targets (Boseovski & Lee, 2006).

While the adults' response can be taken as evidence of a trait-like interpretation of the tendency to prevail over others, it may also reflect the operations of latent structural schemata (e.g., a pyramidal configuration), which presuppose a small set of dominants ruling over a larger set of subordinates. Recruiting such representational schemata would make adults consider novel partners more likely to be sampled from the bottom of the structure than the top. In such case, the dominance generalization, rather than being evidence of trait attribution, would instead be consistent with a relational account which incorporates implicit priors about the shape of the larger hierarchical structures that the inferred relations are embedded in (Fiske, 2012).

Irrespective of which interpretation may best account for the adult data, the present results corroborate the hypothesis that children may first adopt the interpretive stance of naïve sociologists, charting out the social landscape in terms of its constitutive relations (Thomsen & Carey, 2013), and only later develop the stance of naïve personality psychologists, biasing the interpretation of social behaviors as indicative of stable individual dispositions (as famously captured in the "fundamental attribution error", FAE: Ross, 1977).

What could explain such a developmental progression? It has been argued that inductive biases of the dispositional kind (like the FAE) may serve to enhance one's reputation (i.e., by convincing others that our most socially praiseworthy actions stem from personal virtues rather than social obligations: Andrews, 2001) or to shape social behaviors (e.g., by making character-wide ascriptions work as self-fulfilling prophecies about one's socially expected conduit: Alfano, 2013). Under such account, the adoption of trait-based explanations should become a robust phenomenon only once children have begun to appreciate the reputational consequences of their (and others') interpersonal conducts (Silver & Shaw, 2018).

There are however conceivable alternative interpretations of the present findings, especially with regards to the adult data, which should be addressed. First, the adults' response pattern, which we argued to be evidence of trait attribution, may have been supported by a non-interactive interpretation of the events based on relative differences in desire strength. The evidence that one agent (the dominant) consistently acquired the object whereas the other (the subordinate) did not may have induced adults to ascribe a desire for the item only in the former, and thus to reason that the dominant agent should acquire the item again when paired with a familiar or novel agent. To control for this possibility, we plan to run a new experiment with a modified induction phase where both agents are shown acquiring the object an equal number of times prior to the conflict events. This account however does not detract from our main conclusion that children preferred pair-based construals while adults opted for actor-based ones (regardless of whether these were based on desire or trait ascription). Further, we cannot rule out the possibility that the adults' generalization pattern may have stemmed from interpreting the test question ("Who took the ball?") as referring to the previously witnessed conflict event, in which A dominated B. To control for such a confound, we will run a replication in which the temporal reference of the question is made unambiguous ("Which of the two agents behind the screen took the ball?").

On a related note, one may suggest that adults generalized "dominance" because, unlike children, they focused only on the agent who acquired the object, thus failing to encode the partner's identity. This account is however question-begging: why would adults, who unquestionably have more attentional resources to recruit for event processing and have been shown to automatically encode asymmetric thematic roles (Hafri et al., 2018) from rapid scene exposure, omit representing an event participant that even infants readily encode (Mascaro & Csibra, 2012)?

More generally, it should be emphasized that the present study used only one type of interaction (dominance), and one type of diagnostic cue (priority of resource access). Given the manifold ways in which humans relate to each other, it would be premature to extrapolate from these findings which types of inferences children and adults preferentially draw when presented with other types of social behaviors (e.g., helping, imitating, etc.). Future studies employing a broader range of interaction events will contribute to shed light on this issue.

Preregistration & Materials

The preregistration can be accessed at the following link: https://osf.io/eqjch The code and materials are available in the project's OSF repository: https://osf.io/rd36z/

Acknowledgements

We thank the families who participated in this research. We thank Maria Tóth and Agnes Volein for research assistance.

Funding

This research was partly funded by the European Research Council under ERC PARTNERS (Grant agreement ID: 742231).

References

Alfano, M. (2013). *Character as moral fiction*. Cambridge University Press.

Andrews, P. W. (2001). The psychology of social chess and the evolution of attribution mechanisms: Explaining the fundamental attribution error. *Evolution and Human Behavior*, *22*(1), 11-29.

Boseovski, J. J., & Lee, K. (2006). Children's use of frequency information for trait categorization and behavioral prediction. *Developmental psychology*, *42*(3), 500-513.

Fiske, A. P. (2012). Metarelational models: Configurations of social relationships. *European Journal of Social Psychology*, *42*(1), 2-18.

Kalish, C. W. (2002). Children's predictions of consistency in people's actions. *Cognition*, *84*(3), 237-265.

Kalish, C. W., & Shiverick, S. M. (2004). Children's reasoning about norms and traits as motives for behavior. *Cognitive Development*, *19*(3), 401-416.

Hafri, A., Trueswell, J. C., & Strickland, B. (2018). Encoding of event roles from visual scenes is rapid, spontaneous, and interacts with higher-level visual processing. *Cognition, 175,* 36-52.

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

Gülgöz, S., & Gelman, S. A. (2017). Who's the boss? Concepts of social power across development. *Child Development*, *88*(3), 946-963.

Liu, D., Gelman, S. A., & Wellman, H. M. (2007). Components of young children's trait understanding: Behavior-to-trait inferences and trait-to-behavior predictions. *Child development*, *78*(5), 1543-1558.

Mascaro, O., & Csibra, G. (2012). Representation of stable social dominance relations by human infants. *Proceedings of the National Academy of Sciences, 109*(18), 6862-6867.

Mascaro, O., & Csibra, G. (2014). Human infants' learning of social structures the case of dominance hierarchy. *Psychological science, 25*(1), 250-255.

McElreath, R. (2020). Statistical Rethinking: A Bayesian Course with Examples in R and Stan. CRC Press.

Reeder, G. D., Vonk, R., Ronk, M. J., Ham, J., & Lawrence, M. (2004). Dispositional attribution: multiple inferences about motive-related traits. *Journal of personality and social psychology*, *86*(4), 530.

Repacholi, B. M., Meltzoff, A. N., Toub, T. S., & Ruba, A. L. (2016). Infants' generalizations about other people's emotions: Foundations for trait-like attributions. *Developmental Psychology*, *52*(3), 364-378.

Rholes, W. S., & Ruble, D. N. (1984). Children's understanding of dispositional characteristics of others. *Child development*, *55*(2), 550-560.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173-220). Academic Press.

Silver, I. M., & Shaw, A. (2018). Pint-sized public relations: The development of reputation management. *Trends in cognitive sciences*, *22*(4), 277-279.

Surian, L., Ueno, M., Itakura, S., & Meristo, M. (2018). Do infants attribute moral traits? Fourteen-month-olds' expectations of fairness are affected by agents' antisocial actions. *Frontiers in psychology*, *9*, 1649.

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

Thomsen, L., & Carey, S. (2013). Core cognition of relational models. In M. R. Banaji, & S. A. Gelman (Eds.), *Navigating the Social World: What infants, children, and other species teach us* (pp. 16-22). Oxford University Press, New York.

Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A person-centered approach to moral judgment. *Perspectives on Psychological Science*, *10*(1), 72-81.

Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, *11*(12), 3571-3594.

Wynn, K. (2008). Some innate foundations of social and moral cognition. In P. Carruthers, S. Stich, & S. Laurence (Eds.). *The innate mind. Foundations and the future* (Vol. III, pp. 330–347). Oxford University Press.