Exploration is Higher in Social Contexts at the Cost of Rewards

Rista C. Plate[1] (rplate@sas.upenn.edu), Huang Ham[1] (hamhuang@sas.upenn.edu), Adrianna C. Jenkins[1] (acjenk@sas.upenn.edu)

[1]University of Pennsylvania, 3720 Walnut St., Philadelphia, PA, United States

Abstract

In decision-making situations that arise repeatedly, there are tradeoffs between: (i) acquiring new information to facilitate future, related decisions (exploration) and (ii) using existing information to secure expected outcomes (exploitation). Exploration choices have been well characterized in nonsocial contexts, but choices to explore (or not) in social environments are less well understood. Social environments are of particular interest because a key factor that increases exploration in nonsocial contexts is environmental uncertainty, and the social world is appreciated to be highly uncertain. Here, participants searched for rewards in a series of grids that were either described as comprising real people distributing previously-earned points (social context) or as the result of a computer algorithm or natural phenomenon (nonsocial context). Participants explored more, and earned fewer rewards, in the social versus nonsocial context, suggesting that social uncertainty prompted exploration at the cost of task-relevant goals.

Keywords: exploration-exploitation; social context; uncertainty; directed exploration; random exploration

Introduction

Many decisions in human life rely on the consideration of multiple choices with uncertain outcomes. From small decisions that are made on a regular basis (e.g., what to order at a restaurant) to larger decisions made more rarely (e.g., whether or not to stay in one's current job), choices can be broadly characterized as either consistent with exploitation or exploration (see Melhorn et al., 2015 for a review). Exploitation involves leveraging information one already has to garner more certain, desired outcomes. For example, you might choose to invite an old friend for coffee (exploit a known social connection) because you are confident that you will have an enjoyable time. Exploration, on the other hand, involves trying an alternative about which you know less. Choosing to ask a new acquaintance to coffee (explore a new social connection) could result in a more or less enjoyable outcome, but either way it provides you with previously unknown information. Exploration decisions might be especially complex in social situations, like this one, because the factors affecting different outcomes are associated with substantial uncertainty: the internal states that drive other people's behavior are hidden from view, dynamic, and responsive to the behavior of others (Fiske, 1993). While the benefits and costs of exploration have been well studied in nonsocial domains, and uncertainty is known to be a key factor that drives exploration (Gershman, 2018; Gershman, 2019), relatively less is known about how the social landscape affects choices to explore. Here we test predictions about when, and with what consequences, people explore in social contexts. Specifically, we presented participants with a series of grids comprised of individual tiles and told them either that the tiles represented individual people who would give them some reward or that the rewards would be generated in a nonsocial fashion (by a computer in Experiment 1; by a natural phenomenon in Experiment 2). We investigated to what extent this simple framing of the context as social or nonsocial affected exploration behavior and reward receipt within the grids.

Given its relevance to decision-making, both human and nonhuman animal behavior (e.g., foraging), and the cognitive processes that underlie these activities (e.g., memory search; Hills et a., 2015; Todd & Hills, 2020), there has been substantial interest in the factors that influence when, why, and for how long individuals engage in behavioral exploration of nonsocial contexts (see Cohen et al., 2007; Melhorn et al., 2015 for reviews). Research suggests that preferences for exploration are not fixed (Melhorn et al., 2015) but instead reflect the interaction of multiple contextual considerations (Schulz et al., 2018a & b), including the age of the learner (Plate et al., 2019), the time one has to potentially acquire information (Wilson et al., 2014), and an agent's goals (Melhorn et al., 2015). Indeed, uncertainty is one especially important factor known to increase exploration (Speenkenbrink & Konstantinidis, 2015). Exploration of uncertain environments can satisfy curiosity (Kidd & Hayden, 2015; Liquin et al., 2020), reduce boredom (Geana et al., 2016), and support the pursuit of knowledge to inform flexible changes in behavior (Sharot & Sunstein, 2020).

In the social domain, there has been progress in understanding exploration in collaborative and competitive group environments (e.g., Goldstone et al., 2005; Hills et al., 2015). For example, individuals can benefit from social learning to reduce the need to rely on exploration in group

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

contexts (Toyokawa et al., 2014). At the same time, there can be a tendency to over-rely on social information under uncertain learning conditions (Plate et al., 2021; Toyokawa et al., 2017). Learners must maintain sensitivity to structural features of the environment (including reward structure and predictability), thereby balancing—and flexibly updating—the use of social information over time (Wu et al., 2021). There is some evidence that learners can achieve this balancing act, for example, adjusting behavior within a brief experimental session while engaging in both individual exploration and social learning (Krafft et al., 2015). However, much remains unknown about exploration-related choices in social contexts, which are particularly interesting given that individuals may weight uncertainty differently when the uncertainty is social in nature (Blount, 1995; Li et al., 2018; Rilling et al., 2008).

Relatively high unpredictability and ambiguity in the social world has been well documented (e.g., Jenkins & Mitchell, 2010; Feldmann-Hall & Shenhav, 2019; Hertwig & Herzog, 2009). Individuals may need to turn to exploration to resolve this uncertainty in social contexts. In other words, they may need to interact with others to gather information that they could then employ to make inferences in the future. Decisions involving other people can involve heightened uncertainty because not only is the outcome itself uncertain (e.g., whether having coffee with a new acquaintance would be enjoyable), but the social factors influencing that outcome are also uncertain (e.g., whether the acquaintance will be in a pleasant mood when they arrive or whether they are interested in becoming friends). This additional level of uncertainty may trigger exploration to a greater extent than would be expected in nonsocial contexts, in which the primary source of uncertainty is in the outcome itself.

In two experiments, we asked how uncertainty in social (versus nonsocial) contexts influences exploration. Participants searched for rewards across tiles in a series of grids in which rewards were either supposedly generated by other people (social context) or by the physical environment (nonsocial context). To assess the consequences of different search approaches, we measured the rewarding outcomes participants received during their search. Across all experiments, we manipulated the environmental structure (defined as the degree to which rewards cluster together in the search space) to assess whether comparisons between the social and nonsocial contexts mirror comparisons between environments that are higher or lower in uncertainty.

Method – Experiment 1

In Experiment 1, we asked to what extent patterns of exploration differ when the search context is social versus nonsocial, holding constant the underlying reward structure across contexts. We compared the degree of exploration (i.e., how much participants explored choices with unknown rewards versus opted for choices with known rewards) in social and nonsocial contexts. To characterize the type of exploration participants adopted, we used a modeling approach to compare the extent to which participants engaged in *directed exploration* (i.e., selecting options that will reduce the overall uncertainty of the search space), and the extent to which participants engaged in *random exploration* (i.e., selecting unknown options, but not specifically targeting options that will reduce the overall uncertainty) in social and nonsocial contexts. To the extent that participants associate the social context with higher uncertainty than the nonsocial one, they should demonstrate more exploration (specifically, directed exploration, which would indicate targeted uncertainty reduction; Wilson et al., 2014) in the social context.

Participants

Participants were 142 participants (51 female, 87 male, 4 self-described or did not provide gender information; 14 participants were 18-25-years-old, 76 participants were 26-35-years-old, 36 participants were 36-50-years-old, and 16 participants were older than 50-years-old). Thirteen additional participants were excluded for earning a bonus of less than \$0.50 (indicating low engagement in the task) or bot-like responses (i.e., text appeared to be sourced from website content or was identical for multiple participants). Participants received \$0.50 for their participation in the ten-minute task. We restricted participation to MTurk workers with HIT acceptance rates >97% who were located in the United States.

Design & Procedure

The experimental task was adapted from Experiment 2 of Wu and colleagues (2018). While there are many experimental paradigms that set up a tension between exploration and exploitation, we had three additional goals that influenced our task selection, namely to: (1) include a large landscape for possible exploration in order to reflect the scope and variety of options that characterize social contexts; (2) use the same task structure across social and nonsocial contexts; and (3) be well enough established in the literature to support an extension to the social domain, particularly with regard to model based analysis. Wu and colleagues' (2018) "grid task" satisfied these conditions. Participants were randomly assigned to one of four between-subjects conditions in a 2 (context: social, nonsocial) X 2 (environment: rough, smooth) design (N rough, social = 36; N smooth, social = 39; N rough, nonsocial = 29; N smooth, nonsocial = 38). Participants' task was to search for points in a series of eight grids, presented sequentially, each containing 121 tiles. Participants in the social context were told that each tile on the grid represented an MTurk worker who had previously played the game and was able to allocate a proportion of points that they earned on each of their own clicks to someone else, i.e., the current participant. Participants in the nonsocial context were told that the point value for each tile in the grid was determined via a computer algorithm. Grids in the rough condition were sampled from a Gaussian process prior with radial basis having $\lambda = 1$ and grids in the smooth condition were sampled from a Gaussian process prior with $\lambda = 2$, where λ is a length-scale parameter that indicates how quickly the correlation between rewards in the grid decreases across the spatial layout of the grid. Therefore, the social manipulation provided participants explicit information about the task context whereas the uncertainty information had to be gleaned over time in an exclusive "bottom-up" fashion.

We additionally varied the time horizon for exploration across the grids in a within-subjects manner because horizon has been shown to impact the extent of exploration (Wilson et al., 2014), but this manipulation was not central to our hypotheses. Specifically, for four of the grids, participants had 20 clicks ("short search horizon") to search for points before proceeding to the next grid; for the other four grids, participants had 40 clicks ("long search horizon"). The order of the short and long search horizons alternated during the task, and we counterbalanced which search horizon was assigned to the first grid. Participants were instructed to find as many points as possible and told that the magnitude of their bonus (up to $1.50) was dependent on how many points they found. The task was self-paced.

Data Analysis

First, to assess exploration, we regressed the distance (using Manhattan distance) from the previously selected tile on environment (smooth = -.5, rough = .5) and context (nonsocial = -.5, social = .5) using a linear mixed effects model with random intercepts for participant, horizon, and specific search environment (i.e., the specific spatial distribution of underlying rewards). To further disentangle the factors contributing to how people explore differently in social vs non-social contexts, we fit to our data a computational model that decomposes the search behavior into 3 components: generalization, directed exploration, and random exploration (Schulz et al., 2019, Wu et al., 2018).

Generalization aims to capture the mechanism through which people generalize from the rewards of the observed tiles to all tiles. It is formulated as a Gaussian-process regression (Quinonero-Candela et al., 2007) with the radial-basis function as kernel: $k(\boldsymbol{x},\boldsymbol{x}') = \exp\left(-\frac{\|x-x'\|^2}{\lambda}\right)$. The kernel has a *length-scale* parameter λ governing how fast the reward correlation decays as the distance between 2 tiles (x and x') increases, which is a free parameter we fit to the data to capture participants' degree of generalization.

Directed exploration aims to assign a subjective value to each tile that will guide which tile to choose next. From generalization, we can extract 2 pieces of information: the expected value and standard deviation of reward in each tile. The subjective value of a tile is obtained by combining them using upper-confidence-bound (UCB) sampling (Srinivas et al., 2009): $UCB(\boldsymbol{x}) = \mu(\boldsymbol{x}) + \beta\sigma(\boldsymbol{x})$, where a free parameter β controls how much the standard deviation $\sigma(\boldsymbol{x})$ contributes to the subjective value beyond the expected value $\mu(\boldsymbol{x})$. Thus β encodes how much the exploration tendency is directed by the degree of reward uncertainty.

Random exploration converts the values of each tile into a probability distribution over all tiles from which the next choice will be sampled (i.e., a behavioral policy). It is achieved by using the *softmax* function: $p(\boldsymbol{x}) = \frac{\exp\left(UCB(\boldsymbol{x})/\tau\right)}{\sum_{j=1}^{N}\exp\left(UCB(\boldsymbol{x}_j)/\tau\right)}$ with a temperature parameter τ capturing how much exploration happens simply due to behavioral stochasticity. Higher value of τ translates into less value-guided behavior and choosing tiles more randomly.

Each of the parameters described above can range from 0 to infinity theoretically; following prior work, we fit them using the bound $[\exp(-5), \exp(5)]$. We adopted leave-one-out cross validation as the fitting procedure. For each of 8 rounds assigned to a participant, we held out that round and applied Maximum Likelihood Estimation to fit the model to the remaining 3 rounds of the same horizon length. We then averaged estimates from all 8 rounds to obtain the fitted parameter value for that participant. To compare the fitted parameter values between experimental conditions, we chose to use Mann-Whitney U tests, which are robust to extreme values of the estimated parameters.

Finally, we examined the rewards that participants found during the task in two ways. First, we examined whether context and environment influenced participants' overall rewards by regressing the average reward received on environment (smooth = -.5, rough = .5) and context (nonsocial = -.5, social = .5) using a linear mixed effects model with random intercepts for participant, horizon, and specific search environment. We ran a mediation analysis using the mediation package (Tingley et al., 2014) to evaluate whether exploration patterns mediated the relationship between context and reward receipt. Second, we ran the linear mixed effects model with the maximum (rather than average) reward found to test whether there were any differences participants' ability to find the highest reward presented in each of the grids.

Results – Experiment 1

Participants Explore More in Social Contexts

The first question of interest was how the social (vs. nonsocial) context affected exploration behavior. There was a main effect of context on exploration (b = 0.40, $X^2(1)$ = 4.76, p = .029, 95% CI = [0.04, 0.76]; Figure 1), such that participants explored more in the social context (when told that the tiles were comprised of individuals sharing a proportion of their previously-earned rewards) than in the nonsocial context (when told the rewards associated with the tiles were generated by a computer). This pattern is consistent with the idea that uncertainty is heightened in social contexts, and participants turned to behavioral exploration to reduce it.

In line with previous research (Speenkenbrink & Konstantinidis, 2015), there was also a main effect of environment on exploration (b = 0.44, $X^2(1)$ = 5.72, p = .017, 95% CI = [0.08, 0.80]), such that participants explored more in rough than smooth environments across contexts. The context X environment interaction was not significant (b = -

0.18, $X^2(1)$ =0.25, p = .617, 95% CI = [-0.90, 0.54]). These findings also held both for short and long horizon grids, indicating that participants engage in more exploration in the social (vs. non-social) context regardless of the actual level of uncertainty in the environment and even when a short time window for obtaining rewards limits exploration's instrumental benefits.

Model-based analyses made it possible to disentangle two possible sources of exploratory behavior: exploration that reduces uncertainty (directed exploration: β) and choosing tiles at random (random exploration: τ). We found that only the directed exploration parameter β was higher in the social context than the non-social context (U = 3073, p = 0.022; Figure 2), showing that elevated exploration in the social context was driven by motives to reduce uncertainty. The random exploration parameter τ and generalization parameter λ did not differ between contexts (U = 2633, p = 0.624; U = 2566, p = 0.829), meaning that participants did not explore tiles more randomly or infer stronger reward correlation between tiles in one context compared to the other. Together, these results are consistent with the idea that elevated exploration of social contexts occurs in the service of obtaining information that could plausibly be leveraged over longer timescales in the future to obtain desired outcomes.

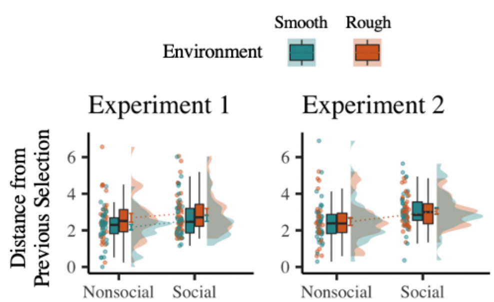

Figure 1: Participants showed more exploration in the social condition in both Experiments 1 and 2.

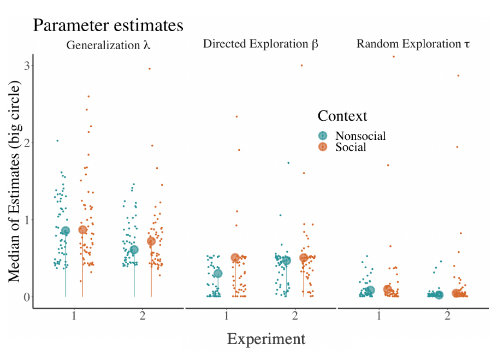

Figure 2: Model parameter estimates. We first averaged the cross-validated parameter estimates for each participant, and then plotted the median across participants by context.

Participants Earn Fewer Rewards in Social Contexts

The second question of interest was how the social (vs. nonsocial) context affected participants' reward earnings. Analysis of the average rewards participants garnered during the task (i.e., across all grids) revealed a main effect of social context (b = -3.15, $X^2(1)$ = 3.81, p = .051, 95% CI = [-6.31, -0.01]), with participants earning lower rewards in the social context. Participants also earned lower rewards in rough (vs. smooth) environments (b = -8.27, $X^2(1)$ = 26.26, p < .001, 95% CI = [-11.43, -5.11]). (The interaction was not significant, b = 2.09, $X^2(1)$ = 0.42, p = .519, 95% CI = [-4.24, 8.42].)

Moreover, the relationship between context and reward receipt was mediated by exploration (b = -0.57, p = 0.017, 95% CI = [-1.05, -0.09]). This provides evidence that the social manipulation acted on patterns of exploration, which in turn influenced reward receipt. In other words, pursuing additional information about the social context to reduce uncertainty was done at the detriment of reward receipt, which is notable considering that participants in this task could earn a monetary bonus based on the rewards they found.

Participants' lower reward earnings in the social context could derive from at least two possible sources: a lower success rate in finding as highly-rewarding tiles in the social context (compared to the non-social context) or a greater tendency to continue exploring even after finding highly-rewarding tiles in the social context (compared to the non-social context). To investigate these possibilities, we examined the magnitude of the highest reward found throughout the task. This analysis revealed no effect of context on the highest reward found by participants (b = 0.07, $X^2(1)$ = 0.007, p = .933, 95% CI = [-1.69, 1.84]). We observed only an effect of environment, such that participants found higher rewards in the smooth, versus rough, environments (b = -2.58, $X^2(1)$ = 8.24, p = .004, 95% CI = [-4.34, -0.82]), which is expected given that the distribution of rewards was less predictable in the rough environments. The interaction between context and environmental uncertainty was not significant (b = -1.46, $X^2(1)$ = 0.66, p = .418, 95% CI = [-4.98, 2.07]). Together, these results show that participants in both contexts had the opportunity to exploit high-reward tiles, but those in the social context were less likely to do so.

Discussion – Experiment 1

In Experiment 1, participants explored more when they thought that the points associated with the grid tiles were determined by human individuals than when they thought the points were determined by a computer algorithm. Additionally, the social context specifically increased directed (as opposed to random) exploration, which targets uncertainty reduction.

One principal limitation in Experiment 1 is the difference in instructions for the social and nonsocial contexts, which leaves open the possibility that the observed difference in exploration between contexts could have arisen due to factors other than the social versus nonsocial context per se. For example, participants might have expected that the points generated from the computer were random but that the points coming from other people were associated with more systematic patterns.

In particular, although we selected the Wu et al. (2018) task for many desirable features, including the vastness of the search space and its established place in the literature, this foraging-style task may have introduced a confound with condition in participants' expectations about the reward distributions. If spatial correlations are less expected in the social context, this could in part explain differences between the conditions. Such a difference is plausible; participants may not expect that individuals associated with adjacent tiles would issue similar rewards, instead inferring that individual motivations could influence the social generation of rewards (Wilke et al., 2015). We address this issue in Experiment 2 by providing an explanation in both conditions as to why there may be spatial relationships between the tiles.

Another limitation is that the experiences of participants completing the task itself might have differed across contexts. For example, participants could have encoded observed rewards differently across conditions. To assess the replicability of the observed patterns of elevated exploration in social (vs. non-social) contexts in a new sample while addressing these considerations, we adapted our social and nonsocial conditions to be even more comparable in structure in Experiment 2.

Method – Experiment 2

The method of Experiment 2 was identical to that of Experiment 1, with the following exceptions. First, we changed the descriptions of the social and nonsocial contexts to make them more parallel. Participants in the social context read that each tile represented an MTurk worker who shared a portion of their points from when the MTurk worker played the game in 2020. Unlike Study 1, participants also read that the tiles were arranged based on the MTurk worker's geographic location. Our intention here was to provide plausibility for the spatial correlations between tiles. Further, in the nonsocial context, participants read that each tile represented a plot of land and that the points corresponded to the crop yield from that plot of land in 2020. As in the social context, participants read that the tiles were arranged based on the plot of land's geographic location.

Second, we wanted to understand whether participants had different expectations and/or experiences of the possible rewards in the social and nonsocial contexts. At the end of the task, we asked participants to report the highest and the lowest reward that they thought was present overall (i.e., across all of the grids they saw). Finally, we revised the bonus structure such that that top 50 point-earning participants would earn a bonus of $0.50.

Participants were 138 adults (N rough, social = 37; N smooth, social = 35; N rough, nonsocial = 28; N smooth, nonsocial = 38) recruited via Amazon's Mechanical Turk (61 female, 76 male, 1 participant did not report gender, 9 participants were 18-25-years-old, 54 participants were 26-35-years-old, 56 participants were 36-50-years-old, and 19 participants were older than 50-years-old). 68 additional participants were excluded for earning a bonus of less than $0.50 (indicating low engagement in the task), estimating the highest reward as < 50 or lowest reward as > 50 (which would be highly inconsistent with the grid displays, where the lowest rewards observed by participants ranged from 3 to 17 and the highest ranged from 69 to 87) or bot-like responses.

Results – Experiment 2

Participants Explore More in Social Contexts

In Experiment 2, we replicated the main effect of the social context (b = 0.57, $X^2(1)$ = 9.54, p = .002, 95% CI = [0.21, 0.93]; Figure 1), with participants exploring more when told that the tiles were comprised of other people versus representing plots of land. The effect of environment was not significant (b = -0.0009, $X^2(1)$ = 0.00, p = .996, 95% CI = [-0.36, 0.36], nor was the interaction, $X^2(1)$ = 0.05, p = .823). It is possible that telling participants that the tiles were arranged geographically made them less sensitive to the observed spatial correlation of the grids. These results were consistent across both the short and long time horizons.

Replicating the modeling results of Experiment 1, the directed exploration parameter β was higher in the social context than nonsocial context (U = 2875, p = 0.034; Figure 2), suggesting participants valued the reduction of uncertainty more if told that rewards were generated from other people than contingent on crop yield. Unlike Experiment 1, there was also a difference in the random exploration τ parameter (U = 3030, p = 0.01; generalization was not significant, U = 2516, p = 0.55).

Participants Earn Fewer Rewards in Social Contexts

As in Experiment 1, there was a main effect of social context (b = -5.18, $X^2(1)$ = 13.24, p < .001, 95% CI = [-7.96, -2.39]) on average reward, with participants earning lower rewards under the social context. Participants also earned lower rewards in the rough environments (b = -7.71, $X^2(1)$ = 29.37, p < .001, 95% CI = [-10.50, -4.92]). The interaction was not significant, $X^2(1)$ = 0.200, p = .655. Moreover, exploration again mediated the relationship between context and reward receipt (b = -0.98, p = .001, 95% CI = [-1.59, -0.36]). Finally, we also replicated the finding from Experiment 1 that environment, but not social context, affected the magnitude of highest reward found. Participants found higher rewards in the smooth environments (b = -3.55, $X^2(1)$ = 25.68, p < .001, 95% CI = [-4.92, -2.18]), but neither the effect of context ($X^2(1)$ = 1.05, p = .306) nor the interaction ($X^2(1)$ = 0.04, p = .839) was significant.

Participant Estimates of Reward do not Differ in Social Contexts

In Experiment 2, we included additional questions at the end of the experiment to test whether participants had different experiences of the rewards between contexts that could explain differences in exploration behavior. Participant estimates of the highest rewards available were related to the highest rewards that they, themselves, earned during the task (r = 0.37, $t(136)$ = 4.63, p < .001, 95% CI [0.22, 0.50]), indicating that participants were accurate at tracking rewards, though notably, the correlation is small to moderate perhaps indicating some additional factors influencing reward estimation. In taking the difference score between the actual and estimated rewards, there were no differences based on environment ($t(134)$ = 0.43, p = .670), social context ($t(134)$ = -1.42, p = .159), or their interaction ($t(134)$ = -0.55, p = .586). This evidence suggests that explicit differences in expectations and/or experiences of reward distributions cannot account for differences in participant behavior across contexts.

General Discussion

The aim of the current investigation was to further scientific understanding of exploration in social contexts. In two experiments, we showed that participants demonstrated more exploration in a social versus nonsocial context. Enhanced exploration is consistent with the idea that social contexts present additional uncertainty that learners attempt to resolve through search and sampling of options. Our model-free results were further backed up by results from a computational model that defines exploration with mathematical precision as a form of information seeking and encapsulates it into one parameter β. Comparing the estimated β parameter yields a significant difference suggesting that participants enhanced exploratory behavior in the social context is in the service of reducing environmental uncertainty. In addition to differences in directed exploration, we did find differences in random exploration between contexts in Experiment 2. While individuals may pursue both types of exploration (e.g., Gershman, 2018; Wilson et al., 2014), future research is needed to further understand when the use of directed and random exploration diverge in social contexts. These results highlight the complex interactions between features of the environment and call for additional research on exploration tradeoffs in social contexts.

Limitations and Future Directions

Given that this is one of the first investigations of exploration in social contexts outside of collaborative/competitive group environments, there are limitations on our conclusions that guide directions for future research. First, in order to best equate social and nonsocial contexts, the social scenarios presented in these experiments were highly pared down and therefore limited in ecological validity. Real-world social contexts undoubtedly provide additional cues that could influence when, why, and for how long individuals explore. For example, if participants were soliciting donations from a group of people, search may be influenced by factors including their relationship to individual social agents, whether social agents present cues that promote approach (e.g., a pleasant facial expression), and norms about whether it would be permissible to "exploit" an individual agent by asking for donations across multiple occasions. Relatedly, it may have been less plausible for participants to expect spatial correlations in the social context (particularly in Experiment 1 when there were no instructions to introduce this idea). To fully map the landscape of exploration in social environments, future research should leverage the many diverse tasks that tap into exploratory behavior and build a knowledge base that parallels, and is integrated with, the vast research on exploration in nonsocial and collaborative contexts. In doing so, future research could also use reward structures that are in fact generated by social agents (e.g., Wilke et al., 2015) and therefore take into account the natural patterns and variations of reward generation in social contexts as well as further measure participant expectations about rewards generated by social agents as compared to those generated by nonsocial means.

Additionally, while the current experiments examined search for monetary rewards from social agents, a related but distinct set of questions concerns how people explore social landscapes for rewards that are themselves social (e.g., relationship value, increased access to resources, emotional rewards) (Cords & Aureli, 2000; de Waal, 1997; Kummer, 1978; Wittig et al., 2008) and navigate potential social costs (e.g., risk of interpersonal aggression, energy expenditure) (Mitani & Amsler, 2003; de Waal & Davis, 2003). For example, research suggests that individuals would rather lose monetary reward because of chance than because of another person (Blount, 1995; Bohnet & Zeckhauser, 2004) and that there may be social-specific risk aversion (Haux et al., 2021). Future research should consider other factors that influence exploration in social contexts beyond uncertainty.

Conclusion

We provide evidence to characterize exploration as a means to uncertainty reduction in social contexts. Specifically, participants demonstrated more exploration of social as compared to nonsocial contexts, in line with patterns of behavior expected in high uncertainty environments. Moreover, we find that the increased exploration in social relative to nonsocial contexts came at the cost of reward receipt. Together, these experiments provide evidence consistent with the view that social context uniquely influences exploration as social beings attempt to resolve uncertainty.

References

Blount, S. (1995). When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences. *Organizational Behavior and Human Decision Processes*, *63*(2), 131–144. https://doi.org/10.1006/obhd.1995.1068

Bohnet, I., & Zeckhauser, R. (2004). Trust, risk and betrayal. *Journal of Economic Behavior & Organization*, *55*(4), 467–484. https://doi.org/10.1016/j.jebo.2003.11.004

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 933–942. https://doi.org/10.1098/rstb.2007.2098

Cords, M., & Aureli, F. (2000). Reconciliation and relationship qualities. In F. Aureli, F.B.M. de Waal (Eds.), Natural Conflict Resolution, University of California Press, Berkeley (2000), pp. 177-198

de Waal, F. B. (1997). The chimpanzee's service economy: food for grooming. *Evolution and Human Behavior*, *18*(6), 375-386. https://doi.org/10.1016/S1090-5138(97)00085-8

de Waal, F. B. M., & Davis, J. M. (2003). Capuchin cognitive ecology: Cooperation based on projected returns. *Neuropsychologia*, *41*(2), 221–228. https://doi.org/10.1016/S0028-3932(02)00152-5

FeldmanHall, O., & Shenhav, A. (2019). Resolving uncertainty in a social world. *Nature Human Behaviour*, *3*(5), 426–435. https://doi.org/10.1038/s41562-019-0590-x

Fiske, S. T. (1993). Social cognition and social perception. *Annual review of psychology*, *44*(1), 155-194.

Geana, A., Wilson, R., Daw, N. D., & Cohen, J. (2016). Boredom, Information-Seeking and Exploration. In *CogSci*.

Gershman S. J. (2018). Deconstructing the human algorithms for exploration. Cognition, 173, 34–42. https://doi.org/10.1016/j.cognition.2017.12.014

Gershman S. J. (2019). Uncertainty and Exploration. Decision (Washington, D.C.), 6(3), 277–286. https://doi.org/10.1037/dec0000101

Goldstone, R. L., Ashpole, B. C., & Roberts, M. E. (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review*, *12*(1), 81–87. https://doi.org/10.3758/BF03196350

Haux, L. M., Engelmann, J. M., Herrmann, E., & Hertwig, R. (2021). How chimpanzees decide in the face of social and nonsocial uncertainty. *Animal Behaviour*, *173*, 177–189. https://doi.org/10.1016/j.anbehav.2021.01.015

Hertwig, R., & Herzog, S. M. (2009). Fast and Frugal Heuristics: Tools of Social Rationality. *Social Cognition*, *27*(5),661–698. https://doi.org/10.1521/soco.2009.27.5.661

Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, *19*(1), 46–54. https://doi.org/10.1016/j.tics.2014.10.004

Jenkins, A. C., & Mitchell, J. P. (2010). Mentalizing under uncertainty: dissociated neural responses to ambiguous and unambiguous mental state inferences. *Cerebral Cortex*, *20*(2), 404-410. https://doi.org/10.1093/cercor/bhp109

Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron*, *88*(3), 449–460. https://doi.org/10.1016/j.neuron.2015.09.010

Krafft, P. M., Hawkins, R. X., Pentland, A., Goodman, N. D., & Tenenbaum, J. B. (2015). Emergent Collective Sensing in Human Groups. In CogSci.

Kummer, H. (1978). On the value of social relationships to nonhuman primates: A heuristic scheme. *Social Science Information*, *17*(4–5), 687–705. https://doi.org/10.1177/053901847801700418

Li, C., Turmunkh, U., & Wakker, P. P. (2019). Trust as a decision under ambiguity. *Experimental Economics*, *22*(1), 51-75. https://doi.org/10.1007/s10683-018-9582-3

Liquin, E., Callaway, F., & Lombrozo, T. (2020). Quantifying Curiosity: A Formal Approach to Dissociating Causes of Curiosity. In *CogSci*.

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191–215. https://doi.org/10.1037/dec0000033

Mitani, J., & Amsler, S. (2003). Social and spatial aspects of male subgrouping in a community of wild chimpanzees. *Behaviour*, *140*(7), 869–884. https://doi.org/10.1163/156853903770238355

Plate, R. C., Fulvio, J. M., Shutts, K., Green, C. S., & Pollak, S. D. (2018). Probability Learning: Changes in Behavior Across Time and Development. *Child Development*, *89*(1), 205–218. https://doi.org/10.1111/cdev.12718

Plate, R. C., Shutts, K., Cochrane, A., Green, C. S., & Pollak, S. D. (2021). Testimony bias lingers across development under uncertainty. Developmental Psychology, 57(12), 2150. https://doi.org/10.1037/dev0001253

Quinonero-Candela, J., Rasmussen, C. E., & Williams, C. K. (2007). Approximation methods for Gaussian process regression. In *Large-scale kernel machines* (pp. 203-223). MIT Press.

Rilling, J. K., King-Casas, B., & Sanfey, A. G. (2008). The neurobiology of social decision-making. *Current opinion in neurobiology*, *18*(2), 159-165. https://doi.org/10.1016/j.conb.2008.06.003

Schulz, E., Wu, C. M., Huys, Q. J., Krause, A., & Speekenbrink, M. (2018). Generalization and search in risky environments. Cognitive science, 42(8), 2592-2620.

Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2018). Putting bandits into context: How function learning supports decision making. *Journal of experimental psychology: learning, memory, and cognition*, *44*(6), 927. https://doi.org/10.1101/081091

Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological science*, *30*(11), 1561-1572. https://doi.org/10.1177/0956797619863663

Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. *Nature Human Behaviour*, *4*(1), 14–19. https://doi.org/10.1038/s41562-019-0793-1

Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and Exploration in a Restless Bandit Problem. *Topics in Cognitive Science*, *7*(2), 351–367. https://doi.org/10.1111/tops.12145

Srinivas, N., Krause, A., Kakade, S. M., & Seeger, M. (2009). Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*.

Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2014). Mediation: R package for causal mediation analysis. *Journal of Statistical Software, 59*(5), 1-38. URL: http://www.jstatsoft.org/v59/i05/.

Todd, P. M., & Hills, T. T. (2020). Foraging in mind. *Current Directions in Psychological Science, 29*(3), 309-315. https://doi.org/10.1177/0963721420915861

Toyokawa, W., Saito, Y., & Kameda, T. (2017). Individual differences in learning behaviours in humans: Asocial exploration tendency does not predict reliance on social learning. Evolution and Human Behavior, 38(3), 325-333. https://doi.org/10.1016/j.evolhumbehav.2016.11.001

Toyokawa, W., Kim, H. R., & Kameda, T. (2014). Human collective intelligence under dual exploration-exploitation dilemmas. PloS one, 9(4), e95789. ttps://doi.org/10.1371/journal.pone.0095789

Wilke, A., Minich, S., Panis, M., Langen, T. A., Skufca, J. D., & Todd, P. M. (2015). A game of hide and seek: Expectations of clumpy resources influence hiding and searching patterns. PloS one, 10(7), e0130976. https://doi.org/10.1371/journal.pone.0130976

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans Use Directed and Random Exploration to Solve the Explore–Exploit Dilemma. *Journal of Experimental Psychology. General*, *143*(6), 2074–2081. https://doi.org/10.1037/a0038199

Wittig, R. M., Crockford, C., Lehmann, J., Whitten, P. L., Seyfarth, R. M., & Cheney, D. L. (2008). Focused grooming networks and stress alleviation in wild female baboons. *Hormones and Behavior*, *54*(1), 170–177. https://doi.org/10.1016/j.yhbeh.2008.02.009

Wu, C. M., Ho, M. K., Kahl, B., Leuker, C., Meder, B., & Kurvers, R. (2021). Specialization and selective social attention establishes the balance between individual and social learning. In Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 43, No. 43).

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, *2*(12), 915–924. https://doi.org/10.1038/s41562-018-0467-4