From seeing to seeking: belief-based exploration in gamified environments

Thirza Dado* thirza.dado@donders.ru.nl

Floris de Lange* floris.delange@donders.ru.nl

Eva Berlot* eva.berlot@donders.ru.nl

*Donders Centre for Cognitive Neuroimaging, Kapittelweg 29, 6525 EN Nijmegen, The Netherlands

Abstract

Contemporary neuroscientific theories postulate that experience is shaped not only by the world as it is, but also by our *beliefs* about it. When beliefs conflict with incoming sensory input, this mismatch should ideally trigger belief updating. However, research on confirmation bias shows that we often discount information that contradicts our prior beliefs. Here, we study whether information sampling after belief formation is biased toward confirmation. To this end, we designed a gamified experiment in which participants decide which of two planets contains more (or fewer) stones. After an initial impression that sets their prior belief, participants continue exploring — first by re-viewing the same visual input, then by actively seeking new evidence. This design allows us to examine how initial beliefs, along with their associated confidence, shape subsequent information gathering across different modes of interaction. This work investigates how we form, update and act on our beliefs, bridging perceptual and action-based exploration.

Keywords: belief formation; exploration; gamified environment

Introduction

The brain is often described as a prediction machine, constantly generating expectations about the world and comparing them to incoming sensory input (Rao & Ballard, 1999; Friston, 2005; Clark, 2013). To generate useful predictions, it constructs internal *beliefs* about the world, which help constrain the space of possible outcomes and guide our responses. Ideally, these beliefs should remain calibrated to the environment, adjusting them as new evidence emerges. But in practice, humans operate under conditions of *bounded rationality*, relying on shortcuts or settling for "good enough" solutions (satisficing) (Pirolli & Card, 1999). As a result, belief updating may become biased: new evidence may be selectively integrated or downweighted after a decision (Bronfman et al., 2015; Talluri, Urai, Tsetsos, Usher, & Donner, 2018). Such post-decision biases may contribute to confirmation bias, often observed in human reasoning and behavior (Lord, Ross, & Lepper, 1979), where the very mechanism intended to update our beliefs instead reinforces them, limiting our ability to adapt. These biases strengthen with higher confidence in initial beliefs (Rollwage et al., 2020).

Recent work shows that humans tend to favor belief-consistent information during perceptual decisions (Kaanders, Sepulveda, Folke, Ortoleva, & De Martino, 2022). In a dot comparison task, participants resampled the patch they initially chose — rather than both equally (the optimal strategy).

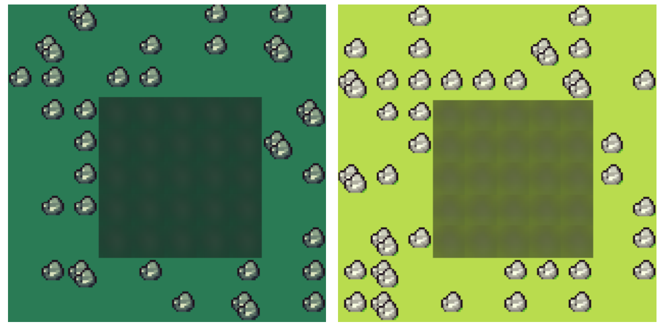

(A) Passive viewing of two planets.

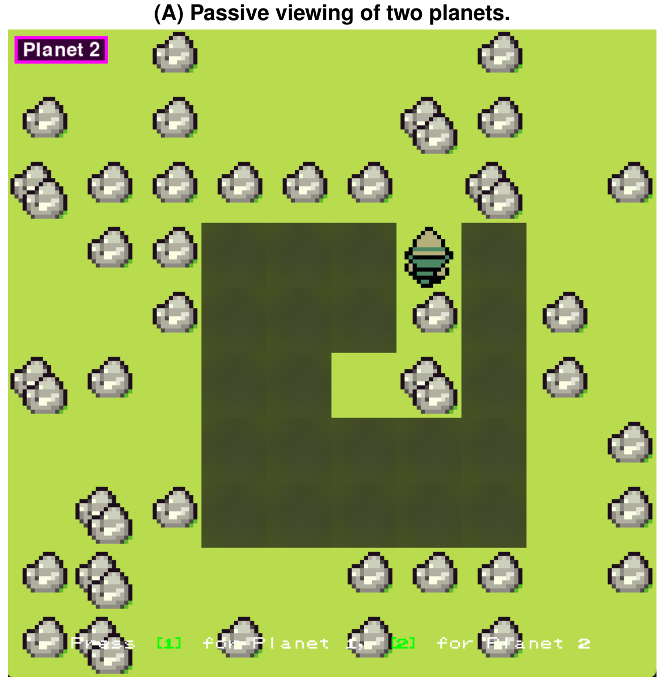

(B) Active exploration with an agent.

Figure 1: Gamified environments. (**A**) Participants briefly viewed each planet (0.5 sec) to form an initial belief about which had more (or fewer) stones based on limited, peripheral information; the center remained hidden. (**B**) After committing to a belief, they explored both planets further (5 sec) by walking around with an agent to uncover new evidence in the previously hidden area (i.e., tiles with potential stones).

Crucially, this bias only emerged when they controlled the sampling themselves. While this shows how biases guide resampling of familiar information, it remains unclear whether similar biases guide the search for *new* information — a more predictive form of belief-driven exploration.

Here, we ask whether initial beliefs shape not only how humans resample previously seen information, but also how they explore new evidence. We move from free resampling (i.e., switching between familiar stimuli to double-check initial impressions — **"Did I see that right?"**) to active exploration of new evidence (**"If I'm right, what else should I see?"**). To test this, we adapted a classic experimental design into a richer, gamified setting in which participants judged which of two planets contained more (or fewer) stones (**Figure 1**). This approach embeds belief updating in a more ecologically-valid setting while maintaining experimental control, and could be extended to artificial agents in future work.

Methods

Participants' task was to determine which of two planets contained more (or fewer) stones. The "baseline" planet always contained 50 stones, while the other contained either 40, 45, 50, 55 or 60 stones. Per trial, the baseline planet was randomly assigned to be either Planet 1 or Planet 2 (left or right on the screen; **Figure 1A**). The experiment consisted of four blocks of 15 trials each, alternating between "more" and "less" conditions. In total, the experiment consisted of 60 trials, each with three phases (**Figure 2**). Phase 3 differed critically from Phase 2 in that participants could actively explore the environment (by making the agent walk around) to reveal new, previously hidden, information (**Figure 1B**).

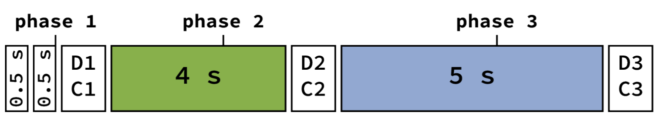

Figure 2: Trial structure. Each trial contained 3 phases, that each ended in a decision (D) and confidence rating (C). In Phase 1, two planets were briefly presented sequentially, followed by D1 and C1. In Phase 2, participants freely resampled the visual input from both planets, after which they again reported D2 and C2. In Phase 3, participants actively explored the planets by navigating an agent to uncover new information, followed by D3 and C3.

Results

Participants successfully identified the correct planet in most trials (**Table 1**). To test how initial belief guides subsequent information gathering, we analyzed how much time participants spent per planet, based on their initial choice and associated

Table 1: Participant accuracy in Phases 2 and 3. Note that percentages reflect correct decisions among non-tied trials.

Participant	% Correct Phase 2	% Correct Phase 3
Participant 1	86.4	79.2
Participant 2	84.2	68.8
Participant 3	84.2	89.6
Participant 4	72.7	68.8

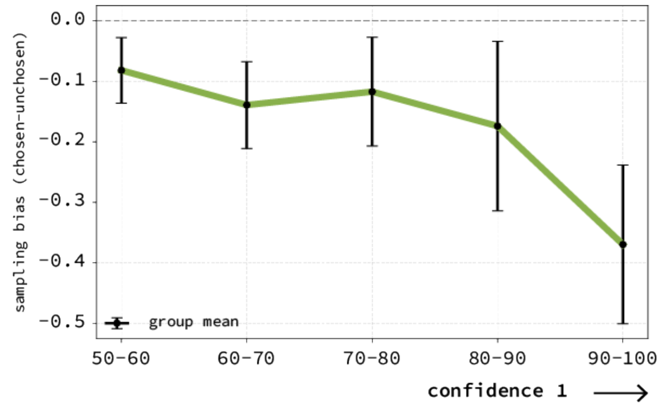

(A) Phase 2. Trend for disconfirmatory sampling.

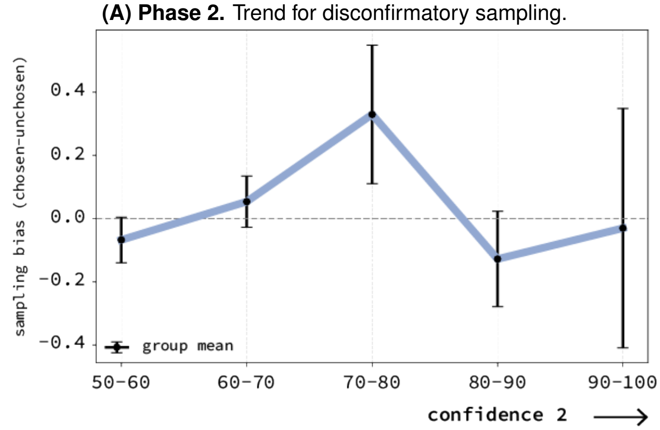

(B) Phase 3. No systematic bias; equal sampling around zero.

Figure 3: Sampling patterns. Sampling bias (difference in seconds) as a function of initial decision confidence.

confidence (**Figure 3**). Sampling bias was defined as the time difference in seconds between the chosen and unchosen planet (so negative values indicate more time spent on the *unchosen* planet). In Phase 2, participants showed a disconfirmation bias: spending more time on the unchosen planet, especially when confident (**Figure 3A**). In Phase 3, no clear bias emerged, although mid-confidence decisions showed a slight preference for the *chosen* planet (**Figure 3A**). These patterns held across both "more" and "less" trial types.

Outlook

These preliminary findings offer a first look at how initial belief and confidence shape sampling behaviour. While we so far did not find the sampling bias in the confirmatory direction, as we hypothesized, we do see an interesting trend for disconfirmatory sampling. This was only observed in the resampling of old evidence, but not in the exploration of new information. We speculate that belief-driven biases may be more likely when evidence is easily accessible (e.g., via the eyes) and reduced when actions require more deliberation and effort (e.g., via button presses). This current work lays the foundation for many exciting avenues for future research into how beliefs shape decision-making and exploration.

References

Bronfman, Z. Z., Brezis, N., Moran, R., Tsetsos, K., Donner, T., & Usher, M. (2015). Decisions reduce sensitivity to subsequent information. *Proceedings of the Royal Society B: Biological Sciences*, *282*(1810), 20150228.

Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, *36*(3), 181–204.

Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences*, *360*(1456), 815–836.

Kaanders, P., Sepulveda, P., Folke, T., Ortoleva, P., & De Martino, B. (2022). Humans actively sample evidence to support prior beliefs. *Elife*, *11*, e71768.

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of personality and social psychology*, *37*(11), 2098.

Pirolli, P., & Card, S. (1999). Information foraging. *Psychological review*, *106*(4), 643.

Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, *2*(1), 79–87.

Rollwage, M., Loosen, A., Hauser, T. U., Moran, R., Dolan, R. J., & Fleming, S. M. (2020). Confidence drives a neural confirmation bias. *Nature communications*, *11*(1), 2634.

Talluri, B. C., Urai, A. E., Tsetsos, K., Usher, M., & Donner, T. H. (2018). Confirmation bias through selective overweighting of choice-consistent evidence. *Current Biology*, *28*(19), 3128–3135.