How does social learning affect stable false beliefs?

Rheza Budiono (rheza@nyu.edu)
Catherine Hartley (cate@nyu.edu)
Todd Gureckis (todd.gureckis@nyu.edu)
Department of Psychology, 6 Washington Pl
New York, NY 10003 USA

Abstract

Learning traps are false beliefs that entrench themselves by discouraging the exploration required to correct them. In previous lab experiments, these learning traps have proven remarkably difficult to prevent. Here, we investigate whether learning traps remain stable in contexts in which both individual and social learning are possible. In two of our three experiments, we found that learners trapped by a false belief were significantly more likely to escape a learning trap when they were able to observe another decision-maker's choices (without observing their outcomes). However, social learning was not a panacea. Social learning was constrained by the challenge of inferring others' beliefs, and trapped learners struggled to learn from partners with sub-optimal decision rules, even when their partner's choices were informative. Collectively, these results suggest that while social learning can help overcome the limits of individual learning, learning from others comes with its own challenges and limitations.

Keywords: learning traps; social learning; observational learning; exploration; selective attention; rule inference

Introduction

False beliefs are difficult to correct when they prevent the exploration needed to correct them. For example, a person might try an Indonesian restaurant for the first time and have a bad experience, leading them to hold the belief that they dislike all Indonesian restaurants. In turn, this belief leads them to avoid Indonesian food, which prevents subsequent updates to their belief. If there exists an Indonesian restaurant that they would really enjoy, their current (false) belief prevents them from trying it. In such cases, we describe the learner as "trapped" because their false belief causes them to avoid potentially corrective experiences (March, 1991; Denrell & March, 2001; Erev, 2014).

Several recent experimental studies have explored the situations under which such "learning traps" emerge, establishing them as robust phenomena in individual learners (Rich & Gureckis, 2018; Li, Gureckis, & Hayes, 2021; Allidina & Cunningham, 2021; Liquin & Gopnik, 2022; Blanco, Turner, & Sloutsky, 2023; Bai, Griffiths, & Fiske, 2024). In these studies, learning traps arise from the links between an individual agent's beliefs, choices, and experiences. However, in real-world environments, people can learn from others' choices as well as their own. Here, we investigate whether learners trapped by a false belief remain trapped when they can observe another person's choices.

Methods

Task

Our task builds upon an existing approach-or-avoid task in which participants often learned a stable false belief which represented a learning trap (Rich & Gureckis, 2018). On each trial, the participant was shown one bee and asked if they would like to approach or avoid it. If they approached a friendly bee, they would harvest honey (+1 point); if they approached a dangerous bee, they would get stung (-5 points). If they avoided a bee, they were neither rewarded nor punished, but they also could not learn anything about the bee. The bees varied along 4 binary features, and there are 16 unique bees. Unbeknownst to the participant, 2 of these 4 features could be used to perfectly predict which bees were friendly and which were dangerous, and a unique conjunction of these 2 features determined whether a bee was dangerous (Fig. 1A).

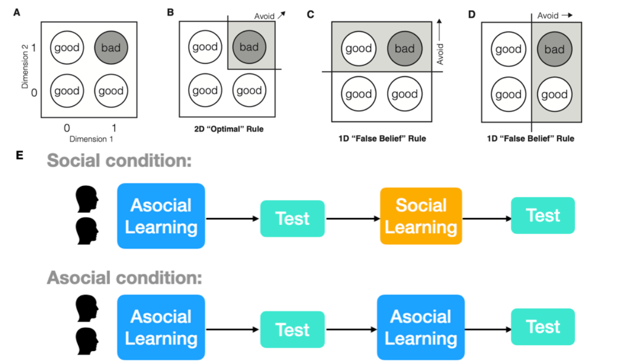

Figure 1: **Task structure. A.** Only stimuli with feature value 1 on both dimension 1 and 2 are "bad"; the rest are "good" **B.** The optimal "2D" decision rule, which approaches all good stimuli and avoids all bad stimuli. **C & D**. Two sub-optimal "1D" decision rules. **E.** Schematic of the experimental design.

Participants first completed an asocial learning phase and a test phase during which participants made choices without observing choice outcomes. Then, we split participants into a social learning condition, in which participants could observe a partner's choice at the end of each trial, or an asocial control condition in which they completed a second asocial learning phase(Fig. 1E). Finally, all participants completed another test phase.

Participants

Participants were recruited from Prolific and compensated at a rate of $15 per hour, plus a performance-based bonus of up to $4. In total, we collected data from 552 participants across 3 experiments.

Results

In each test phase, we classified participants as following either the optimal 2D decision rule, a sub-optimal 1D decision rule (a learning trap), or neither. Our classification threshold allowed for two deviations from the rule over the two full passes of the stimulus set.

In all three experiments, less than 20% of participants learned the optimal 2D decision rule (Fig. 1B), while about 40% of participants learned a sub-optimal 1D decision rule (Fig. 1C&D). These sub-optimal 1D rules are learning traps, because they prevent the exploration required to realize that they are sub-optimal. Indeed, a second asocial learning

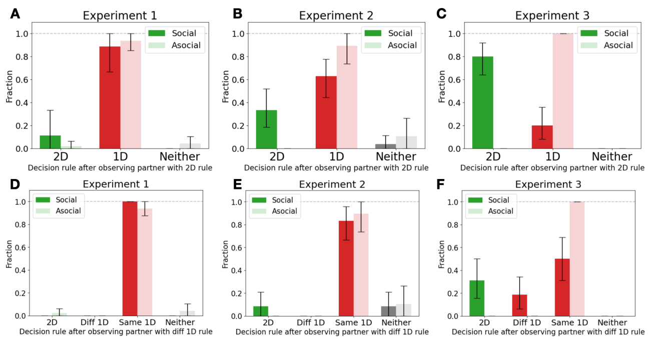

Figure 2: **Effect on trapped learners of observing a partner following either the optimal 2D decision rule (first row) or a sub-optimal 1D decision rule (second row).** We selected participants who displayed a sub-optimal 1D decision rule in the first test phase, and plot their decision rule distribution in the second test phase (95% CIs, bootstrapped with $1,000$ resamples). We compared this distribution for those in the social condition versus those in the asocial condition.

phase did not significantly reduce the number of asocial control participants who were trapped (Fig. 2). Thus, learning traps in the asocial learning setting were both prevalent and stable.

We next asked whether or not learners who were trapped by a false belief could break free by observing the choices of another decision-maker. A trapped learner could in principle escape by learning from either (1) a learner following the optimal 2D decision rule, or (2) a learner trapped by a different sub-optimal 1D decision rule. In both cases, their partner's choices would provide evidence against the learning trap because their partner would approach the bees the trapped observer is mistakenly avoiding.

Experiment 1

In Experiment 1, "social" participants observed the choices of another human participant playing alongside them in real-time. For learners trapped by a 1D decision rule, we found no significant effect of observing a partner with either the optimal 2D decision rule (Fig. 2A; $P > 0.3$[1]; $N_{\text{social}} = 7$, $N_{\text{asocial}} = 48$) or a different sub-optimal 1D decision rule (Fig. 2D; $P > 0.3$; $N_{\text{social}} = 8$, $N_{\text{asocial}} = 48$). However, our analyses were limited by the number of naturally occurring dyads of interest (i.e., trapped participants who were paired with either an optimal 2D decision-maker or someone with a different sub-optimal 1D decision rule).

Experiment 2

In Experiment 2, "social" participants observed the choices of a bot partner with a programmed decision rule. A third of trapped learners (who learned a sub-optimal 1D decision rule) learned the optimal 2D rule after observing a partner who followed the optimal 2D rule. In contrast, no trapped learners from the asocial control condition were able to learn the optimal 2D rule. This difference was significant (Fig. 2B; $P < 0.001$; $N_{\text{social}} = 27$, $N_{\text{asocial}} = 19$).

However, less than 10% of participants of trapped learners learned the 2D rule after observing a partner who followed a different 1D rule. Although this fraction is larger than the 0% of trapped learners from the asocial control condition who learned the 2D rule, this difference was not significant (Fig. 2E; $P > 0.1$; $N_{\text{social}} = 24$, $N_{\text{asocial}} = 19$).

Experiment 3

We hypothesized that social learning in Experiment 2 was limited by the difficulty of inferring another decision-maker's decision rule. We tested this hypothesis in Experiment 3, in which "social" participants additionally read a natural language description of the bot's decision rule (e.g. "avoid bees with both spots and antennae"). About 80% of trapped learners (who learned a sub-optimal 1D decisoin rule) learned the optimal 2D rule after reading a partner's optimal 2D rule and observing their choices. This was a significant improvement over the asocial condition (Fig. 2C; $P < 0.001$; $N_{\text{social}} = 25$, $N_{\text{asocial}} = 19$).

Furthermore, over 25% of trapped learners learned the 2D rule after reading a partner's different 1D rule and observing their choices. This was also a significant improvement over the asocial condition (Fig. 2F; $P = 0.002$; $N_{\text{social}} = 32$, $N_{\text{asocial}} = 19$). Compared to previous experiments, trapped learners were much more likely to escape the learning trap in Experiment 3, suggesting that social learning in previous experiments was indeed limited by the difficulty of inferring another decision-maker's decision rule.

Discussion

We found that people could escape learning traps by way of social learning. However, social learning was not a panacea. We found that participants were limited in their ability to learn from their partner by the difficulty of inferring their partner's decision rule.

Moreover, across all three experiments, trapped learners were much more likely to learn from the choices of a partner following an optimal decision rule versus a different sub-optimal one that was nonetheless informative. We identified two possible explanations for why this might be the case: participants could have (1) employed a limited "copy or not" social learning strategy, or (2) refused to learn from someone they deemed generally incompetent (even when their beliefs were informative). In future work, we aim to further investigate these two possible explanations. In particular, we plan to build computational models to better understand how people with stable false beliefs may learn (or fail to learn) from observing others' choices (Toyokawa, Whalen, & Laland, 2019; Witt, Toyokawa, Gaissmaier, Lala, & Wu, 2024; Hawkins et al., 2023; Wisdom, Song, & Goldstone, 2013) and considering others' advice (Yaniv & Kleinberger, 2000; Hawthorne-Madell & Goodman, 2019).

[1] All significance tests are one-sided bootstrap tests (1, 000 resamples).

References

Allidina, S., & Cunningham, W. A. (2021). Avoidance begets avoidance: A computational account of negative stereotype persistence. *Journal of Experimental Psychology: General*, *150*(10), 2078–2099. doi: 10.1037/xge0001037

Bai, X., Griffiths, T., & Fiske, S. (2024, January). Multidimensional Stereotypes Emerge Spontaneously When Exploration is Costly. doi: 10.31234/osf.io/mbuhv

Blanco, N. J., Turner, B. M., & Sloutsky, V. M. (2023, February). The benefits of immature cognitive control: How distributed attention guards against learning traps. *Journal of Experimental Child Psychology*, *226*, 105548. doi: 10.1016/j.jecp.2022.105548

Denrell, J., & March, J. G. (2001). Adaptation as Information Restriction: The Hot Stove Effect. *Organization Science*, *12*(5), 523–538.

Erev, I. (2014, October). Recommender systems and learning traps. In *Proceedings of the first international workshop on decision making and recommender systems* (pp. 38–41). Bolzano, Italy: Free University of Bozen-Bolzano.

Hawkins, R. D., Berdahl, A. M., Pentland, A. S., Tenenbaum, J. B., Goodman, N. D., & Krafft, P. M. (2023, October). Flexible social inference facilitates targeted social learning when rewards are not observable. *Nature Human Behaviour*, *7*(10), 1767–1776. doi: 10.1038/s41562-023-01682-x

Hawthorne-Madell, D., & Goodman, N. D. (2019, January). Reasoning about social sources to learn from actions and outcomes. *Decision*, *6*(1), 17–60. doi: 10.1037/dec0000088

Li, A. X., Gureckis, T. M., & Hayes, B. K. (2021). Can losses help attenuate learning traps? *Cognitive Science Society*.

Liquin, E. G., & Gopnik, A. (2022, January). Children are more exploratory and learn more than adults in an approach-avoid task. *Cognition*, *218*, 104940. doi: 10.1016/j.cognition.2021.104940

March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science*, *2*(1), 71–87.

Rich, A. S., & Gureckis, T. M. (2018, September). The limits of learning: Exploration, generalization, and the development of learning traps. *Journal of Experimental Psychology: General*, *147*(11), 1553. doi: 10.1037/xge0000466

Toyokawa, W., Whalen, A., & Laland, K. N. (2019, February). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nat Hum Behav*, *3*(2), 183–193.

Wisdom, T. N., Song, X., & Goldstone, R. L. (2013). Social learning strategies in networked groups. *Cognitive Science*, *37*(8), 1383–1425. doi: 10.1111/cogs.12052

Witt, A., Toyokawa, W., Gaissmaier, W., Lala, K., & Wu, C. M. (2024, February). Social learning with a grain of salt. doi: 10.31234/osf.io/c3fuq

Yaniv, I., & Kleinberger, E. (2000, November). Advice Taking in Decision Making: Egocentric Discounting and Reputation Formation. *Organizational Behavior and Human Decision Processes*, *83*(2), 260–281. doi: 10.1006/obhd.2000.2909