Agency boosts sensitivity to social interactions through active information seeking

Qi Liang[a], Zishan Su[a,b], Rekha S. Varrier[a,c], Eshin Jolly[a,d] and Emily S. Finn[a,1]

[a]Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH.

[b]Current affiliation: School of Computer Science, Wuhan University, Wuhan, China.

[c]Current affiliation: Center for Economics and Neuroscience, University of Bonn, Bonn, Germany.

[d]Current affiliation: Department of Psychology, University of California San Diego, La Jolla, CA.

[1]To whom correspondence should be addressed. E-mail: Emily.S.Finn@dartmouth.edu

Keywords: social perception | agency | perspective difference | information seeking

Abstract

Humans use visual information to detect social interactions not only between others (third-person perspective), but also—and more importantly—between themselves and others (first-person perspective). Previous work in social perception has mostly focused on the former, leaving first-person scenarios underexplored. In this study, we created simple visual scenes with varying levels of evidence for a social interaction between two agents, represented as dots. Participants either actively participated in the scene by controlling one of the dots with their mouse or passively observed the two dots' movements, then reported whether one dot was chasing the other (social) or the two were moving independently (nonsocial). Active participation biased participants toward reports of socialness and lowered their threshold for detecting an interaction, indicating heightened sensitivity to social information; these perspective differences persisted even when visual input was perfectly matched. Drift-diffusion modeling on self-paced responses revealed a significantly higher drift rate in the active condition, suggesting that faster evidence accumulation was related to the boost in sensitivity. Further analysis of mouse trajectories showed that participants who engaged in more exploratory, information-seeking movements (i.e., heading towards the other agent during a possible chase) exhibited lower thresholds for social signal detection in the first-person perspective and larger perspective differences. These findings demonstrate that agency boosts sensitivity to social interaction by enabling active information sampling and accelerating evidence accumulation, highlighting the importance of the first-person perspective for understanding human social cognition.

Significance Statement

Humans are especially attuned to social information in their environment. While people can quickly and accurately detect social interactions using limited visual information, most past research on social perception has adopted the third-person perspective—that is, participants judge potential social interactions between others, but not between others and themselves. Here, we directly compare social perception between first- and third-person perspectives. We find that actively participating in, compared to passively observing, a potential interaction lowers the threshold to declaring something “social”. People accumulate evidence for socialness faster when they themselves are part of the interaction, and individuals who engage in more active exploration are even more sensitive to social information. Findings highlight the importance of the first-person perspective in social perception.

Introduction

Humans are social beings; understanding and engaging in social interactions is essential to our survival. The first step in social cognition is to detect social information in the environment. Since Fritz Heider and Marianne Simmel’s landmark 1944 study (1) showing that humans spontaneously perceive social interactions among simple animations of geometric shapes, extensive work has shown that humans are primed to perceive “socialness” even from minimal visual displays (2–7), suggesting that basic components of social interactions are extracted quickly and automatically by the human visual system (8). Furthermore, social information, once perceived as such, enjoys a range of cognitive advantages, including enhanced attention, faster and more efficient processing, and greater memory capacity and precision (9–15).

While existing work convincingly demonstrates a perceptual and cognitive privilege for social information, most past studies of social vision adopt the third-person perspective (3PP), using passive-viewing paradigms to measure how participants detect social signals—i.e., the presence and the nature of a social interaction—among other agents when they are outside observers. However, while people can and do learn from observing others, in our daily lives, interactions that we ourselves are part of are usually more salient and significant to us (16, 17). Although passive-observer paradigms lend themselves to a high degree of experimental control, enabling precise inferences about the visual features and mental computations over those features that engender social percepts, they leave a gap in our understanding of how we perceive social interactions in what is arguably a much more important perspective: first-person perspective (1PP).

A key difference between third- and first-person perspectives is the presence of agency. In a first-person context, people engage in the interaction as active agents, meaning they have control over what information gets sought, sampled and encoded via their own actions (18, 19). In other domains, including learning and memory, agency has been shown to improve information processing and retention: people showed faster and more efficient learning as well as better episodic and temporal memory when given the opportunity to control or select the information being encoded (20–24). This agency advantage aligns with the self-reference effect, a well-established phenomenon whereby information related or relevant to oneself is prioritized for encoding and retrieval (25, 26). These findings, together with extensive work documenting the phenomenological differences between first- and third-person perspectives (27–29), motivate an investigation of how agency affects perception of social interactions: it stands to reason that detecting whether two agents in your environment are interacting with one another may differ systematically from detecting whether an agent is trying to interact with *you*.

Although some studies have explored social signal detection in first-person contexts (30–33), direct comparison with third-person contexts has remained challenging due to the absence of controlled paradigms. Ideally, participants would experience identical—or at least extremely similar—sensory input under both active and passive conditions, to isolate the effect of perspective above and beyond any differences in low- and mid-level visual features. Previous research often employed handcrafted or naturalistic stimuli, which are difficult to manipulate quantitatively and even harder to compare across perspectives while controlling for visual information (7, 8, 17, 34, 35). Here, we overcame this challenge using a recently developed parameterized approach to study social signal detection (31, 36). We created simple visual scenes consisting of two agents, represented as moving dots, with varying objective levels of evidence for a social interaction. Critically, these scenes could be experienced commensurately across both the first- and third-person perspectives: participants either actively participated in the scene by controlling one of the dots with their mouse or passively observed the two dots' movements, then reported whether or not they perceived a social interaction. Results showed that, compared to passive observation, active participation enhanced bias towards social percepts and boosted sensitivity to social information. A series of control experiments confirmed that the perspective differences persisted when visual input was perfectly matched, when agent identity was disclosed, and when responses were self-paced, ruling out confounds due to low-level visual information, uncertainty, or cognitive load, respectively. Further, drift-diffusion modeling and analysis of participants' agent-movement trajectories indicated that agency may boost sensitivity to social signals by enabling active information sampling and a faster evidence accumulation rate. Overall, results speak to the importance of centering the first-person perspective in the earliest stages of social perception.

Results

To test if and how perspective affects social signal detection, we ran one main experiment consisting of two sessions (n = 268 participants for session 1, with n = 108 returning for session 2) as well as several follow-up experiments to rule out the influence of possible confounds on observed results. All data were collected online via Prolific.

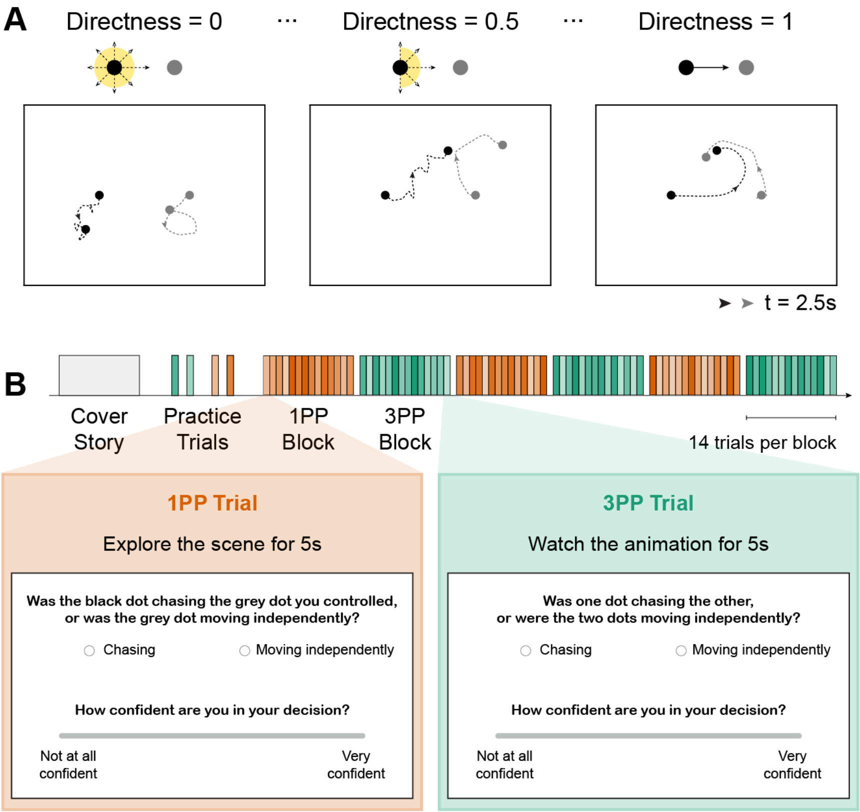

Was the black dot chasing the grey dot you controlled, or was the grey dot moving independently?

○ Chasing ○ Moving independently

How confident are you in your decision?

Not at all confident — Very confident

Was one dot chasing the other, or were the two dots moving independently?

○ Chasing ○ Moving independently

How confident are you in your decision?

Not at all confident — Very confident

Fig. 1. Experimental design and tasks. (A) Socialness manipulation. We created simple visual scenes consisting of two agents, represented as moving dots. To vary the amount of objective evidence for a social interaction between the two dots, we manipulated chase directness, which governs the fidelity with which the predator (black in these examples) chased the prey (gray), across 7 levels between 0 and 1. When chase directness = 0, the predator randomly chooses a heading direction on each timestep with no directed movement toward the prey; when chase directness = 0.5, the predator chooses a heading direction within ±90° of the prey; and when chase directness = 1, the predator heads directly toward the prey. Example trajectories (dashed lines) illustrate the movement patterns over a 5 s trial period (arrows showed the dots locations and headings at the 2.5 s moment). (B) Experimental procedure and trial structures. After a cover story and practice trials, participants completed alternating blocks of 1PP (first-person perspective, orange) and 3PP (third-person perspective, green) trials, with 14 trials per block, 3 blocks of each perspective. The gradient color of the trials represents the chase directness level; note that directness level was randomized across the whole experiment, not within blocks. In 1PP trials, participants actively controlled the prey dot while exploring the scene for 5 s, then judged whether the other dot was chasing their dot or moving independently. In 3PP trials, participants passively watched an animation for 5 s showing two dots, then judged whether one dot was chasing the other or both were moving independently. Participants were asked to rate their confidence in both types of trials.

All scenes consisted of two agents, represented as a gray and a black dot. We varied the amount of objective evidence for a social interaction between the two dots using the parameter “chase directness”, which governs the fidelity with which one dot (the “predator”) chased the other dot (the “prey”) (31). At the start of the experiment, participants were given a cover story that the two agents represented children in a park; their task was to determine whether one agent was chasing the other (perceived social) or whether the two dots were moving independently (perceived nonsocial). We included this cover story to promote uniform percepts of animacy across levels of chase directness and perspectives, since our goal here was to isolate percepts of *socialness*, for which we considered animacy necessary but not sufficient.

We presented these scenes in two perspectives, first-person (1PP) and third-person (3PP), for 48 trials each (7 levels of chase directness [Fig. 1A] × 6 trials per level) (Fig. 1B). In 3PP trials, participants passively watched the two moving dots. In 1PP trials, participants controlled the “prey” agent with their mouse and were instructed to move around the scene; the potential “predator” dot was programmed to chase the participant-controlled dot with a given level of chase directness. All scenes across both perspectives were 5 s in duration. The initial set of animations shown to participants in the passive condition in session 1 of the main experiment were generated by an independent sample of human participants in pilot experiments. At the end of each trial, in a two-alternative forced choice (2AFC) task, participants indicated whether one dot was chasing the dot they controlled (1PP) or the other dot (3PP) versus moving independently and reported their confidence in their choice using a continuous slider.

Active participation increases bias toward social percepts

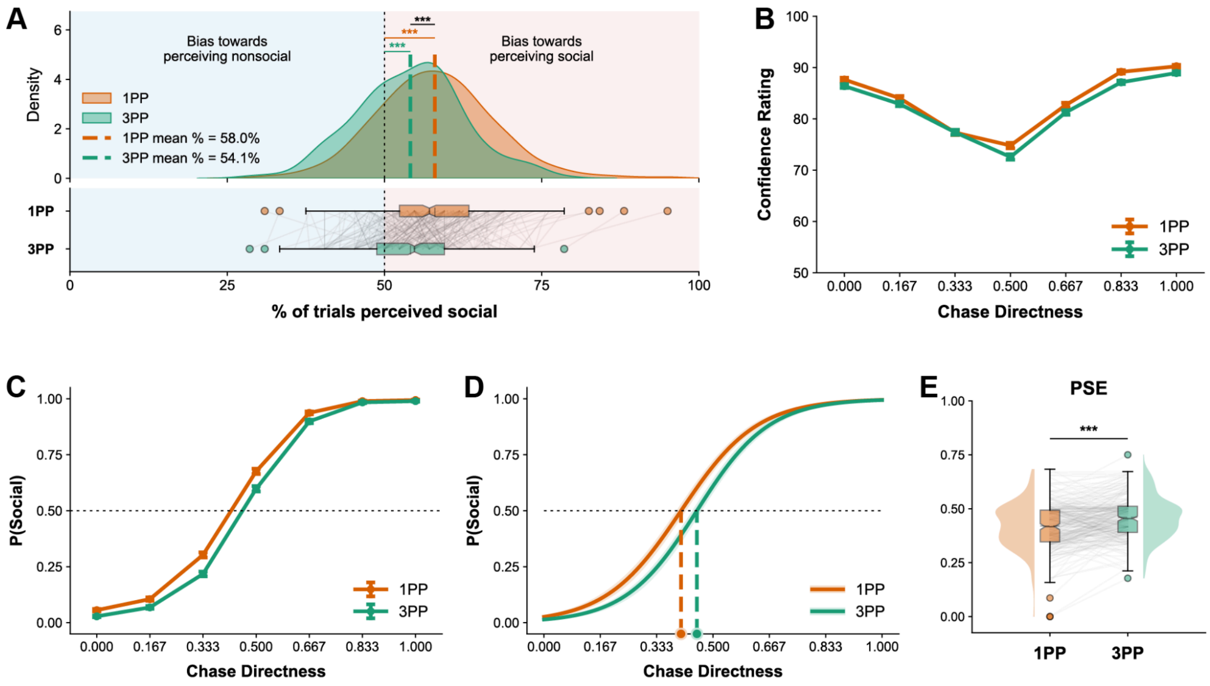

Fig. 2. Main experiment results showed social biases and boosts in sensitivity to social information during active participation (22328 total responses from n = 268 participants). (A) Distributions of two-alternative forced choice (2AFC) responses in the first session of the main experiment. Density plots show the distribution of individual participant means for the percentage of trials perceived as social (chasing) in 1PP (orange) and 3PP (green) conditions. Box plots below show median and quartiles, with dashed lines showing the mean and connecting lines indicating paired data from the same participants across conditions. The dotted vertical line at 50% indicates an even split between responses (no bias). Participants' percepts were biased towards social in both perspectives and this bias was stronger in 1PP scenarios. (B) Confidence ratings across chase directness levels. Error bars indicate standard error. Confidence followed a U-shaped pattern, indicating greater certainty when evidence is clearer. The confidence was in general higher in 1PP than 3PP scenarios. (C) 2AFC responses across chase directness levels. Trials perceived as social were coded as 1 and nonsocial coded as 0, and responses were averaged at each chase directness level per participant. The dashed horizontal line at 0.5 indicates the flipping point of social perception and error bars indicate standard error. P(Social) means probability of judging trials as social. Compared to passively watching, participants tended to respond seeing social more when they were part of the interaction and this increment was largest in middle ranges of chase directness. (D) Psychometric curves with perceptual thresholds. Same data as panel C with vertical dashed lines indicating the point of subjective equality (PSE)—the directness level at which

participants judged trials as social 50% of the time. Shading indicates standard error. (E) PSE comparison. Individual participant PSEs are shown as dots with connecting lines, box plots show median and quartiles (dashed lines show mean), violin plots show distributions. A significant difference was found between conditions, indicating that the threshold for perceiving social interactions was lower in the first-person perspective. ***$p < .001$.

First, we examined if participants exhibited a response bias in either or both perspectives. We coded participants' binary responses as 0 for “moving independently" and 1 for "chasing" and calculated the probability of reporting seeing a social interaction in each perspective (Fig. 2A). In general, participants’ judgments were biased towards “chasing” (one-sample t-test against 0.5; 1PP: mean = 0.580, *SE* = 0.006, *t*(267) = 14.175, $p < .001$; 3PP: mean = 0.541, *SE* = 0.005, *t*(267) = 7.923, $p < .001$) in both conditions. Moreover, compared to passively watching, participants’ judgments were biased even more towards seeing a social interaction when they were part of the interaction (paired t-test, 1PP vs. 3PP: *t*(267) = 7.364, $p < .001$). This result aligns with and extends previous studies supporting the hypothesis that the cost of failing to recognize social cues (misses) is higher than mistakenly perceiving some cues as social (false alarms) (7, 33), suggesting that this trade-off might be even more salient in first-person scenarios.

Active participation boosts sensitivity to social information

We then examined how subjective social percepts tracked objective evidence for socialness across perspectives by fitting a generalized linear mixed-effect model (GLME) to predict response as a function of our independent variable *chase directness*. Aligned with previous work (31–33, 36), we found that changing a simple motion parameter shaped participants’ subjective social perception: participants were more likely to report seeing a social interaction when the chase was more direct (Fig. 2C, fixed effect of directness: *OR* = 2.008×10^4, *SE* = 3.617×10^3, $z = 55.000$, $p < .001$). Compared to passively watching, participants tended to respond seeing social more when they were part of the interaction and this increment was largest in middle ranges of chase directness (perspective: *OR* = 0.496, *SE* = 0.058, $z = -5.995$, $p < .001$; perspective x chase directness: *OR* = 1.701, *SE* = 0.421, $z = 2.147$, $p = .032$).

As expected, we found a quadratic effect of chase directness on confidence, such that participants became more confident toward the extremes of objective evidence when there was clear evidence for or against a social interaction (i.e., chase directness = 1 or 0, respectively. Fig. 2B, linear term of directness: β = -251.558, *SE* = 19.817, *t* = 12.694, $p < .001$; quadratic term of directness: β = 679.244, *SE* = 19.829, *t* = 34.255, $p < .001$). Compared to passively watching, participants tended to have higher confidence when actively participating in an interaction (perspective: β = -1.339, *SE* = 0.187, *t* = -7.168, $p < .001$), but there was no interaction between perspective and chase directness on confidence (perspective x linear term of directness: β = -39.073, *SE* = 28.000, *t* = -1.395, *p* = .163; perspective x quadratic term of directness: β = 4.540, *SE* = 28.015, *t* = 0.162, *p* = .871).

To further quantify participants' social perception, adopting methods from psychophysics, we fitted each participant's average responses across directness levels with a psychometric function (Fig. 2D, see Methods for details) and calculated their point of subjective equality (PSE) in both perspectives. PSE quantifies the stimulus intensity (i.e., value of chase directness) at which the stimulus is equally likely to be perceived as social or nonsocial. Compared to passively watching (3PP), participants had a lower PSE in active (1PP) trials (Fig. 2E, paired t-test, *t*(267) = -6.363, $p < .001$), indicating that participating in social interaction boosts sensitivity to social information and lowers the threshold for declaring something "social".

Perspective differences persist even under matched visual input

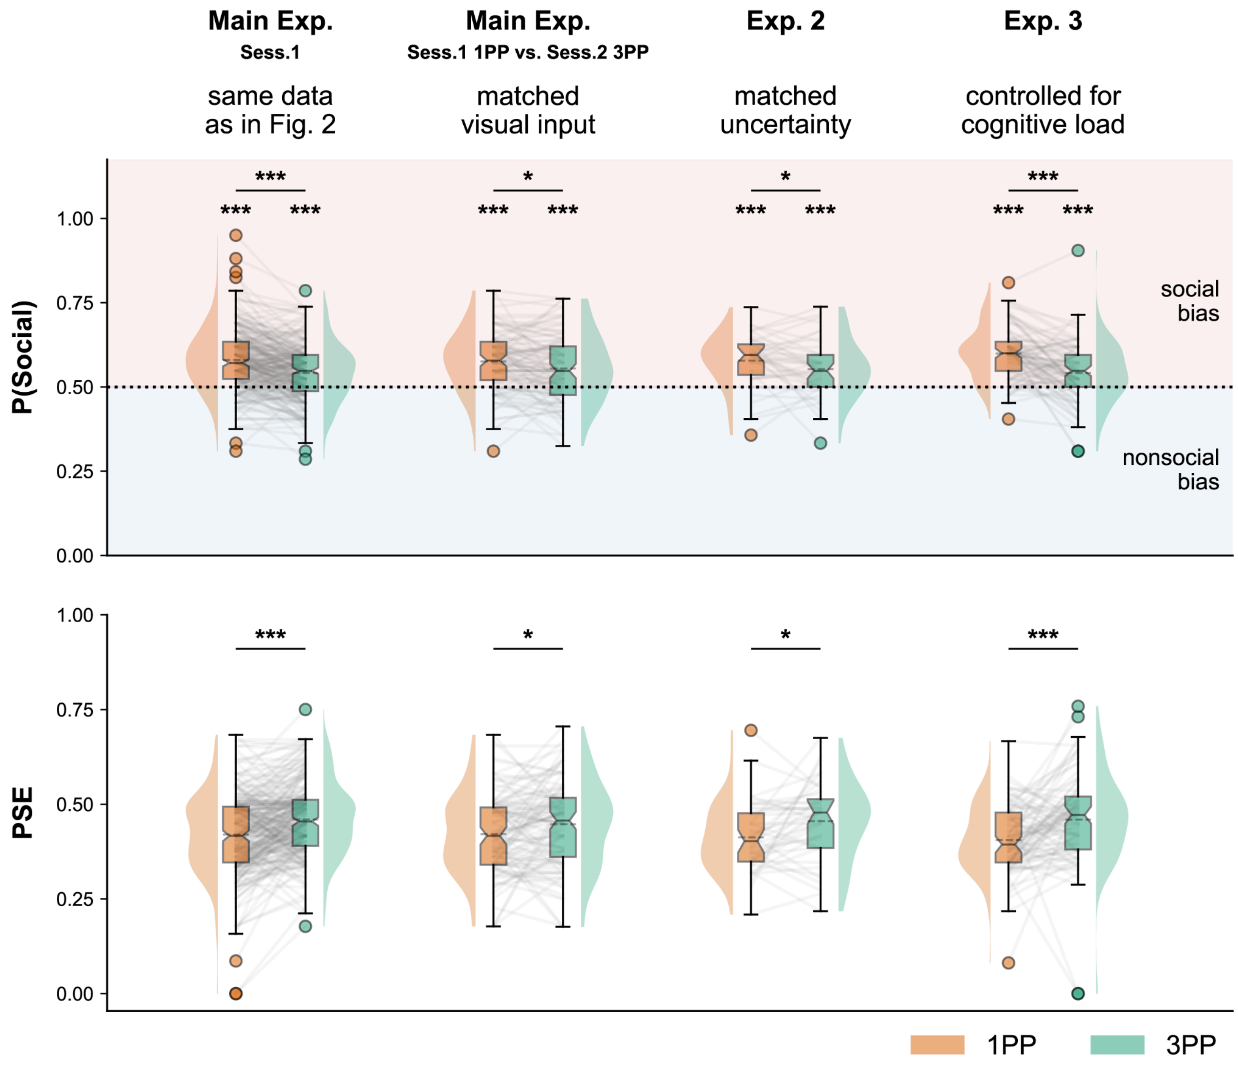

Fig. 3. Increased social bias and boost in sensitivity during active participation persisted across control experiments. Top row: P(Social), probability of judging trials as social; asterisks above each box indicate one-sample t-tests against no bias probability (0.5); asterisks between boxes indicate paired t-tests between perspectives. Bottom row: PSE, point of subjective equality; asterisks indicate paired t-tests between perspectives. From left to right: Main Experiment Session 1 (same data as in Fig. 2 A and E), n = 268 participants; Main Experiment Session 1 vs. Session 2, a matched visual input comparison since the 1PP trials from session 1 were used as 3PP stimuli in session 2, n = 108; Experiment 2 controlled for uncertainty about agent roles by informing participants of the potential predator dot's color before the stimulus was displayed in 3PP trials, n = 43; Experiment 3 mitigated for cognitive demand differences by letting participants respond as soon as they formed a sense of social percepts, n = 93. *$p < .05$, ***$p < .001$.

Even though our experimental setup allowed us to create first- and third-person perspective scenes that were extremely similar in their visual properties, for participants in the first session of experiment 1, 3PP stimuli came from a bank of 1PP trials generated by other human participants in pilot experiments, meaning the exact trajectories of the agents were not identical between perspectives. To test whether perspective differences persisted even when the visual input itself was perfectly matched, we invited participants back for a second session in which all 3PP stimuli were playbacks of each individual participant's own 1PP trials in session 1. Thus, comparing participants' behavior in session 1 1PP trials with session 2 3PP trials strictly controls for the visual trajectories of the agents. We found a smaller but still significant increase in social bias, sensitivity, and confidence for trials experienced in the active perspective (Fig. 3, second column, top row, P(Social): 1PP vs. 0.5, $t(107) = 8.810$, $p < .001$; 3PP vs. 0.5, $t(107) = 5.785$, $p < .001$; 1PP vs. 3PP, $t(107) = 2.620$, $p = .010$; bottom row, PSE: 1PP vs. 3PP: $t(107) = -2.559$, $p = .012$. Fig. S2, top right, perspective: $\beta = -1.745$, $SE = 0.318$, $t = -5.492$, $p < .001$), highlighting that agency is the main driver of perspective differences above and beyond visual input. By comparing responses within session 2, we replicated our previous set of results (Fig. S1), suggesting that perspective differences persist into a second session in which all trajectories are generated by the same human participant, which further controls for any differences in exploration styles between individuals.

Role uncertainty or cognitive load are not driving factors of perspective differences

We next sought to rule out two other possible factors that could confound the observed perspective differences.

One difference between 1PP and 3PP trials in the main experiment was the uncertainty surrounding each agent's identity. In 1PP trials, participants were always instructed that they were the "prey" agent (i.e., the agent possibly being chased), meaning the other agent was by default the possible "predator"; in 3PP trials, en route to deciding whether a chase was present, they had to figure out which of the agents was which. Not knowing the possible predator-prey identity ahead of time may have led to higher uncertainty in 3PP trials, which could mean that participants needed more evidence to make judgments. To test the extent to which this difference could account for our observed results, we performed an auxiliary "identity-known" experiment. In this experiment, participants (n = 43) were told the color of the agent possibly being chased before the 3PP trials. After watching each 5 s animation, participants were asked, e.g., "Was the gray dot chasing the black dot, or were the two dots moving independently?". We found that, after eliminating the uncertainty difference, there were still significant differences between perspectives in social bias, sensitivity and confidence (Fig 3, third column, top row, P(Social): 1PP vs. 0.5, $t(42) = 6.543$, $p < .001$; 3PP vs. 0.5, $t(42) = 4.069$, $p < .001$; 1PP vs. 3PP, $t(42) = 2.055$, $p = .046$; bottom row, PSE: 1PP vs. 3PP: $t(42) = -2.619$, $p = .012$. Fig. S2, bottom left, perspective: $\beta = -2.723$, $SE = 0.509$, $t = -5.348$, $p < .001$).

A second possible confounding difference between perspectives is the degree of task demands, or cognitive load. In 1PP trials, participants had to move their mouse to control one of the agents while simultaneously judging the presence versus absence of a social interaction, which likely required more mental resources than simple passive viewing. While the possible influence of this difference on sensitivity to social information is not necessarily straightforward—one might predict, for example, that higher cognitive load could lead to *less* sensitivity in active scenarios as participants would need more evidence to make a socialness judgment—we nonetheless sought to explore the extent to which this factor affected our results. To this end, we performed a self-paced response time experiment (n = 93), in which participants were allowed to respond anytime within the 5 s stimulus duration once they arrived at their judgment. We once again found that there were still significant differences between perspectives in social bias, sensitivity and confidence (Fig 3 fourth column, top row, P(Social): 1PP vs. 0.5, $t(92) = 12.326$, $p < .001$; 3PP vs. 0.5, $t(92) = 4.432$, $p < .001$; 1PP vs. 3PP, $t(92) = 5.784$, $p < .001$; bottom row, PSE: 1PP vs. 3PP: $t(92) = -3.843$, $p < .001$. Fig. S2, bottom middle, confidence: $\beta = -1.290$, $SE = 0.380$, $t = -3.393$, $p < .001$), suggesting the differences between perspectives is not due to lack of time in 1PP trials. Furthermore, while response times were slower in the first-person perspective, they were still comfortably within the 5 s trial duration (Fig. 4A), indicating that participants in the main experiment should not be experiencing different speed-precision tradeoffs across perspectives.

Faster evidence accumulation in active trials explains increased social bias and sensitivity

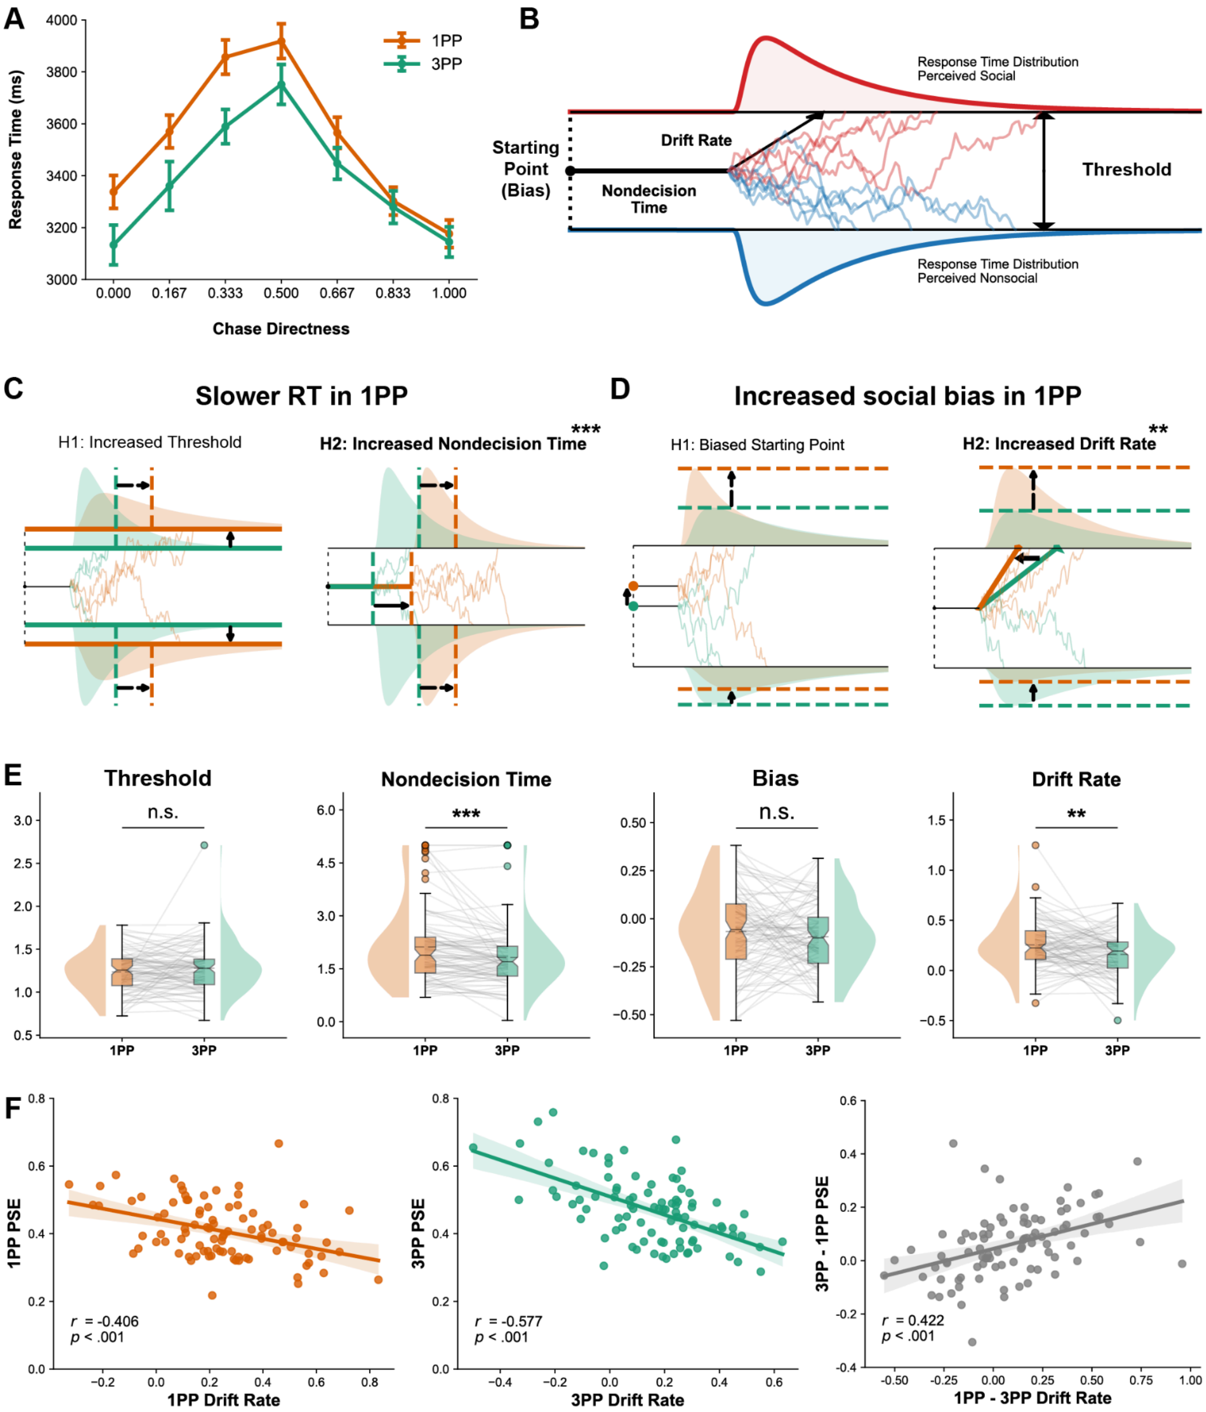

Fig. 4. Drift diffusion modeling revealed faster evidence accumulation during active participation, which explained increased social bias and sensitivity. (A) Response times across chase directness levels (7717 total trials from n = 93

participants). Mean response times (ms) plotted as a function of chase directness for 1PP (orange) and 3PP (green) conditions. Both conditions showed an inverted U-shaped pattern, with the longest response times at intermediate directness levels where evidence was most ambiguous. 1PP trials showed consistently longer response times than 3PP trials. Error bars represent standard error of the mean. (B) Drift diffusion model schematic. The model framework illustrates how evidence accumulates over time from a starting point toward one of two decision boundaries (perceived social vs. nonsocial). Key parameters include bias (initial bias), drift rate (speed of evidence accumulation), nondecision time (encoding and motor response), and threshold (amount of evidence required). Response time distributions for each decision type are shown at the top and bottom boundaries. (C) Two hypotheses that could account for slower response times in 1PP. Hypothesis 1 proposes an increased threshold (requiring more evidence before committing to a decision), shown by wider separation between boundaries; hypothesis 2 proposes increased nondecision time (longer encoding or motor preparation), shown by a longer duration before evidence accumulation begins. Dashed vertical lines indicate the shift in response times. (D) Two hypotheses that could account for increased social bias in 1PP. Hypothesis 1 proposes a biased starting point closer to the "social" boundary (indicated by orange dot shifted upward relative to green); Hypothesis 2 proposes increased drift rate toward the "social" boundary in 1PP (steeper slope of evidence accumulation trajectories in orange). (E) Drift diffusion model parameter estimates. Threshold and bias show no significant differences between perspectives, while nondecision time is significantly longer in 1PP and drift rate is significantly higher in 1PP. Individual participant data shown as dots with connecting lines; box plots show median and quartiles; violin plots show distributions. (F) Correlations between drift rate and sensitivity. Each dot is a participant. Drift rates negatively correlated with PSE in both 1PP and 3PP, indicating faster evidence accumulation toward "social" judgments corresponds to more sensitivity. The difference in drift rate between conditions (1PP minus 3PP) correlated with the difference in PSE (3PP minus 1PP), suggesting that the difference in evidence accumulation speed might drive the boost in sensitivity. **$p < .01$, ***$p < .001$, n.s. = not significant.

Data from the response time experiment revealed that, in addition to slower overall response times in the first-person perspective (perspective: $\beta = -140.814$, $SE = 26.034$, $t = -5.409$, $p < .001$), in both perspectives, participants were slowest to respond at the middle levels of chase directness (Fig. 4A, linear term of directness: $\beta = -8.278 \times 10^3$, $SE = 1.634 \times 10^3$, $t = -5.065$, $p < .001$; quadratic term of directness: $\beta = -1.954 \times 10^4$, $SE = 1.636 \times 10^3$, $t = -11.942$, $p < .001$); this was unsurprising given that these levels are most ambiguous as to whether or not a social interaction is present. Collapsing across directness levels, we next tested whether response times differed as a function of response ("chasing" versus "moving independently"). We hypothesized that, in line with the idea of socialness as a prepotent percept, "chasing" responses might be faster. However, we found no consistent difference between response times for the two alternatives (Fig. S4, response: $\beta = -21.221$, $SE = 38.660$, $t = -0.549$, $p = .583$).

How can we explain both the overall slowing of response times and the increased bias toward "social" responses in the first-person perspective? To further investigate the underlying cognitive mechanisms driving the different response patterns between active and passive scenarios, we applied a drift-diffusion model (DDM) to the data from the response time experiment (Fig. 4B). The DDM jointly models response time and responses themselves using four parameters: (1) bias, which quantifies the initial bias to favor one response or the other at the start of the evidence accumulation process; (2) threshold, which quantifies the level of evidence required to make a decision; (3) drift rate, which quantifies the speed and direction of evidence accumulation toward a decision boundary (in our case, more positive drift rates would indicate a faster pull toward "social" responses); and (4) nondecision time, which represents the length of all processes unrelated to decision-making (e.g., encoding, motor response time). We first confirmed that letting perspective affect all of the four parameters was the best model by performing model comparisons (Fig. S5). We then fitted the DDM to each participant's response time distributions (collapsed across all chase directness levels) to estimate these four parameters separately for 1PP and 3PP trials. In these models, the upper boundary corresponded to "social" responses, while the lower boundary corresponded to "nonsocial" responses.

First, we hypothesized that the overall slower response times in 1PP trials could be explained by a longer nondecision time (Fig. 4C, right), which would encompass the additional task demands in the form of motor execution required in this perspective (i.e. move the mouse to explore). Consistent with this hypothesis, the nondecision time was significantly longer in the 1PP condition than in the 3PP condition (Fig. 4E, second, $t(92) = 4.113$, $p < .001$). An alternative explanation for this effect could have been an increased threshold in the first-person perspective (Fig. 4C, left), but there was no significant perspective difference in the threshold (Fig. 4E, first, $t(92) = -1.403$, $p = .164$).

Second, we reasoned that the enhanced social bias in first-person perspective could arise from two non-mutually potential exclusive mechanisms (Fig. 4D): (1) a shift in the starting point towards the “social” decision boundary, indicating a stronger *a priori* response bias, and/or (2) an increase in the drift rate, reflecting more efficient evidence accumulation for social cues. Consistent with the second possibility, we found a significantly higher drift rate in 1PP condition compared to 3PP condition (Fig. 4E, fourth, $t(92) = 3.003$, $p = .003$) but no significant bias differences (Fig. 4E, third, $t(92) = 1.038$ $p = .302$). Taken together, these findings indicated that the increased social bias observed in the first-person perspective do not stem from a pre-existing response bias (an altered starting point), but rather from a more efficient evidence accumulation process (a higher drift rate). This suggests that active involvement (agency) in an interaction enhances the processing of social signals, increasing their salience and more effectively driving the decision-making process towards a social percept.

Furthermore, we found that across participants, drift rates negatively correlated with PSE in both first- and third-person perspectives (Fig. 4F, left and middle plots; 1PP: $r = -0.406$, $p < .001$; 3PP: $r = -0.577$, $p < .001$), indicating that faster evidence accumulation toward "social" judgments leads to higher sensitivity (indexed by lower PSE) regardless of perspective. The difference in drift rate between conditions (1PP minus 3PP) positively correlates with the difference in PSE (3PP minus 1PP) (Fig. 4F, rightmost plot; $r = 0.422$, $p < .001$), suggesting that the difference in evidence accumulation speed might contribute to the boost in sensitivity for the first-person perspective.

First-person sensitivity boosts are related to active information-seeking behaviors

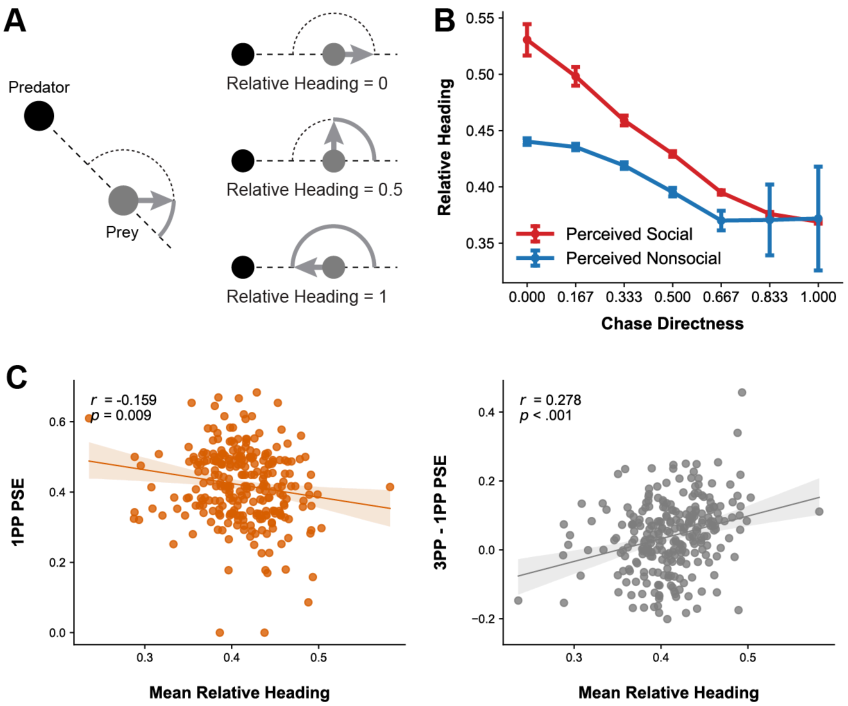

Fig. 5. Analysis of participants' mouse-movement trajectories reveals the relationship between active information-seeking behavior and sensitivity to social information. (A) Relative heading schematic. The relative heading quantifies how directly the prey orients toward the predator. When the prey moves directly toward the predator, the relative heading = 1. When the prey moves directly away from the predator direction, the relative heading = 0. Higher values indicate more exploratory, less escape-like movement. (B) Relative heading differs between trials perceived as social vs. nonsocial in the first-person perspective, especially at low chase directness. Error bars indicate standard error (11118 total trials across n = 268 participants). (C) Individual differences in mean relative heading correlate with social perception: negatively with 1PP PSE (left) and positively with PSE differences (right, 3PP minus 1PP PSE), suggesting more information seeking behavior is related to enhanced sensitivity in the first-person perspective.

In our final set of analyses, we returned to data from the main experiment to explore relationships between participants' active motor behavior during first-person trials and their ultimate percepts. Using participants' prey movement trajectories, we extracted the relative heading at each moment during the trial (Fig. 5A). A smaller relative heading (moving directly away from the predator's current location) suggest an “escape” strategy, whereby participants simply move as far away from the predator as quickly as possible; conversely, larger relative headings suggest an “explore” strategy, whereby participants move laterally or even toward the possible predator to stay closer, perhaps to test the other dot’s reaction. We found that, across chase directness levels, on trials when participants showed more “explore”-like movements, they tended to ultimately subjectively respond seeing a social interaction (Fig. 5B, response: $\beta = 0.050$, $SE = 0.001$, $t = 69.135$, $p < .001$). Visualizing the time courses of relative heading revealed that explore-like behavior emerged early in the trial when the scene was ultimately perceived as social (Fig. S6, response x time: $\beta = -1.518 \times 10^{-4}$, $SE = 7.622 \times 10^{-6}$, $t = -19.922$, $p < .001$).

To confirm the relationship between explore-like behavior and subjective social percepts, we correlated participants’ relative heading (“explore index”) with their overall social perception tendencies. Collapsing across all trials and directness levels, we found that participants who showed more exploratory behavior on average tended to have lower PSE in the first-person perspectives (Fig. 5C, left, 1PP PSE: $r = -0.159$, $p = .009$) and a greater difference in PSE between perspectives (Fig. 5C, right, 3PP minus 1PP PSE: $r = 0.278$, $p < .001$). Importantly, neither first-person PSEs nor PSE differences were correlated with simpler movement indices such as participants’ total moving time or mean speed, acceleration, or jerk (Fig. S7), suggesting that this effect is not due to general engagement in the task, but rather specific to a more explore-like (versus escape-like) strategy. We replicated the relationship between information-seeking behavior and sensitivity in the response time experiment (Fig. S8). Together, results indicate that participants who engaged in more exploratory information-seeking had a sense of agency that enhanced their sensitivity to social information specifically in active scenarios.

Finally, having identified two factors that partially drive enhanced sensitivity in the first-person perspective (increased drift rate as revealed by drift-diffusion modeling and increased explore-like behavior revealed by trajectory analysis), we wondered if these were part of the same underlying mechanism or if they operate independently. Returning to data from the response time experiment, we tested whether mean relative heading correlated with drift rate across participants, but found no such relationship ($r = -0.076$, $p = .468$). Multiple linear regressions including both relative heading and drift rate as predictors of PSE revealed that both factors had unique and significant contributions (1PP PSE: 1PP drift rate $\beta = -0.103$, $SE = 0.035$, $t = -2.917$, $p = .004$; 1PP mean relative heading $\beta = -0.237$, $SE = 0.088$, t

= -2.683, *p* = .009. 3PP minus 1PP PSE: 1PP minus 3PP drift rate *β* = 0.177, *SE* = 0.041, *t* = 4.343, *p* < .001; 1PP mean relative heading *β* = 0.367, *SE* = 0.127, *t* = 2.887, *p* = .005). This suggests that, while more explore-like behavior does not predict an increased drift rate, both factors can independently contribute to sensitivity boosts in the first-person perspective.

Discussion

Here, we systematically compared social signal detection when participants actively participated in versus passively watched a possible interaction. Across multiple experiments, we observed two consistent effects of agency: increased bias toward social percepts and enhanced sensitivity to social information. These effects persisted even when visual input was perfectly matched, when uncertainty about agent identity was eliminated, and when differences in cognitive load between perspectives were mitigated, indicating that agency—not low-level stimulus differences, task-related uncertainty, or difficulty—drives the observed perspective differences. Further drift-diffusion modelling and analysis of participants' movement trajectories suggested that agency boosts sensitivity to social information by hastening evidence accumulation, and that participants who actively sought information benefited most from agency.

Social psychologists have increasingly recognized the need for interactive experimental frameworks (37–42), also called "second-person" social cognition and neuroscience (16, 17). Previous work has shown, for example, that the mere opportunity for social interaction can alter gaze dynamics (43, 44), synchronize affect (45), and boost activity in brain regions associated with mentalizing (46). We uncover evidence for these phenomena even earlier in the processing pipeline: before engaging in a social interaction, one must first detect the presence of social information in the environment; our results show that agency has an effect even at the most basic level of classifying whether information is social or nonsocial. In contrast to the growing number of studies in the interactive tradition, studies probing the role of the visual system in social perception (reviewed in (8)) have largely continued to rely on the detached, third-person perspective. Our results raise the possibility that past work in social vision, while having convincingly demonstrated enhanced processing for social information at some of the earliest stages of the perceptual-cognitive hierarchy (11–14, 47), by relying only on passive observation may have in fact *underestimated* the perceptual advantage conferred by socialness.

Several avenues might enrich further investigation. First, our parametric psychophysics-inspired framework, which allows for perfect control over visual information, lends itself well to neuroimaging. Past work has shown that mirror-neuron and mentalizing systems engage more strongly when participants are or simply believe they are in real-time social engagements (48–51). The perspective-dependent boost in social signal detection observed here raises the hypothesis that activity in the brain's third visual pathway (52, 53), which includes middle temporal and posterior temporal regions thought to be specialized for social perception, may be tuned not only to visual cues but also to the motor-feedback loop that accompanies agency, consistent with theories of active inference (54, 39). Second, this framework could be extended to characterize individual differences in social signal detection across perspectives. Indeed, recent work (reviewed in (55)) has shown that interactive paradigms provide a much richer characterization of social (dys)function in developmental and psychiatric conditions such as autism (56, 57) and social anxiety disorder (50, 51). Using our approach, one could probe whether aberrant social processing in these and other conditions (e.g., mood disorders, schizophrenia) arises as early as if and how individuals classify information as social in the first place: for example, whether social signal detection is preferentially altered in the first-person perspective, or whether individuals with these conditions engage in different amounts and patterns of information-seeking behavior when given agency to explore a potentially social scene. Third, this framework could be used to trace the developmental trajectory of agency-driven social signal detection across the lifespan. Social perception undergoes significant refinement from infancy through adolescence, and even very young children show sensitivity to interacting agents (60, 61). However, it remains unknown whether the agency advantage observed here emerges early in development. With this framework, one could examine whether agency effect depends on more protracted maturation of the systems that integrate motor feedback with social inference.

While our results convincingly establish an agency advantage for detecting interactions in simple visual scenes, this work has several limitations. First, moving dots are highly abstract relative to real-world visual information. While this choice of stimuli gave us precise control over agents' appearance and movements, allowing us to perfectly match visual information between perspectives, it will be important to replicate and extend these results using increasingly lifelike agents (e.g., faces, bodies, and/or avatars), as well as multi-agent scenarios, to determine if and how the agency advantage scales with interaction complexity. Additionally, the social interaction signal investigated here is chasing, which is a single interaction type with clear perceptual signature. Future work should examine whether the agency advantage extends to a broader range of interaction types (e.g., cooperation, or competing approach behaviors) (62, 54, 63–65), and whether active participation helps people for discriminate not only the presence but also the valence or nature of a social interaction(36)—for instance, distinguishing threatening from friendly encounters. Relatedly, participants always controlled the prey agent, limiting both the nature of the possible interaction and their role within that specific interaction; follow-up experiments could allow more open-ended behaviors to probe how people detect both the presence and the nature of a social interaction in less constrained scenarios. Finally, the form of agency here is purely motor-based; future work could examine how other types of agency (e.g., eye gaze, communicative gestures or utterances) affect the earliest stages of social signal detection.

Overall, our work demonstrates that agency—embodied in the first-person perspective—acts as a catalyst for social perception, lowering the threshold to detect social interactions and enhancing sensitivity to subtle social cues. These results add to a growing call to study social cognition from the participant's own viewpoint and suggest that incorporating agency into experimental designs may yield a more ecologically valid understanding of how humans navigate social environments.

Materials and Methods

Participants

All participants were recruited using the online platform Prolific (http://www.prolific.com/). To make sure participants understood the cover story and instructions while controlling for their familiarity with our stimuli, we recruited participants who self-reported being fluent in English, had their location set as the United States or the United Kingdom, and had no prior participation in studies from our laboratory that used similar stimuli. Numbers of participants varied by experiment and are reported in each respective section in the Experimental Design and Procedures section below. All participants provided informed consent before the experiments and all procedures were approved by the Committee for the Protection of Human Subjects at Dartmouth College.

After exclusion (Table S1), there were 268 valid participants for the main experiment (157 female, 109 male, 2 nonbinary; age [mean ± SD, same below] = 41.86 ± 13.68 yr), 43 for the identity known experiment (18 female, 20 male, and 5 not reported; age = 39.68 ± 13.18 yr), and 93 for the response time experiment (41 female, 50 male, 2 nonbinary; age = 40.23 ± 12.45 yr). 119 main experiment participants returned for a second session (interval = 203.82 ± 4.75 days), resulting in 108 valid participants for session 2 (56 female, 42 male, 2 nonbinary, and 8 not reported; age = 44.69 ± 13.57 yr).

Stimuli

All stimuli were generated with psyanim, a custom JavaScript-based software (https://github.com/thefinnlab/psyanim-2). To manipulate the strength of social information, we adopted the chasing paradigm of Gao et. al (31), and created scenes in which one agent (the "predator") was programmed to chase another agent (the "prey") to varying degrees. Social information was manipulated by altering the *chase directness* of the predator across seven levels spanning from 0 to 1. Chase directness was defined as a normalized transformation of the chase subtlety parameter ($\frac{180^\circ - subtlety}{180^\circ}$), defined in the original paper as the maximal angular deviation of the predator's heading direction compared to perfect heat-seeking (Fig. 1A). A chase directness of 0 indicates that, on each time step, the predator randomly selects a heading direction (±180°) regardless of the current location of the prey, while a chase directness of 1 indicates that the predator heads directly towards (±0°) the current location of the prey. Thus, there should be no evidence supporting the existence of social interaction when the chase directness was 0, and maximum evidence when the chase directness was 1. To make the movements smoother, the predator's heading direction was recomputed every 350 ms.

In the first-person (active) condition, to enable participants to control the movement of the prey while minimizing the difference in low-level visual features (e.g. speed and acceleration) across participants and perspectives, the prey was set as an agent directly chasing participant's mouse pointer (i.e. chase directness = 1), with the heading direction recomputed every frame (frame rate = 60Hz) and the maximum speed and acceleration set as 1.8px/frame and 0.15px/frame^2. To leave room for the prey to escape and minimize the chance that the predator would actually "catch" (i.e., make contact with) the prey, the max speed and acceleration for the predator were set slightly lower as 1.5px/frame and 0.1px/frame^2. Both agents were represented as dots (radius 12px), one black one gray, within a rectangular scene (800px by 600px). At the trial onset, these two agents were positioned symmetrically along the horizontal axis, each located 150 px from the center (i.e., at x = ±150 px relative to center).

With the exception of session 2 of the main experiment (described further below), all third-person stimuli were animations reconstructed from the active trials collected in a prior pilot experiment (n = 288). In that experiment, each participant (hereafter, generator) completed six trials at each chase directness level, using the same stimulus settings described above. All the trials were trimmed to 5 s and then trials in which (i) agents collided (predator-prey distance < 32px), (ii) the prey stayed stationary for more than 33% the 5 s, or (iii) the prey headed towards the predator for more than 80% of the 5 s were excluded from the stimuli sampling pool. Generators with fewer than two valid trials per chase directness level were also removed. After exclusion, the sampling pool consisted of 7425 trials from 209 generators, with a minimum of 935 trials per level. For each trial, we created four copies representing all possible combinations of color assignment between predator and prey and/or starting location to enable counterbalancing. Then for each participant in the present experiments, we first randomly selected 42 generators from the pool, assigned six generators to each of the seven chase-directness levels, and then randomly selected one animation per level per generator. This two-stage sampling procedure was used to reduce the influence of idiosyncratic trajectories from any single generator on participants' responses.

Experimental Design and Procedures

Cover Story

To isolate the concept of socialness above and beyond animacy, we presented a cover story as follows at the beginning of all experiments: "We recently videotaped a public park where local children go. Our goal is to understand children's behavior in a familiar park setting. To protect their privacy, we used an algorithm that represents each child as a dot and converts their trajectories in the park into movements of the dots. Throughout the experiment, there will be two dots in each trial, representing two children in the park.”. To ensure that participants read and remembered this cover story, we asked the following multiple-choice question both before and after the experimental trials: “What did the dots represent?” Participants who failed to select “children” among options (“animals”, “balls”, “adults”, “children”, “magnets”) for this question before the experimental trials were excluded from the recruitment.

Main Experiment

The main experiment was designed to directly compare social perception between first- and third-person perspectives. We used a 2 (two perspectives, active interacting vs. passive viewing) by 7 (seven chase directness levels) within-participant design (Fig. 1B). Each condition was repeated six times, resulting in a total of 84 trials per participant. Active and passive trials were grouped into three blocks each (14 trials per block), and the six total blocks were presented in interleaved order. The identity of the starting block (active versus passive) was randomized across participants. Levels of chase directness were randomized across the entire session rather than within each block to prevent participants from self-normalizing within blocks. In passive trials, participants watched a 5 s animation of two moving dots, reported whether one dot was chasing the other dot or moving independently, and rated their confidence on a continuous scale. The colors and the starting locations of the predator and prey dots were counterbalanced. In active trials, participants controlled one of the dots with their mouse, explored the scene for 5 seconds, reported whether the other dot was chasing the dot they controlled or moving independently, and also reported their confidence. The color of the participant-controlled dot (gray or black) was kept constant throughout each first-person block, but varied across the three first-person blocks within each participant. Participants were told at the start of each first-person block which color dot they would be controlling. The starting locations of the dots were counterbalanced within participants, and the color assignment (black-gray-black or gray-black-gray) was randomly selected and counterbalanced across participants. The 2AFC social judgment and confidence rating were presented on the same page, and the page timed out in 10s. Participants performed two practice trials per task at the beginning of the session (i.e., prior to the first experimental block).

After finishing the experimental trials, participants were asked to fill out six trait questionnaires:, (i) autism quotient questionnaire (AQ, 50 items) (66), (ii) Positive and Negative Affect Schedule (PANAS, 20 items) (67), (iii) NEO five-factor inventory for multidimensional personality (NEO-FFI; a measure of the "Big Five", 60 items) (68, 69), (iv) UCLA loneliness scale (version 3, 20 items) (70), (v) Revised Green et al. Paranoid Thought Scale (R-GPTS, 18 items) (71), (vi) Belief in the Purpose of Events (BPE, 22 items) (72), plus a single question asking them to report their number of friends. The sequence of the questionnaires was counterbalanced across participants. We inserted one attention-check question per questionnaire (e.g. "Select 'Neutral' for this question.") and used responses on trait questionnaires as part of our basic quality checks and screening criteria (see "Quality checks and exclusion criteria" subsection of Data Analysis section below) but did not analyze these data further in this study.

We recruited 350 participants for this experiment, of which 268 passed quality checks (Table S1).

Session 2

To rule out any effect of slight differences in visual input between active and passive trials, we invited participants who passed the screening criteria back for a second session. We launched the session 2 experiment 6.7 months after the first session and left recruitment open for 1 month (mean ± SD interval between sessions = 203.82 ± 4.75 days). We adopted the same experimental design as session 1, except that we replaced the animation stimuli for passive trials with the animations reconstructed from participants' own active trials in session 1. Therefore, by comparing participants' social judgments and confidence rating in session 1 active trials versus session 2 passive trials, we were able to estimate perspective differences with the visual input fully matched.

Of the 268 participants invited, 119 returned for a second session, of which 108 passed quality checks (Table S1).

Identity-Known Experiment

In the main experiment, agent identity (i.e., which dot was the possible predator and which was the prey) was uncertain in third-person trials but unambiguous in first-person trials (since participants always controlled the prey, they knew by default that the other agent was the possible predator). The identity- known experiment aimed at ruling out (un)certainty about agent identity as a possible confound in perspective comparisons. We used the same experimental design as the main experiment, except that we informed participants the color assigned to the prey and predator before the passive trials. Specifically, instead of randomizing the color of the dots, we kept the color of the dots the same within blocks, and let participants know of the agent roles by presenting the instruction beforehand as e.g. “You will be asked ‘Was the black dot chasing the gray dot, or were the two dots moving independently?’ in the following 14 trials”.

We recruited 50 participants for this experiment, of which 43 passed quality checks (Table S1).

Response Time Experiment

The response time experiment aimed at mitigating any effects of the difference in cognitive load between perspectives. We used the same 2x7 within-participant design, except that participants were allowed to respond whenever they formed a percept within each trial. We reasoned that if participants responded within 5 seconds, there shouldn’t be cognitive load differences in the main experiment as 5 seconds is sufficient for participants to form the social percepts in both perspectives. We collected data in two batches (60 participants recruited in each batch). To encourage participants to respond as soon as they made a judgment, in batch 1, we displayed the message “You didn’t respond within 5 seconds” for trials on which participants failed to respond before the stimulus ended. We removed this message in batch 2 to ensure we were not unwittingly introducing strong speed-accuracy tradeoffs and to more closely mimic conditions of the main experiment. Participants in both batches showed similar perspective differences and response time patterns (Fig. S3). Therefore, to increase statistical power, we pooled data from the two batches together for analysis.

We recruited 120 participants in total, of which 93 passed quality checks (Table S1).

Data Analysis

We performed similar analyses across all experiments unless noted.

Quality checks and exclusion criteria

To ensure data quality, we excluded participants who met any of these criteria: (i) did not experience all 84 trials (indicating a possible software or internet glitch), (ii) missed (i.e., failed to provide a response on) >5% of trials, (iii) failed an attention-check item regarding the cover story, (iv) remained on the stimulus page for >20 s on >5% of trials (which could indicate lack of attention due to multitasking or slow browser speeds), (v) gave a reported perceived “nonsocial” response when chase directness = 1 for more than one-third of trials in either perspective (which could indicate random responding and/or failures of attention), (vi) (main experiment only) failed an attention check in the trait questionnaires or omitted >10% of responses on any trait questionnaire. The numbers of participants failing each of these exclusion criteria, along with the total remaining participants, are given for each experiment in Table S1.

Judgment

Bias in Judgments

We coded participants' binary responses as 0 for “moving independently” and 1 for “chasing” and performed the same analyses for all experiments. To examine if participants exhibited a response bias in either or both perspectives, we collapsed the binary choices across all chase directness levels and calculated the probability of reporting seeing a social interaction in each perspective (Fig. 2A, Fig. 3 top row, and Fig. S1A). As the objective probability of chasing was balanced by having the chase directness change from 0 (having no information of chasing) to 1 (having maximum information of chasing), we performed one-sample t-tests on participants’ probabilities against 0.5 (50%). We also performed paired t-tests to test for a difference in the degree of social bias across perspectives.

Linear Mixed-Effect Modelling

To examine how subjective social percepts tracked objective evidence for socialness across perspectives, we fit participants’ binary choices using a generalized linear mixed-effect model (GLME) with a binomial family and logit link function. The model was specified in R formula syntax as:

$$response \sim chase\ directness \times perspective\ \ +\ \ (1\ |\ participant)$$

where chase directness and perspective were entered as fixed effects, including their interaction; a random intercept per participant was included to account for baseline differences in response preference across individuals (Fig. 2C and Fig. S1C).

Psychometric Function Fitting

To further quantify sensitivity to social information, for each perspective separately, we fitted each participant's average response at each chase directness level with the following psychometric function:

$$P(x) = \gamma + (1 - \gamma - \lambda)\frac{1}{1+e^{-\beta(x-\alpha)}}$$

where x is the chase directness; γ and λ are the lower and upper asymptote of the curve, calculated based on participants' minimum and maximum probability of responses; α and β are the horizontal location and the rate of change of the curve, respectively (Fig. 2D and Fig. S1D).

Based on the fitted function, we calculated participants' point of subjective equality (PSE) as $PSE = P^{-1}(0.5)$, which represented how direct a chase needed to be for someone to switch their perception from nonsocial to social. We then performed paired t-tests on PSEs from the first- versus third-person perspective to test for perspective differences in sensitivity to social information (Fig. 2E, Fig. 3 bottom row, and Fig. S1E).

Confidence Rating

We coded participants' confidence ratings on a continuous scale from 0 ("not at all confident") to 100 ("very confident"). To examine whether confidence varied differently with objective evidence for socialness across perspectives, we fit the confidence rating with a linear mixed-effect model (LME) as:

$$confidence\ rating \sim poly\ (\ chase\ directness,\ 2\ ) \times perspective\ +\ (1\ |\ participant)$$

where linear and quadratic terms of chase directness, perspective, and their interactions were entered as fixed effects; a random intercept per participant was included to account for baseline differences in rating scale usage across individuals (Fig. 2B, Fig. S1B, and Fig. S2).

Response Time

To examine if response time changed as a function of chase directness level and/or perspectives, we fit the response times (ms) with the same LME as for confidence:

$$response\ time \sim poly\ (\ chase\ directness,\ 2\ ) \times perspective\ \ +\ \ (1\ |\ participant)$$

where linear and quadratic terms of chase directness, perspective, and their interactions were entered as fixed effects; a random intercept per participant was included to account for baseline differences in response time usage across individuals (Fig. 4A).

Then, to examine if response time changed as a function of binary responses ("chasing" vs. "moving independently") and/or perspectives, we fit the response times (ms) with another LME as:

$$response\ time \sim response \times perspective\ \ +\ \ (1\ |\ participant)$$

where response, perspective, and their interactions were entered as fixed effects; a random intercept per participant was included to account for baseline differences in response time across individuals (Fig. S4).

Drift Diffusion Modelling

We first used the hierarchical drift-diffusion modeling (HDDM) package (73, 74) in Python to perform initial model comparison (Fig. S5), as it allowed explicitly modeling individual differences by estimating participant-level parameters as samples from a group-level distribution, leading to more robust results. We tested six models: (i) baseline (no perspective effects), (ii) only threshold varies by perspective;, (iii) only bias varies by perspective, (iv) only drift rate varies by perspective, (v) only nondecision time varies by perspective, and (vi) all parameters vary by perspective (full). All models were specified as hierarchical group models, meaning that parameters were estimated for each participant, which were drawn from group-level distributions. To account for possible lapses in attention, we included an outlier model, assuming that 5% of trials (p_outlier=0.05) were generated from a uniform distribution spanning the range of observed response times rather than from the DDM process. We then calculated the Deviance Information Criterion (DIC) to evaluate the goodness of the model. Lower DIC values indicate better model fit.

Then for individual-participant fits, we used the PyDDM package (75) in Python to perform drift diffusion modelling. As we found that the full model performed best when modelling all participants, we fitted each participant's response and response time data in each perspective with a full-model DDM and extracted the four free parameters of the model: threshold, nondecision time, bias and drift rate. We performed paired t-tests on these parameters to investigate the possible mechanism that drives the observed perspective differences (Fig. 4E).

To further investigate the relationship between evidence accumulation speed and sensitivity measured via psychometric function fitting, we correlated the drift rate with participants' PSE (Fig. 4F).

Movement Trajectories

In first-person trials, the (x,y) locations within the scene of the predator and prey dots were recorded at 60Hz sampling rate. The relative heading, the degree to which the participant-controlled prey dot was moving directly towards the possible predator's current location, was calculated at each time step as $relative\ heading\ =\ \frac{\phi_{prey} - \phi_{predator \rightarrow prey}}{\pi}$. Therefore, a relative heading equal to 1 means the prey dot was heading directly towards the predator's current location, while a relative heading equal to 0 means the prey dot was heading directly away (Fig. 5A).

We then calculated the average relative heading within each trial and tested for any relationship between movement patterns (indexed by relative heading) and perceived socialness (i.e., response) across objective evidence for socialness using LME as:

$$relative\ heading \sim chase\ directness \times response\ +\ (1\ |\ participant)$$

where directness, binary response, and their interactions were entered as fixed effects; a random intercept per participant was included to account for baseline differences in relative heading across individuals (Fig. 5B).

To further investigate the increased relative heading values when perceived "social" in temporal domain, we fitted the relative heading using a LME as:

$$relative\ heading \sim chase\ directness \times\ response \times time\ +\ (1\ |\ participant)$$

where directness, binary response, time, and their interactions were entered as fixed effects; a random intercept per participant was included to account for baseline differences in relative heading across individuals (Fig. S6).

To examine the relationship between information-seeking behavior and sensitivity to social signals, we calculated each participant's mean relative heading collapsed across all first-person trials and correlated it with their PSE (Fig. 5C).

To examine whether information-seeking behavior and drift rate independently predicted PSE in the first-person perspective and/or differences in PSE between perspectives, we conducted two linear regressions:

$$PSE_{1PP} \sim drift\ rate_{1PP} + mean\ relative\ heading_{1PP}$$

$$PSE_{diff.} \sim drift\ rate_{diff.} + mean\ relative\ heading_{1PP}$$

where $PSE_{diff.} = PSE_{3PP}$ - PSE_{1PP} and $drift\ rate_{diff.} = drift\ rate_{1PP}$ - $drift\ rate_{3PP}$. All predictors were entered in a single step to assess their unique contributions.

Code and Data Availability

All stimuli, data and code will be available upon publication on GitHub.

Acknowledgments

We acknowledge grant R01MH129648 (E.S.F.) that supports this work. We thank Tory Benson, Dallas Brodersen, Lenny Perez, Walker Ball and Kalina Duncheva for their help with stimuli quality-controlling and experiment piloting; and Luke Chang, Thomas Botch, Clara Sava-Segal, Kathryn O'Nell, and Zachary Paris for their thoughtful comments during the development of the project and the manuscript.

References

1. F. Heider, M. Simmel, An Experimental Study of Apparent Behavior. *Am. J. Psychol.* **57**, 243–259 (1944).

2. F. Abell, F. Happé, U. Frith, Do triangles play tricks? Attribution of mental states to animated shapes in normal and abnormal development. *Cogn. Dev.* **15**, 1–16 (2000).

3. A. Klin, Attributing Social Meaning to Ambiguous Visual Stimuli in Higher-functioning Autism and Asperger Syndrome: The Social Attribution Task. *J. Child Psychol. Psychiatry* **41**, 831–846 (2000).

4. B. J. Scholl, P. D. Tremoulet, Perceptual causality and animacy. *Trends Cogn. Sci.* **4**, 299–309 (2000).

5. J. Zwickel, S. J. White, D. Coniston, A. Senju, U. Frith, Exploring the building blocks of social cognition: spontaneous agency perception and visual perspective taking in autism. *Soc. Cogn. Affect. Neurosci.* **6**, 564–571 (2011).

6. C. E. Rasmussen, Y. V. Jiang, Judging social interaction in the Heider and Simmel movie. *Q. J. Exp. Psychol.* **72**, 2350–2361 (2019).

7. R. S. Varrier, E. S. Finn, Seeing Social: A Neural Signature for Conscious Perception of Social Interactions. *J. Neurosci.* **42**, 9211–9226 (2022).

8. E. McMahon, L. Isik, Seeing social interactions. *Trends Cogn. Sci.* **0** (2023).

9. T. Vestner, S. P. Tipper, T. Hartley, H. Over, S.-A. Rueschemeyer, Bound together: Social binding leads to faster processing, spatial distortion, and enhanced memory of interacting partners. *J. Exp. Psychol. Gen.* **148**, 1251–1268 (2019).

10. J. Yin, H. Xu, J. Duan, M. Shen, Object-Based Attention on Social Units: Visual Selection of Hands Performing a Social Interaction. *Psychol. Sci.* **29**, 1040–1048 (2018).

11. L. Papeo, N. Goupil, S. Soto-Faraco, Visual Search for People Among People. *Psychol. Sci.* **30**, 1483–1496 (2019).

12. Y. Fu, M. Zhou, J. Zhou, M. Shen, H. Chen, Unconscious prioritization for face-to-face people. *J. Exp. Psychol. Gen.* **153**, 1268–1280 (2024).

13. L. Papeo, T. Stein, S. Soto-Faraco, The Two-Body Inversion Effect. *Psychol. Sci.* **28**, 369–379 (2017).

14. L. Papeo, E. Abassi, Seeing social events: The visual specialization for dyadic human–human interactions. *J. Exp. Psychol. Hum. Percept. Perform.* **45**, 877–888 (2019).

15. X. Ding, Z. Gao, M. Shen, Two Equals One: Two Human Actions During Social Interaction Are Grouped as One Unit in Working Memory. *Psychol. Sci.* **28**, 1311–1320 (2017).

16. L. Schilbach, *et al.*, Toward a second-person neuroscience. *Behav. Brain Sci.* **36**, 393–414 (2013).

17. E. Redcay, L. Schilbach, Using second-person neuroscience to elucidate the mechanisms of social interaction. *Nat. Rev. Neurosci.* **20**, 495–505 (2019).

18. M. Emirbayer, A. Mische, What Is Agency? *Am. J. Sociol.* **103**, 962–1023 (1998).

19. M. Liljeholm, Agency and goal-directed choice. *Curr. Opin. Behav. Sci.* **41**, 78–84 (2021).

20. D. B. Markant, T. M. Gureckis, Is it better to select or to receive? Learning via active and passive hypothesis testing. *J. Exp. Psychol. Gen.* **143**, 94–122 (2014).

21. J. Code, Agency for Learning: Intention, Motivation, Self-Efficacy and Self-Regulation. *Front. Educ.* **5** (2020).

22. X. Li, *et al.*, Agency personalizes episodic memories. [Preprint] (2024). Available at: https://osf.io/preprints/psyarxiv/7evwj_v1/ [Accessed 12 January 2026].

23. W. Yu, B. Silver, K. Ochsner, L. Davachi, Causality and agency shape memory for a real-life dynamic social experience. [Preprint] (2025). Available at: https://osf.io/preprints/psyarxiv/zprmw_v2/ [Accessed 12 January 2026].

24. L. I. Rait, E. A. Horwath, S. DuBrow, V. P. Murty, Agency alters memory organization during free recall. *Learn. Mem.* **33**, a054121 (2026).

25. C. S. Symons, B. T. Johnson, The self-reference effect in memory: A meta-analysis. *Psychol. Bull.* **121**, 371–394 (1997).

26. J. Sui, G. W. Humphreys, The Integrative Self: How Self-Reference Integrates Perception and Memory. *Trends Cogn. Sci.* **19**, 719–728 (2015).

27. T. E. Horgan, J. L. Tienson, G. Graham, “The Phenomenology of First-Person Agency” in *Physicalism and Mental Causation: The Metaphysics of Mind and Action*, S. Walter, H.-D. Heckmann, Eds. (Imprint Academic, 2003), p. 323.

28. T. Fuchs, The phenomenology and development of social perspectives. *Phenomenol. Cogn. Sci.* **12**, 655–683 (2013).

29. H. G. Debarba, *et al.*, Characterizing first and third person viewpoints and their alternation for embodied interaction in virtual reality. *PLOS ONE* **12**, e0190109 (2017).

30. P. W. Blythe, P. M. Todd, G. F. Miller, “How motion reveals intention: Categorizing social interactions” in *Simple Heuristics That Make Us Smart*, Evolution and cognition., (Oxford University Press, 1999), pp. 257–285.

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

32. T. Gao, G. McCarthy, B. J. Scholl, The Wolfpack Effect: Perception of Animacy Irresistibly Influences Interactive Behavior. *Psychol. Sci.* **21**, 1845–1853 (2010).

33. S. Castiello, J. D. K. Ongchoco, B. van Buren, B. J. Scholl, P. R. Corlett, Paranoid and teleological thinking give rise to distinct social hallucinations in vision. *Commun. Psychol.* **2**, 117 (2024).

34. S. Quadflieg, K. Koldewyn, The neuroscience of people watching: how the human brain makes sense of other people’s encounters. *Ann. N. Y. Acad. Sci.* **1396**, 166–182 (2017).

35. J. Zaki, K. Ochsner, The Need for a Cognitive Neuroscience of Naturalistic Social Cognition. *Ann. N. Y. Acad. Sci.* **1167**, 16–30 (2009).

36. R. S. Varrier, *et al.*, Shared and individual tuning curves for social perception. [Preprint] (2025). Available at: https://www.biorxiv.org/content/10.1101/2025.01.19.633772v3 [Accessed 26 December 2025].

37. H. De Jaegher, E. Di Paolo, S. Gallagher, Can social interaction constitute social cognition? *Trends Cogn. Sci.* **14**, 441–447 (2010).

38. T. Froese, S. Gallagher, Getting interaction theory (IT) together: Integrating developmental, phenomenological, enactive, and dynamical approaches to social interaction. *Interact. Stud.* **13**, 436–468 (2012).

39. K. Lehmann, *et al.*, An Active-Inference Approach to Second-Person Neuroscience. *Perspect. Psychol. Sci.* **19**, 931–951 (2024).

40. R. Hari, L. Henriksson, S. Malinen, L. Parkkonen, Centrality of Social Interaction in Human Brain Function. *Neuron* **88**, 181–193 (2015).

41. A. Gomez-Marin, A. A. Ghazanfar, The Life of Behavior. *Neuron* **104**, 25–36 (2019).

42. E. F. Risko, K. E. Laidlaw, M. Freeth, T. Foulsham, A. Kingstone, Social attention with real versus reel stimuli: toward an empirical approach to concerns about ecological validity. *Front. Hum. Neurosci.* **6** (2012).

43. K. E. W. Laidlaw, T. Foulsham, G. Kuhn, A. Kingstone, Potential social interactions are important to social attention. *Proc. Natl. Acad. Sci.* **108**, 5548–5553 (2011).

44. T. Foulsham, E. Walker, A. Kingstone, The where, what and when of gaze allocation in the lab and the natural environment. *Vision Res.* **51**, 1920–1931 (2011).

45. J. H. Cheong, Z. Molani, S. Sadhukha, L. J. Chang, Synchronized affect in shared experiences strengthens social connection. *Commun. Biol.* **6**, 1099 (2023).

46. K. Rice, E. Redcay, Interaction matters: A perceived social partner alters the neural processing of human speech. *NeuroImage* **129**, 480–488 (2016).

47. J. Su, J. J. A. van Boxtel, H. Lu, Social Interactions Receive Priority to Conscious Perception. *PLOS ONE* **11**, e0160468 (2016).

48. E. Redcay, *et al.*, Live face-to-face interaction during fMRI: A new tool for social cognitive neuroscience. *NeuroImage* **50**, 1639–1647 (2010).

49. E. Bilek, *et al.*, Information flow between interacting human brains: Identification, validation, and relationship to social expertise. *Proc. Natl. Acad. Sci.* **112**, 5207–5212 (2015).

50. E. Redcay, K. R. Velnoskey, M. L. Rowe, Perceived communicative intent in gesture and language modulates the superior temporal sulcus. *Hum. Brain Mapp.* **37**, 3444–3461 (2016).

51. C. A. Hill, *et al.*, A causal account of the brain network computations underlying strategic social behavior. *Nat. Neurosci.* **20**, 1142–1149 (2017).

52. D. Pitcher, L. G. Ungerleider, Evidence for a Third Visual Pathway Specialized for Social Perception. *Trends Cogn. Sci.* **25**, 100–110 (2021).

53. E. Abassi, *et al.*, Neural Encoding of Bodies for Primate Social Perception. *J. Neurosci.* **44** (2024).

54. A. Netanyahu, T. Shu, B. Katz, A. Barbu, J. B. Tenenbaum, Phase: Physically-grounded abstract social events for machine social perception in *Proceedings of the Aaai Conference on Artificial Intelligence*, (2021), pp. 845–853.

55. K. Lehmann, L. Maliske, A. Böckler, P. Kanske, Social Impairments in Mental Disorders: Recent Developments in Studying the Mechanisms of Interactive Behavior. *Clin. Psychol. Eur.* **1**, 1–15 (2019).

56. E. A. H. von dem Hagen, R. S. Stoyanova, J. B. Rowe, S. Baron-Cohen, A. J. Calder, Direct Gaze Elicits Atypical Activation of the Theory-of-Mind Network in Autism Spectrum Conditions. *Cereb. Cortex* **24**, 1485–1492 (2014).

57. M. Freeth, T. Foulsham, A. Kingstone, What Affects Social Attention? Social Presence, Eye Contact and Autistic Traits. *PLOS ONE* **8**, e53286 (2013).

58. A. Myllyneva, K. Ranta, J. K. Hietanen, Psychophysiological responses to eye contact in adolescents with social anxiety disorder. *Biol. Psychol.* **109**, 151–158 (2015).

59. V. S. Chib, R. Adachi, J. P. O'Doherty, Neural substrates of social facilitation effects on incentive-based performance. *Soc. Cogn. Affect. Neurosci.* **13**, 391–403 (2018).

60. S. Liu, T. D. Ullman, J. B. Tenenbaum, E. S. Spelke, Ten-month-old infants infer the value of goals from the costs of actions. *Science* **358**, 1038–1041 (2017).

61. S. Yasuda, W. Li, D. Martinez, B. M. Lake, M. R. Dillon, 15-Month-Olds' Understanding of Imitation in Social and Instrumental Contexts. *Infancy* **30**, e70002 (2025).

62. T. Shu, Y. Peng, L. Fan, H. Lu, S.-C. Zhu, Perception of Human Interaction Based on Motion Trajectories: From Aerial Videos to Decontextualized Animations. *Top. Cogn. Sci.* **10**, 225–241 (2018).

63. T. Shu, Y. Peng, S.-C. Zhu, H. Lu, A unified psychological space for human perception of physical and social events. *Cognit. Psychol.* **128**, 101398 (2021).

64. A. Salatiello, M. Hovaidi-Ardestani, M. A. Giese, A Dynamical Generative Model of Social Interactions. *Front. Neurorobotics* **15**, 648527 (2021).

65. K. Tsutsui, R. Tanaka, K. Takeda, K. Fujii, Collaborative hunting in artificial agents with deep reinforcement learning. *eLife* **13**, e85694 (2024).

66. S. Baron-Cohen, S. Wheelwright, R. Skinner, J. Martin, E. Clubley, The Autism-Spectrum Quotient (AQ): Evidence from Asperger Syndrome/High-Functioning Autism, Males and Females, Scientists and Mathematicians. *J. Autism Dev. Disord.* **31**, 5–17 (2001).

67. D. Watson, L. A. Clark, A. Tellegen, Development and validation of brief measures of positive and negative affect: The PANAS scales. *J. Pers. Soc. Psychol.* **54**, 1063–1070 (1988).

68. P. T. Costa, R. R. McCrae, *Revised NEO personality inventory (NEO PI-R) and NEP five-factor inventory (NEO-FFI) : professional manual* (Psychological Assessment Resources, 1992).

69. H. Hoekstra, J. Ormel, F. De Fruyt, NEO Personality Questionnaires NEO-PI-R, NEO-FFI, Manual. *Lisse Swet Zeitlinger BV* (1996).

70. D. W. Russell, UCLA Loneliness Scale (Version 3): reliability, validity, and factor structure. *J. Pers. Assess.* **66**, 20–40 (1996).

71. D. Freeman, *et al.*, The revised Green et al., Paranoid Thoughts Scale (R-GPTS): psychometric properties, severity ranges, and clinical cut-offs. *Psychol. Med.* **51**, 244–253 (2021).

72. M. Lindeman, K. Aarnio, Superstitious, magical, and paranormal beliefs: An integrative model. *J. Res. Personal.* **41**, 731–744 (2007).

73. T. V. Wiecki, I. Sofer, M. J. Frank, HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Front. Neuroinformatics* **7** (2013).

74. W. Pan, *et al.*, dockerHDDM: A User-Friendly Environment for Bayesian Hierarchical Drift-Diffusion Modeling. *Adv. Methods Pract. Psychol. Sci.* **8**, 25152459241298700 (2025).

75. M. Shinn, N. H. Lam, J. D. Murray, A flexible framework for simulating and fitting generalized drift-diffusion models. *eLife* **9**, e56938 (2020).

Supporting Information

Table S1. Details of data exclusion

Exclusion Criteria*	Main Experiment		Identity -Known Experiment (n = 50 recruited)	Response Time Experiment	
	Session 1 (n = 350 recruited)	Session 2 (n = 119 out of 268 returned)		Batch 1 (n = 60 recruited)	Batch 2 (n = 60 recruited)
Incomplete data	1	0	1	3	0
Missed more than 5% of experimental trials	28	8	1	6	7
Failed the attention check question about cover story	2	0	0	3	3
Stayed on stimuli page for longer than 20 seconds for more than 5% of experimental trial	2	0	0	0	0
Failed to report seeing interaction at chase directness equaled 1 for more than ⅓ of the trials in either perspective	15	5	5	6	8
Failed trait questionnaires' attention check or missed more than 10% responses of any trait questionnaire	56	NA	NA	NA	NA
Final valid n	268	108	43	47	46

* Note that participants might fail multiple exclusion criteria

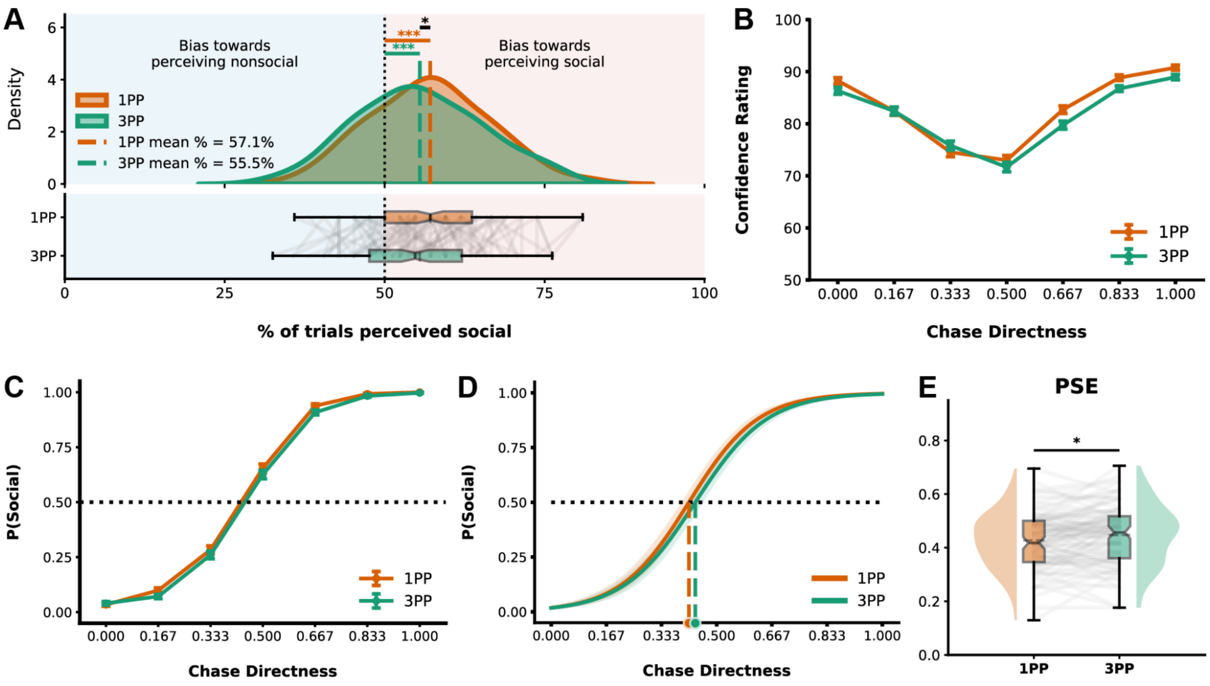

Fig. S1. Main experiment session 2 replicated the increased social biases, sensitivity and confidence in active scenarios. (A) Distribution of two-alternative forced choice responses. P(Social): 1PP vs. 0.5, *t*(107) = 7.956, *p* < .001; 3PP vs. 0.5, *t*(107) = 5.785, *p* < .001; 1PP vs. 3PP, *t*(107) = 2.418, *p* = .017. (B) Confidence ratings across chase directness levels. Fixed effect of perspective: *β* = -1.299, *SE* = 0.318, *t* = -4.089, *p* < .001. (C) Two-alternative forced choice responses across chase directness levels. (D) Psychometric curves with perceptual thresholds. (E) PSE comparison. *t*(107) = -2.376, *p* = .019. **p* < .05, ***p* < .01, ****p* < .001.

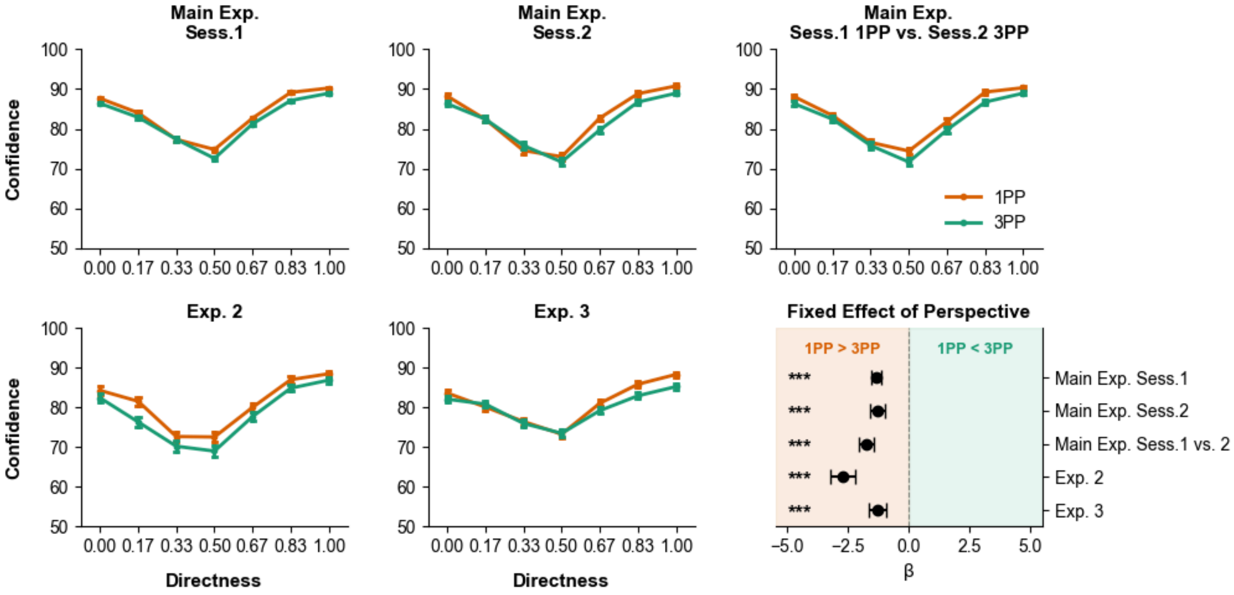

Fig. S2. Increased confidence during active participation persisted across control experiments. Left and center: Mean confidence ratings plotted as a function of chase directness for 1PP (orange) and 3PP (green) conditions across five experimental comparisons. From top left to bottom right: Main Experiment Session 1 (n = 268), Main Experiment Session 2 (n = 108), Main Experiment Session 1 vs. Session 2 (matched visual input, n = 108), Experiment 2 (matched uncertainty, n = 43), and Experiment 3 (controlled for cognitive load, n = 93). Bottom right: Main effect of perspective on confidence in LME. Error bars indicate the standard errors. Negative values (orange shaded region) indicate higher confidence in 1PP; positive values (green shaded region) indicate higher confidence in 3PP. ***$p < .001$.

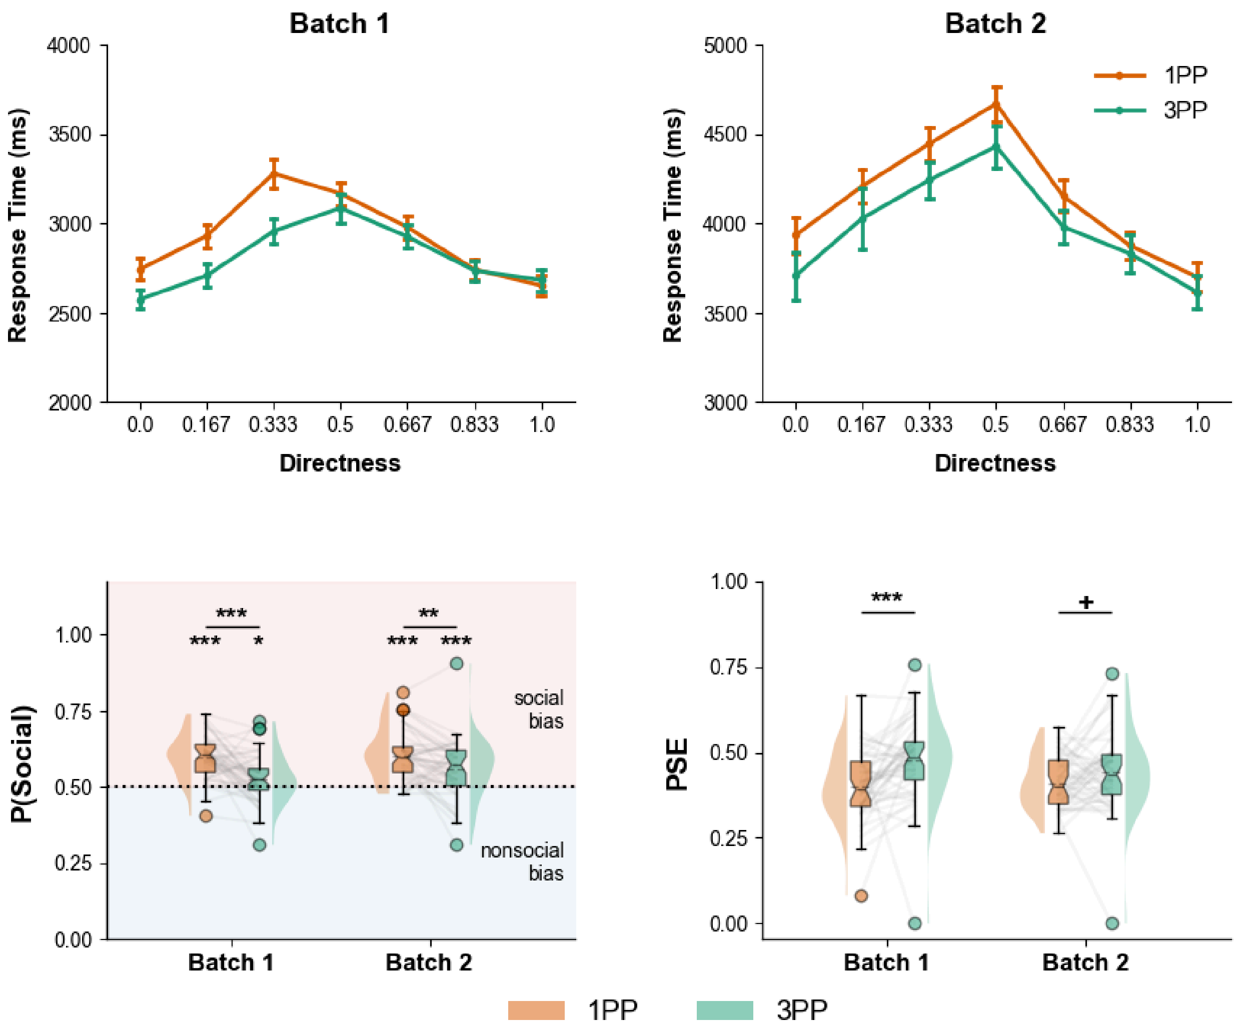

Fig. S3. Response time patterns and perspective differences were similar across batches. Top row: mean response times (ms) plotted as a function of chase directness for 1PP (orange) and 3PP (green) conditions in Batch 1 (left, n = 47, with "didn't response" message) and Batch 2 (right, n = 46, without message). Error bars represent standard error. Bottom row: social bias and sensitivity for both perspectives in two batches. $^{+}p < .1$, $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$.

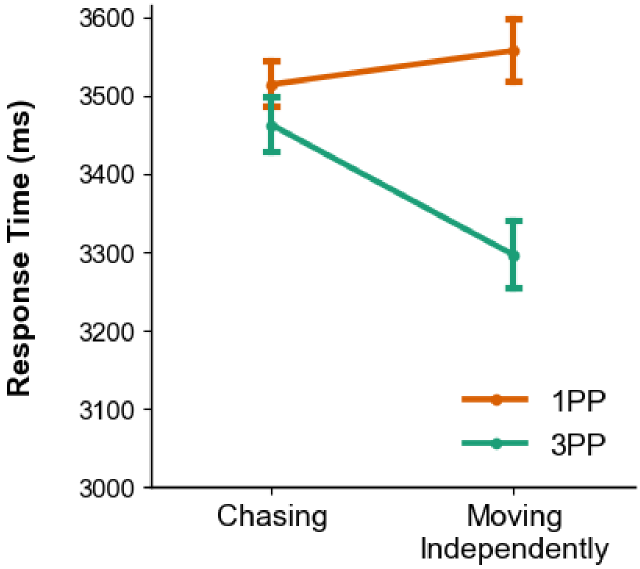

Fig. S4. No significant difference between response times for the two binary choices in self-pace response time experiment (n = 93). Mean response times (ms) for trials judged as "chasing" versus "moving independently" in 1PP (orange) and 3PP (green) conditions. Error bars represent standard errors.

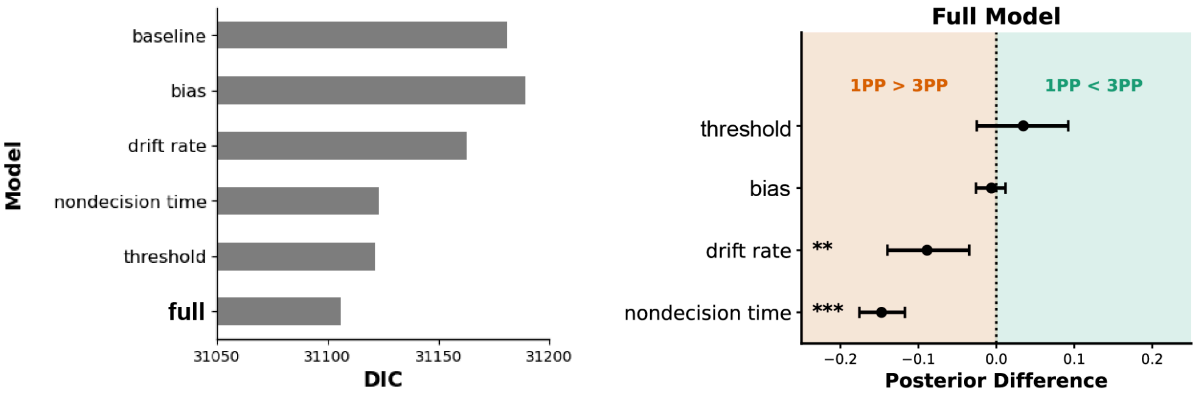

Fig. S5. Hierarchical drift diffusion model comparison showed the full model (i.e. allowing perspective to affect all parameters) performed the best. Left: model comparison results. Models tested from top to bottom: baseline (no perspective effects); drift rate, bias nondecision time, and threshold affected by perspective respectively; and full model (all parameters vary by perspective). Lower Deviance Information Criterion (DIC) indicates better model fit. The full model shows the lowest DIC (3.109 x 10^4). Right: posterior differences for each parameter in the full model. Error bars indicate 95% CI. Results showed the same longer nondecision time and higher drift rate in 1PP trials effect as revealed by individual fits (Fig. 4E).

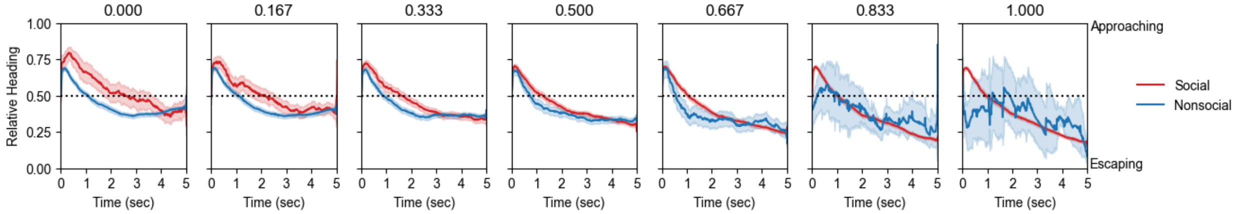

Fig. S6. The relative heading direction of the participant-controlled prey evolves over time across chase directness levels. The red line is the average over trials ultimately perceived as social; the blue line is the average over trials ultimately perceived as non-social. The difference in relative heading between response types emerges relatively early in the trial. Shaded regions represent standard error across participants.

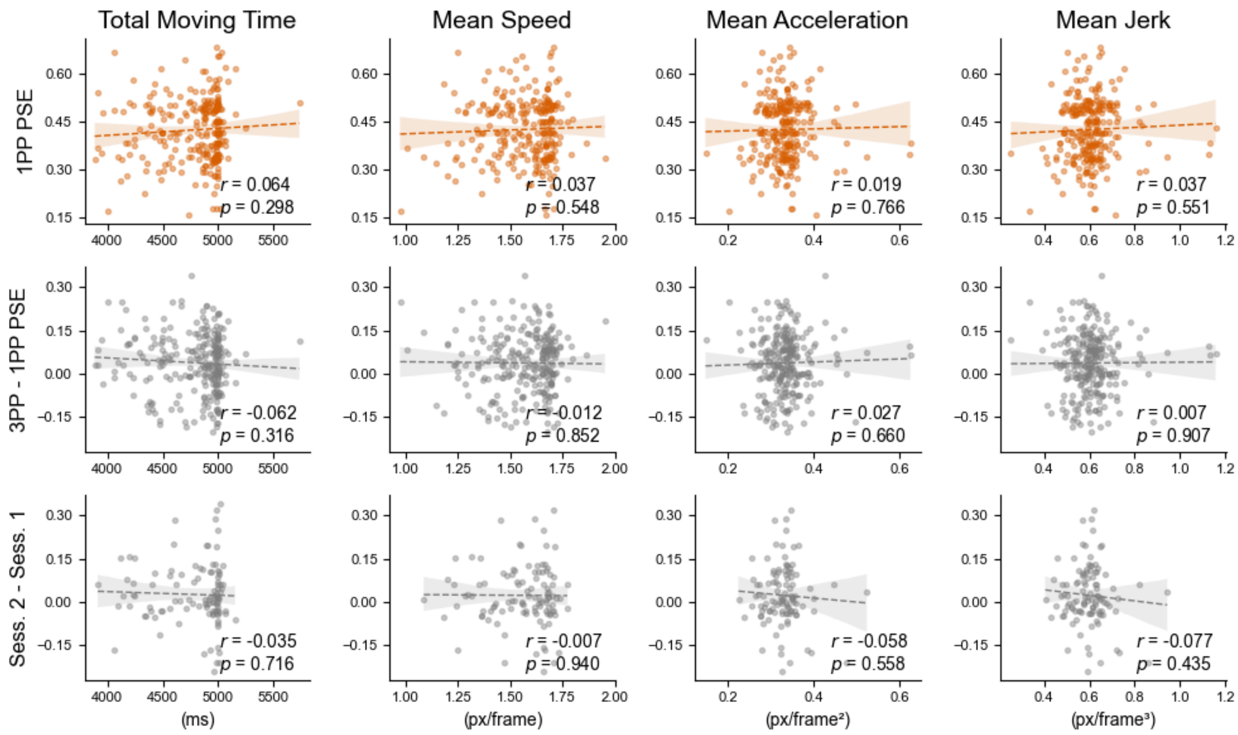

Fig. S7. Sensitivity to social interactions was not correlated with simple prey trajectory measures. Scatter plots showed correlations between the point of subjective equality (PSE) and four trajectory variables. Rows from top to bottom show: (1) session 1 1PP PSE (n = 268), (2) session 1 3PP minus 1PP PSE (n = 268), (3) session 2 3PP PSE minus session 1 1PP PSE (n = 108). Columns from left to right show: Total Moving Time (total time participants spent on moving the prey dot in ms), Mean Speed (average velocity), Mean Acceleration (average rate of velocity change), and Mean Jerk (average rate of acceleration change).

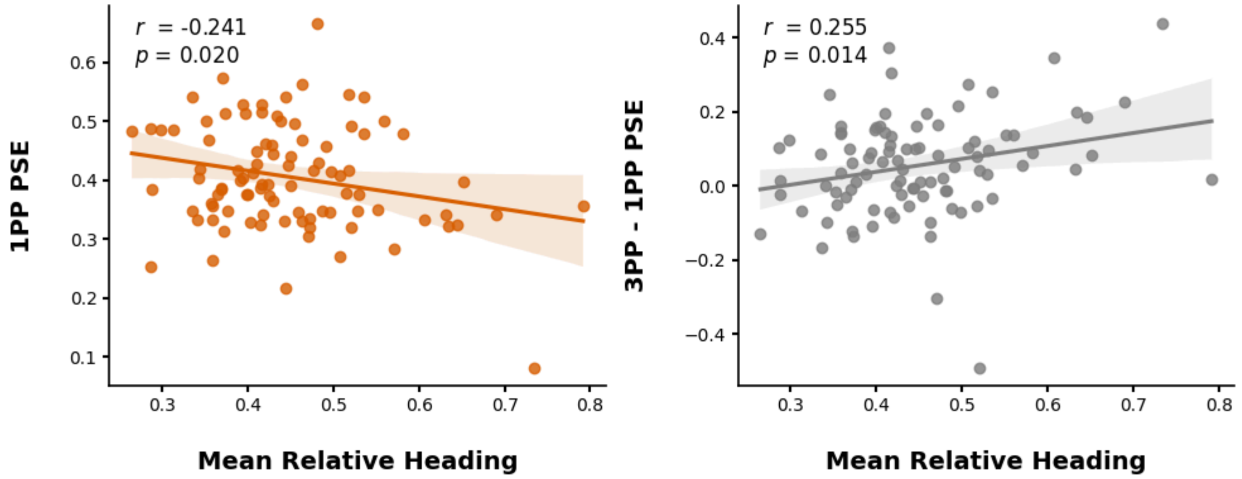

Fig. S8. The relationship between information seeking behavior and sensitivity was replicated in the response time experiment. Scatter plots examining correlations between 1PP PSE (left) / 3PP minus 1PP PSE (right) and mean relative heading (n = 93).