Curiosity is linked to information seeking in the moral domain

Jane Acierno[1]; Nathan Liang[2]; Isaac Hadley-Miner[2]; Sara Constantino[3]; Max Kleiman-Weiner[4]; Liane Young[2]; Jordan Wylie[2]

[1] Northeastern University; [2] Boston College; [3] Stanford University; [4] University of Washington

Abstract

Curiosity motivates information seeking and knowledge acquisition across the lifespan, yet its role in the moral domain remains unexplored. To address this gap, we developed a task to examine how people seek information about real-life ambiguous moral scenarios. Across two experiments, we found that curiosity predicts greater information seeking, regardless of whether the action was seen as morally good or bad. In Experiment 1, we also found that moral goodness was associated with less information seeking, suggesting an asymmetry in how people engage with moral information. Experiment 2 replicated these findings and showed reputational concerns also influenced information-seeking, and that judgments of action justification help explain the link between perceived morality and information-seeking behavior. Together, these results demonstrate the role of curiosity in moral information seeking, and highlight the influence of perceived morality, justification, and reputational concerns in this process.

Keywords: curiosity; morality; information seeking

Introduction

Curiosity is a universal human drive that plays a central role in our cognition—it is often referred to as the "impulse towards better cognition" (James, 1899/1983; Kidd & Hayden, 2015). It is piqued by salient information gaps (Berlyne, 1966; Loewenstein, 1994) and perceived information utility (Liquin & Lombrozo, 2020; Dubey, Griffiths, & Lombrozo, 2022), and it often motivates knowledge gathering for its own sake (see Poli et al., 2024 for review). In other words, curiosity fuels the desire to explore, discover, and make sense of one's surroundings and oneself, including when information seeking is non-instrumental (see e.g., Kobayashi et al., 2019; Sharot & Sunstein, 2020; van Lieshout, de Lange, & Cools, 2019). While a substantial body of research has focused on identifying the triggers and outcomes of curiosity in domains like learning, problem-solving, and decision making, its role and elicitors within the moral domain remain underexplored (but see Wylie & Gantman, 2023).

This notable gap is surprising given the centrality of morality in human cognition and behavior. Moral reasoning underpins decisions ranging from who we trust (e.g., Brambilla et al., 2011; Everett, Pizarro, & Crockett, 2016) to how we navigate social dilemmas (e.g., Awad et al., 2018; Gawronski et al., 2017; Greene et al., 2004)—where the stakes often involve both personal and societal consequences. For example, having a good reputation affords access to important resources (see e.g., Sperber & Baumard, 2012) and is critical to maintaining cooperative social interactions (Boyd & Richerson, 1989; Sylwester & Roberts, 2010). The costs and benefits associated with knowing 'the right thing to do' or 'the right person to trust' tend to be high (Rand & Nowak, 2013).

Despite its central role in everyday life, ambiguity pervades the moral domain. Moral dilemmas emerge everywhere from the workplace (LaMontagne, 2016) and interpersonal relationships (Zhan et al., 2018), to larger societal challenges like how best to program autonomous vehicles (Awad et al., 2018). Though this ambiguity can sometimes present a challenge, it can also spark curiosity and motivate engagement, especially about other people. For example, popular television shows like *The Sopranos* feature morally complex protagonists who often do both good and bad things. Viewers are drawn in not in spite of this ambiguity, but because of it (see Krause & Rucker, 2020; Wylie & Gantman, 2022; 2023). Such examples suggest that, despite the seeming objectivity (e.g., Goodwin & Darley, 2008) or black and white nature (e.g., Turiel, 1998) of the moral domain, people are frequently curious about moral ambiguity. Paradoxically, this lack of objectivity may make moral discourse particularly engaging, as an absence of clear answers can allow individuals to explore alternative viewpoints without necessarily needing to resolve them.

While curiosity reliably drives learning in many areas, its impact on integrating moral information may be more complex. Consider the difference between learning when the dinosaurs went extinct—a question of historical fact—and evaluating whether an action is morally justified. The former invites relatively straightforward belief updating in response to new evidence. In contrast, integrating new moral information can be challenging (see Skitka et al., 2021); deeply held values are often at stake, and beliefs can have significant social and personal repercussions. Moreover, moral questions frequently involve evaluating intentions, actions, and outcomes under conditions of incomplete information (Cushman, Young, & Hauser, 2006). As such, the stakes of forming and updating moral beliefs are often substantially higher than updating beliefs about non-social information, and they are uniquely powerful compared to other social information (e.g., sociability, competence; see Brambilla et al., 2019).

Additionally, although curiosity often fuels the pursuit of knowledge, it does not always translate directly into active information seeking (Sharot & Sunstein, 2020). For instance, someone may be deeply curious about their genetic predisposition to a disease but choose to avoid medical test results due to fear of their implications (Horn, Litovsky, & Loewenstein, 2024). Conversely, people often seek mundane information without experiencing much curiosity, such as checking a grocery store's hours of operation (see Sharot & Sunstein, 2020 for additional examples). This distinction becomes particularly salient in the moral domain, where the interplay between curiosity and information seeking is influenced by factors such as personal beliefs, reputational

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

concerns, and the risk of cognitive dissonance. Some individuals may be curious about morally ambiguous situations or agents, but hesitate to engage with new information if it threatens their existing moral beliefs. Others may seek moral information not out of curiosity, but to reinforce preexisting views or manage their social reputation. Understanding how curiosity translates, or fails to translate, into information seeking in moral contexts is therefore critical to understanding the cognitive and motivational processes underlying moral reasoning.

In this paper, we develop and implement a novel paradigm to examine the links between curiosity and information seeking in the moral domain. Specifically, we explore how people engage with ambiguous moral agents, and whether the motivation to acquire information is influenced by the perceived morality of ambiguous moral agents. In Experiment 1, we first test whether curiosity about ambiguous agents influences information seeking, and we explore the role that personal moral judgments and uncertainty play. In Experiment 2, we replicate and extend the findings of Experiment 1 to examine whether beliefs about action justification or concerns about one's reputation influence information seeking. All experiments were approved by the Boston College institutional IRB and were pre-registered on OSF.[1] All analyses were conducted using R programming software (R Core Team) and the *lme4* (Bates et al., 2015), *lmerTest* (Kuznetsova, Brockhoff, & Christensen, 2017), and *sjPlot* (Lüdecke, 2023) packages for running linear mixed-effects models.

Norming Study

To identify ambiguous moral claims for use in the main study, we collected norming data on the perceived morality and veracity for 80 potential claims, of which 39 were true and 41 were false. Ratings were collected for both true and false claims for use in a future study; however, for the present study, we selected only true claims to ensure that any measures related to confidence in judgments were attributable only to confidence about the morality of the claim, not its veracity.

Method

Participants We recruited a total of 403 English speaking U.S. participants from Connect, an Amazon Mechanical Turk recruitment platform (202 male, 200 female, 1 non-binary; 65.01% White; Mean age = 45.33, *SD* = 15.59).

Procedure Participants viewed a random subset of 20 (out of 80) claims, written by the study team, describing actions taken by prominent people such as celebrities or historical figures. Claims described an action taken by the person, and detailed one positive and one negative outcome from the action, making the claim morally ambiguous. For example, participants might see the claim "Akon created the 'Akon Lighting Africa' initiative, which supported authoritarian regimes but initiated renewable energy projects in Africa." For all claims, the valence order of outcomes was counterbalanced.

After viewing each claim, participants rated the extent to which they believed the claim to be morally good or bad using a 100-point sliding scale. Participants also rated their familiarity with the person on a 50-point sliding scale. We collected additional judgments related to politics, truth, and emotions to pilot for a separate project. Finally, participants completed demographic measures and were debriefed and compensated for their time.

Results

Of the 39 true claims, eight were selected based on having difference scores of less than 5 points between average morality and veracity ratings. This selection criterion supports a planned future study comparing morality- and veracity-driven information seeking. However, three of the claims were excluded because they received fewer than the 100 ratings required to construct the sampling grid in the main study. In total, we selected five claims rated slightly more moral than immoral (Mean morality = 63.42, *SD* = 10.12), and slightly more likely to be true than false (Mean veracity = 66.87, *SD* = 23.36). Additionally, participant familiarity with the people in the five claims was below the midpoint of the 50-point scale (Mean familiarity = 19.82, *SD* = 18.81).

Experiment 1

Experiment 1 sought to examine the links between curiosity and information seeking in the moral domain. Drawing on prior work in curiosity (Kidd & Hayden, 2015; Loewenstein, 1994; Sharot & Sunstein, 2020), we explored whether curiosity and prior beliefs influence information seeking and moral belief updating. We hypothesized that higher self-reported curiosity, and weaker priors (i.e., greater uncertainty), would predict greater information seeking and moral belief updating. We also explored whether personal moral judgments influenced subsequent information seeking behavior.

Method

Participants We recruited a nationally representative sample of 545 U.S. participants on Prolific, stratified by gender, race, and political ideology. An *a priori* power analysis using G*Power (Faul et al., 2007) indicated that a sample of 459 participants would provide 90% power to detect a correlation between curiosity and avatar selection of r = .15 (based on pilot data), with alpha set at .05 for a two-tailed test. To account for potential exclusions, we preregistered the recruitment of approximately 100 additional participants. After excluding individuals who failed attention checks or sampled more than 3 standard deviations above the mean in

[1] Study 1 pre-registration is available here: osf.io/xmj4s; Study 2 pre-registration is available here: osf.io/6szyu.

the sampling task, the final sample consisted of 525 participants (Mean age = 46.1, *SD* = 16.0, 65.90% White, 13.71% Black, 20.38% Other).

Information Seeking Task We developed a novel information seeking task in which participants viewed claims and had the opportunity to seek information about others' moral judgments about those claims. Participants were informed that they would view three claims that described actions taken by real historical figures. For each claim, they were asked to estimate the percentage of people in the U.S. who would judge the action to be morally good or bad. They could then view responses from a nationally representative sample—collected in a previous study—by clicking on different avatars. After each selection, participants viewed the moral judgment for a minimum of 3 seconds before being prompted to either continue sampling or stop. This process repeated until the participant chose to stop sampling. Our primary dependent variable was the number of unique avatars that participants clicked on, out of a grid of 100 possible avatars (see Figure 1). Thus, information seeking was a continuous, numeric variable, ranging from 1 to 100.

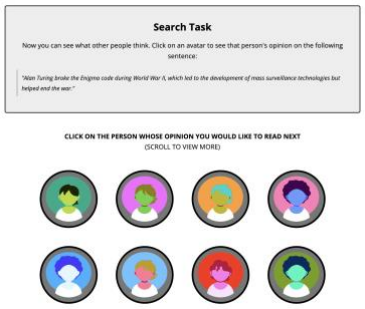

Figure 1: Example screen for the information seeking task. Each grid contained a total of 100 avatars.

Curiosity The primary independent variable of interest was self-reported curiosity. Curiosity was collected pre-sampling by asking participants to rate their curiosity to learn what other people think about each claim on a 100-point sliding scale with endpoints "*not at all curious*" and "*extremely curious*".

Moral Belief Updating To measure moral belief updating, participants were asked pre- and post-sampling to rate the morality of each action on 100-point sliding scales from "*definitely morally bad*" to "*definitely morally good*". Belief updates were then calculated as difference scores between the pre- and post-sampling morality ratings for each action.

Meta-perceptions As part of the study cover story, we asked participants to provide pre- and post-sampling meta-perception ratings estimating what percentage of people in the United States believe each action to be morally good (vs bad). Additionally, the strength of participants' meta-perceptions was measured by asking participants to report their confidence in their estimate on a 100-point sliding scale from "*not at all confident*" to "*completely confident*".

Behavioral Curiosity At the end of the study, we included an exploratory behavioral measure of curiosity (i.e., not based on self-report). Participants were given the option to view the average moral judgment rating (i.e., the average judgment from the representative sample of participants) for each of the three claims they saw during the sampling task. This average represents the average rating of the 100 avatars. For each claim, participants could choose whether or not to reveal the average rating. We used the total number of times participants viewed the average ratings as a behavioral indicator of curiosity about the claims.

Design & Procedure The study employed a fully within-subjects design. Each participant completed three randomly selected trials from a pool of five, presented in random order. At the beginning of each trial, participants viewed a single claim and reported their initial moral judgments about the agent and action, their curiosity about others' views, their meta-perception estimating the percentage of people in the U.S. who would view the action as morally good versus bad, and their confidence in that estimate. Participants then engaged in the information seeking task, during which they could sample as many avatars as they wished. Once they chose to stop sampling, participants again reported their moral judgments about the agent and action, their confidence in those judgments, and their meta-perception of how others would morally evaluate the action. Finally, participants completed demographic questions and were debriefed and compensated for their time.

Results

Our primary hypotheses centered around whether self-reported curiosity and prior beliefs predict information seeking. All models were linear mixed effects regressions with by-participant and by-stimulus random intercepts. For models in which the count of avatars was the outcome, we used a generalized linear mixed effects model (poisson distribution) and included random by-participant intercepts. Models that also include random slopes are noted. Bayesian hierarchical models were conducted for singular models using the same specifications.

We first investigated the role of curiosity: As hypothesized, curiosity predicted greater information seeking. The more curious people reported being about what others thought, the more information they sought, $b = 0.48$, $SE = 0.06$, $z = 7.93$, *95%CI* [0.36, 0.60], $p < .001$ (see Figure 2, left panel). We also hypothesized that higher curiosity ratings would predict greater personal moral belief updates and greater meta-perception (what the crowd thinks) updates in units of the change between pre- and post-sampling ratings, but found no evidence to support those predictions ($p = .28$ and $p = .38$, respectively).

Next, we investigated how confidence in one's beliefs predicted information seeking. Here again, our results

matched our predictions. The less confident people were, the more information they sought, b = -0.25, SE = 0.07, z = -3.63, *95%CI* [-0.39, -0.12], $p < .001$ (see Figure 2, right panel). We also hypothesized that weaker confidence ratings would predict greater personal belief updating and greater meta-perception updating. Results supported these predictions: the less confident people were, the more they updated their personal moral judgments, b = -5.92, SE = 1.11, t = -5.32, 95% CI [-8.11, -3.74], $p < .001$, and the more they updated their moral meta-perceptions, b = -11.20, SE = 1.33, t = -8.43, 95% CI [-13.81, -8.60], $p < .001$.

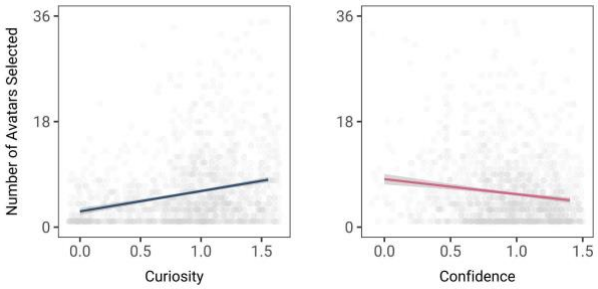

Figure 2: Left panel shows the relationship between curiosity and avatar selections. Right panel shows the relationship between confidence and avatar selections. Shading represents 95% CIs.

Critically, we also sought to examine whether initial moral action ratings affected information seeking. Results suggested that moral goodness ratings decreased information seeking: the more morally good participants thought a particular action was, the less they sought information about what others thought about it, b = -0.23, SE = 0.04, z = -5.52, *95%CI* [-0.32, -0.15], $p < 0.001$.

We also explored whether the information participants encountered influenced their personal moral beliefs or their meta-perceptions by analyzing changes in these judgments. Results provided evidence that the task affected participants' evaluations. Specifically, greater information seeking was associated with more negative changes in meta-perceptions, b = -0.13, SE = 0.05, t = -2.64, *95%CI* [-0.23, -0.03], p = 0.008. That is, the more information people sought, the more they reduced their estimate of how moral others perceived the action to be. In contrast, information seeking was not significantly associated with changes in participants' own moral judgments, although the trend was in the same direction.

Finally, we explored whether self-reported curiosity predicted our behavioral curiosity measure. Indeed, we found that greater self-reported curiosity significantly predicted the behavioral measure of looking at the average ratings more times, b = 0.19, SE = 0.06, z = 2.96, 95%CI [0.07, 0.32], $p < .001$.

Study 1 Discussion

Overall, our findings suggest moral curiosity—curiosity toward moral content—influences information seeking. People were more motivated to gather information about moral agents they were curious about, even in an effortful search task that offered no financial incentives. These results align with prior research showing that curiosity serves as a motivational driver for information seeking behavior (e.g., Gottlieb et al., 2013), and they suggest that similar mechanisms undergird information seeking both inside and outside of the moral domain. That is, here we see that both self-reported curiosity and uncertainty motivate information seeking about moral ambiguity.

However, these findings also point to some potentially unique aspects of curiosity in the moral domain. While past research has highlighted a negativity bias (e.g., Rozin & Royzman, 2001; Vaish, Grossmann, & Woodward, 2008), other studies have found that people often avoid negative information in favor of positive information (Marvin & Shohamy, 2016; van Lieshout, de Lange, & Cools, 2020). Here, we show that people are more active in seeking out negative moral information. Notably, while curiosity predicted increased information seeking, it did not predict significant updates in participants' personal moral beliefs. This may suggest that while curiosity drives the acquisition of moral information, it may not directly translate into belief change.

Taken together, these findings highlight the complexity that underlies information seeking in the moral domain. However, it remains unclear why people seek out more information about immoral agents than good ones. In Experiment 2, we replicate and extend the patterns evidenced here to explore possible motivations behind information gathering in the moral domain.

Experiment 2

Experiment 2 had two aims: (1) to replicate the findings from Experiment 1, and (2) to examine the motivations behind asymmetric moral information seeking. We considered two explanations for heightened negative moral information seeking. First, individuals may feel less compelled to seek additional information when the means justify the ends. If an action leads to both good and bad outcomes, but the means ultimately seem morally acceptable, there may be less unresolved ambiguity to explore. Second, reputational concerns may also influence moral information seeking. People may perceive greater social cost in labeling someone as morally bad than as morally good (or vice versa), prompting more thorough information gathering in morally negative contexts. To test these possibilities, we examined whether judgments of how justified an action seems, and participants' willingness to share their moral judgment with others, predict information seeking.

Overall, our predictions for Experiment 2 remained consistent with Experiment 1. We expected that higher self-reported curiosity and weaker prior beliefs would be associated with greater information seeking. For scenario justification, we hypothesized that perceiving an action as more justified would predict reduced information seeking. We also explored whether justification moderated the

relationship between personal moral beliefs and information seeking. Finally, we predicted that greater willingness to share one's moral judgment would be associated with increased information seeking.

Participants We recruited a total of 550 participants from Prolific in a gender balanced sample, based on the power analysis conducted in Experiment 1. We used the same inclusion criteria as Experiment 1. After removing those who did not pass our pre-registered exclusion criteria (i.e., failed attention and bot checks or sampled more than 3 standard deviations above the mean), we had a final sample of 544 participants (Mean age = 45.5, *SD* = 15.9, range: [18, 86]; Race: 66.18% White, 13.71% Black, 20.38% Other).

Primary Measures We used the same primary measures as in Experiment 1, including self-reported curiosity, moral belief updating, and meta-perceptions. All measures were assessed using the same scales as in Experiment 1.

Secondary Measures To better understand the impact of curiosity and moral valence on information seeking, we included two new measures. First, we assessed action justification—the extent to which participants believed the action described in each claim was justified—using a 100-point sliding scale ranging from "*not at all justified*" to "*completely justified*". Second, we measured participants' willingness to share their moral evaluation of the action with others, also on a 100-point scale ranging from "*not at all willing*" to "*extremely willing*". Both measures were collected prior to the information seeking task.

Finally, consistent with Experiment 1, we included a behavioral measure of curiosity in which participants could choose to view the average moral judgment rating for each claim.

Design & Procedure Experiment 2 used the same basic design as Experiment 1, with the added measures of action justification and willingness to share one's moral judgment. After completing three trials, participants again provided demographics, and were debriefed and compensated.

Results

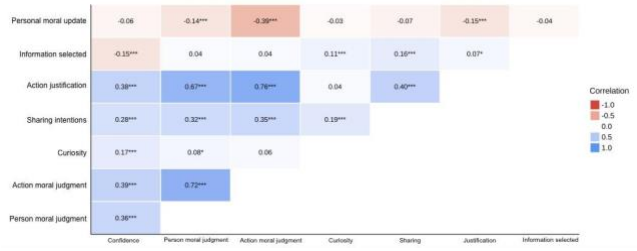

Figure 3: Correlations between measured variables.

Using the same mixed-effects models as in Experiment 1, we first tested whether self-reported curiosity predicted information seeking. As hypothesized, we replicated our previous finding: higher curiosity ratings were associated with greater information seeking, b = 0.22, SE = 0.05, z = 4.13, *95%CI* [0.12, 0.33], p < .001. Next, we examined whether confidence in one's moral judgments was associated with information seeking. As in Study 1, and in line with our predictions, we found that weaker confidence was associated with greater information seeking, b = -0.21, SE = 0.07, z = 3.19, *95%CI* [-0.34, -0.08], p = .001. Weaker confidence also predicted more updating of personal beliefs, b = -3.81, SE = 1.21, t = 3.16, *95%CI* [-6.18, -1.44], p = .002, and of meta-perceptions, b = -8.18, SE = 1.33, t = 6.13, *95%CI* [-10.79, -5.56], p < .001. That is, the less confident people were about their own judgments, the more negative their updating. We also again found that people sought more information about actions they saw as negative; however, this effect was marginal (b = -0.10, SE = 0.05, z = -1.92, p = 0.05) and should be interpreted with caution.

Critically, we next tested whether perceived action justification or willingness to share predicted information seeking. Contrary to our hypothesis, action justification did not directly predict information seeking, p = .17. However, action justification significantly moderated the relationship between personal moral ratings and information seeking, b = 0.57, SE = 0.16, z = 3.53, *95%CI* [0.25, 0.89], p < .001. Simple slopes analysis indicated that this moderation effect was driven by actions rated low in justification: for these actions, lower moral evaluations (i.e., more negative judgments) elicited greater information seeking than those seen as more morally good, b = -0.22, SE = 0.08, t = -2.68, p = 0.01 (see Figure 4). This pattern was not observed for actions rated as moderately or highly justified, ps > .34. Finally, we found no evidence that action justification moderated the relationship between self-reported curiosity and information seeking.

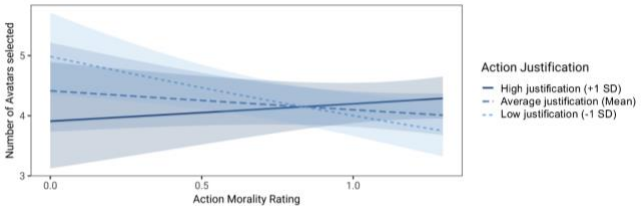

Figure 4: Negative moral judgments led to greater information seeking when action justification was low. Shading represents 95% Confidence Intervals.

We conducted parallel analyses to examine the relationship between willingness to share one's moral judgment and information seeking. As predicted, greater willingness to share was associated with increased information seeking, b = 0.26, SE = 0.09, z = 2.95, *95%CI* [0.09, 0.44], p = .003. However, unlike action justification, willingness to share did not moderate the relationship between personal moral beliefs and information seeking, p = .66.

We also included the behavioral curiosity measure at the end of the study, in which participants could choose to view average moral ratings for each claim. As expected, higher self-reported curiosity predicted increased viewing of the

average moral ratings, $b = 0.21$, $SE = 0.07$, $z = 2.93$, *95%CI* [0.07, 0.35], $p = .003$. In addition, we explored whether willingness to share one's moral judgments predicted viewing more average ratings. The results confirmed that greater willingness to share was associated with increased viewing of the average ratings, $b = 0.19$, $SE = 0.06$, $z = 2.96$, *95%CI* [0.07, 0.32], $p < .001$. Critically, however, this effect was moderated by personal moral judgments: the relationship between willingness to share and viewing the average ratings was stronger among participants who made more negative moral evaluations, $b = 0.19$, $SE = 0.06$, $z = 2.96$, *95%CI* [0.07, 0.32], $p < .001$. A simple slopes analysis revealed that this association was significant for those who judged the action as immoral ($b = .47$, $SE = .14$, $t = 3.30$, $p < .001$) or average ($b = .26$, $SE = .10$, $t = 2.50$, $p = .01$), but not for those who judged the action as moral ($p = .70$).

Study 2 Discussion

Replicating our findings from Experiment 1, we found that people who report higher moral curiosity were more likely to seek additional information. Additionally, weaker priors were associated with increased information seeking and greater belief updating, suggesting that people who were less confident were more open to revising their beliefs when presented with new information.

Importantly, while we do not find direct evidence that justification predicts greater information seeking overall, justification significantly moderated the relationship between moral judgments and information seeking behavior. Specifically, for actions low in justification, individuals sought more information when they initially judged the action to be immoral. This suggests that moral violations that lack clear justifications may trigger a greater drive to seek out additional information, possibly in an effort to confirm or reassess one's initial moral stance. In contrast, when actions were seen as more justified, information seeking did not vary systematically with moral judgments. This finding highlights the potential role of justification in shaping when and why individuals engage in moral information seeking, and points to the possibility that people may be particularly motivated to learn about morally ambiguous or seemingly unjustified behavior.

Another key finding is the role of willingness to share one's moral judgments in predicting information seeking behavior. As hypothesized, people who reported greater willingness to share their moral evaluations were also more likely to seek additional information. In addition, an exploratory analysis revealed that willingness to share was not only associated with greater information seeking, but also moderated engagement with our behavioral measure of curiosity: viewing average moral ratings. Participants who judged actions as more immoral (vs moral) were especially likely to examine average ratings and display greater willingness to share. This suggests that individuals may seek external validation when preparing to disclose moral judgments to others, especially when those judgments are negative.

General Discussion

How do curiosity-driven information seeking processes unfold in the moral domain? Across two experiments, we found that curiosity consistently predicted greater information seeking, regardless of whether the agent was perceived as morally good or bad. This finding supports the idea that curiosity, often piqued by salient information gaps (Dubey et al., 2022; Liquin & Lombrozo, 2020; Loewenstein, 1994), motivates individuals to resolve ambiguity (Wylie & Gantman, 2022), even in complex moral contexts.

However, our results also revealed important asymmetries. Specifically, participants sought less information about agents they judged to be morally good, suggesting that the drive to resolve moral ambiguity is stronger when evaluating potentially immoral actions. Although action justification did not directly predict information seeking, it did moderate the relationship between moral judgments and information seeking: when actions were perceived as unjustified, participants who judged them as immoral sought significantly more information. This asymmetry aligns with prior findings showing that negative moral judgments like blame tend to be more extreme and sensitive to context than positive judgments like praise (Guglielmo & Malle, 2019), and may demand greater cognitive elaboration. Similarly, our findings align with research showing that people are more cautious and slower to make negative moral judgments—especially when expecting contextual updates (Andrejević et al., 2022). Together, these findings suggest that moral evaluations involving perceived immorality and low justification evoke a heightened need for information. Additionally, the link between information seeking and willingness to share evaluations—especially negative ones—suggests that individuals may also seek information as a means of securing social validation before expressing potentially condemning views.

These findings have important real-world implications, particularly in domains in which moral information shapes reputation and decision-making. In contexts ranging from politics and business to everyday social interactions, individuals seek out moral information to evaluate others' trustworthiness and credibility—a process underscored by past work on the importance of moral reputation (Sperber & Baumard, 2012). The asymmetry we observed, where moral goodness reduces information seeking, suggests that once someone is perceived as good, people may feel less compelled to learn more about their character. Indeed, prior work has suggested that beliefs about moral goodness are more stable than moral badness (Siegel et al., 2018). Our findings suggest that the initial drive to seek out information may contribute to this stability—people are more inclined to seek out negative information to update their beliefs, even if it leads to more negative conclusions. As a result, it is possible that, for example, leaders with historically good reputations may receive less critical examination, while those with negative reputations may struggle to regain trust. These dynamics warrant further investigation in future research.

References

Andrejević, M., White, J. P., Feuerriegel, D., Laham, S., & Bode, S. (2022). Response time modelling reveals evidence for multiple, distinct sources of moral decision caution. *Cognition*, *223*, 105026.

Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., ... & Rahwan, I. (2018). The moral machine experiment. Nature, 563(7729), 59-64.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. Journal of Statistical Software, 67(1), 1–48.

Berlyne, D. E. (1966). Curiosity and exploration. Science (New York, N.Y.), 153(3731), 25–33.

Boyd, R., & Richerson, P. J. (1989). The evolution of indirect reciprocity. Social Networks, 11(3), 213-236.

Brambilla, M., Rusconi, P., Sacchi, S., & Cherubini, P. (2011). Looking for honesty: The primary role of morality (vs. sociability and competence) in information gathering. *European journal of social psychology*, *41*(2), 135-143.

Brambilla, M., Carraro, L., Castelli, L., & Sacchi, S. (2019). Changing impressions: Moral character dominates impression updating. Journal of Experimental Social Psychology, 82, 64-73.

Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. Psychological Science, 17(12), 1082-1089.

Dubey, R., Griffiths, T. L., & Lombrozo, T. (2022). If it's important, then I'm curious: Increasing perceived usefulness stimulates curiosity. Cognition, 226, 105193.

Everett, J. A., Pizarro, D. A., & Crockett, M. J. (2016). Inference of trustworthiness from intuitive moral judgments. Journal of Experimental Psychology: General, 145(6), 772.

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behavior Research Methods, 39(2), 175-191.

Gawronski, B., Armstrong, J., Conway, P., Friesdorf, R., & Hütter, M. (2017). Consequences, norms, and generalized inaction in moral dilemmas: The CNI model of moral decision-making. Journal of Personality and Social Psychology, 113(3), 343.

Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: exploring objectivism. Cognition, 106(3), 1339–1366.

Gottlieb, J., Oudeyer, P. Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. Trends in Cognitive Sciences, 17(11), 585–593.

Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. Neuron, 44(2), 389-400.

Guglielmo, S., & Malle, B. F. (2019). Asymmetric morality: Blame is more differentiated and more extreme than praise. PloS one, 14(3), e0213544.

Horn, S., Litovsky, Y., & Loewenstein, G. (2024). Using curiosity to counter health information avoidance. Social Science & Medicine, 340, 116383.

James, W. (1899/1983). Talks to Teachers on Psychology and to Students on Some of Life's Ideals (Vol. 12). Harvard University Press.

Jordan, J. J., Hoffman, M., Bloom, P., & Rand, D. G. (2016). Third-party punishment as a costly signal of trustworthiness. Nature, 530(7591), 473–476.

Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. Neuron, 88(3), 449-460.

Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. Nature Human Behaviour, 3(6), 587–595.

Krause, R. J., & Rucker, D. D. (2020). Can bad be good? The attraction of a darker self. Psychological Science, 31(5), 518-530.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: tests in linear mixed effects models. Journal of Statistical Software, 82(13).

LaMontagne, R. M. (2016). Ethical dilemmas in the workplace: HR managers' perceptions of behavioral influences with implications for building an ethical culture in organizations. International Journal of Knowledge, Culture & Change Management: Annual Review, 15.

Liquin, E. G., & Lombrozo, T. (2020). A functional approach to explanation-seeking curiosity. Cognitive Psychology, 119, Article 101276.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological Bulletin, 116(1), 75–98.

Marvin, C. B., & Shohamy, D. (2016). Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. Journal of Experimental Psychology: General, 145(3), 266.

Poli, F., O'Reilly, J. X., Mars, R. B., & Hunnius, S. (2024). Curiosity and the dynamics of optimal exploration. *Trends in Cognitive Sciences*, *28*(5), 441-453.

R Core Team (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.

Rand, D. G., & Nowak, M. A. (2013). Human cooperation. Trends in Cognitive Sciences, 17(8), 413-425.

Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. Personality and Social Psychology Review, 5(4), 296-320.

Siegel, J. Z., Mathys, C., Rutledge, R. B., & Crockett, M. J. (2018). Beliefs about bad people are volatile. Nature human behaviour, 2(10), 750-756.

Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. Nature Human Behaviour, 4(1), 14–19.

Skitka, L. J., Hanson, B. E., Morgan, G. S., & Wisneski, D. C. (2021). The psychology of moral conviction. Annual Review of Psychology, 72(1), 347-366.

Sperber, D., & Baumard, N. (2012). Moral reputation: An evolutionary and cognitive perspective. Mind & Language, 27(5), 495-518.

Sylwester, K., & Roberts, G. (2010). Cooperators benefit through reputation-based partner choice in economic games. Biology Letters, 6(5), 659-662.
Turiel, E. (2008). The development of morality. Child and adolescent development: An advanced course, 473-514.
Vaish, A., Grossmann, T., & Woodward, A. (2008). Not all emotions are created equal: The negativity bias in social-emotional development. Psychological Bulletin, 134(3), 383.
van Lieshout, L. L., de Lange, F. P., & Cools, R. (2019). Motives underlying human curiosity. Nature Human Behaviour, 3(6), 550-551.
van Lieshout, L. L., de Lange, F. P., & Cools, R. (2020). Why so curious? Quantifying mechanisms of information seeking. Current Opinion in Behavioral Sciences, 35, 112-117.
Wylie, J., & Gantman, A. (2022). What is moral ambiguity and when does it trigger curiosity?. In Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 44, No. 44).
Wylie, J., & Gantman, A. (2023). People are curious about immoral and morally ambiguous others. Scientific Reports, 13(1), 7355.
Zhan, Y., Xiao, X., Li, J., Liu, L., Chen, J., Fan, W., & Zhong, Y. (2018). Interpersonal relationship modulates the behavioral and neural responses during moral decision-making. Neuroscience Letters, 672, 15-21.