

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Belief updating patterns and social learning in stable and dynamic environments

Permalink

<https://escholarship.org/uc/item/4bh1939r>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Papakonstantinou, Trisevgeni

Raihani, Nichola J

Lagnado, David

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Belief updating patterns and social learning in stable and dynamic environments

Trisevgeni Papakonstantinou¹, Nichola Raihani^{1,2} & David Lagnado¹

¹University College London
Department of Experimental Psychology
26 Bedford Way, London WC1H 0APA

²University of Auckland
School of Psychology
23 Symonds St, Auckland, 1010, New Zealand

Abstract

Humans are resistant to changing their beliefs even in the face of disconfirming evidence. The Bayesian brain theory suggests that we should update our beliefs optimally in light of new evidence, but recent research indicates that belief formation is far from the Bayesian ideal. Individuals can exhibit “stronger-than-rational” updating or be resistant to revising their beliefs. The present study proposes a novel paradigm to explore perceptions and preferences for belief updating patterns in stable and dynamic stochastic environments, using an advice-taking paradigm. In an experiment (N=567) based on a fishing task, we introduce three advisor characters representing formal updating models: Bayesian, Volatile and Rigid. We find that participants exhibit higher trust for the Bayesian advisor than the Rigid advisor, in the stable but not changeable environment conditions. In the changeable environment, participants exhibit higher trust for the Volatile advisor, compared to both the Bayesian and Rigid advisors. The findings also suggest that participants’ own learning closely mimics the pattern of the Volatile model. This study illustrates that people can differentiate between Bayesian updating, and its “stronger-than” and “weaker-than” variations, and exhibit preferences for these updating patterns, in different environment structures.

Keywords: belief updating; advice taking; social learning; mental models; Bayesian inference

Introduction

It’s now received wisdom that evidence alone is insufficient to drive belief revision (Ecker et al., 2022). Altogether, both behavioural experiments and studies using naturalistic data seem to suggest that humans are averse to epistemic change and that our beliefs are remarkably immovable, even in the face of disconfirming evidence (e.g., Kube & Rozenkrantz, 2021). This phenomenon could be largely attributed to the fact that beliefs are not held in isolation but are rather part of a cohering system. These systems can be conceptualised as “intuitive theories” (Gerstenberg & Tenenbaum, 2020) or “cognitive models” (Goodman, Tenenbaum, & Contributors, 2016). Analogous to scientific theories, intuitive theories are explanatory frameworks composed of an ontology of concepts and causal links relating them, that represent and model the structure of the world within a domain. The ability to form, maintain, and adapt theories or cognitive models of the world is a crucial function of cognition. Determining how we should update these models when confronted with new (especially conflicting) evidence could be the key to unlocking the puzzle of learning.

A normative framework for belief revision

In recent years, Bayesian inference has been proposed as the underlying principle of neural computation, representing an “optimal” account of integrating new evidence into existing cognitive models (Knill & Pouget, 2004). The Bayesian brain theory proposes that the brain behaves like an intuitive scientist, inferring the causes of observed data and constantly updating its models of the world based on it. Grounded in Bayes’ theorem, the Bayesian brain theory suggests a procedure for arriving at optimally updated posterior beliefs. This account provides a formal normative framework of inference, where the probability of a cause (or hypothesis, h), given the evidence (e), is proportional to the probability of the evidence, given the cause, weighted by the probability of the cause prior to observing the evidence:

$$P(h|e) = P(h) \frac{P(e|h)}{P(e)}$$

Over the past two decades, this theory has gained considerable traction within the scientific community, with many researchers exploring and expanding upon its implications in many domains, including perception, planning, action, and cognition (Vilares & Kording, 2011; Friston, 2012). Particular foci of this research have been the areas of predictive processing and belief updating, which appeal to the notion of the Bayesian brain to describe evidence integration in stable and uncertain environments (Knill & Pouget, 2004; Bennett, 2015)

Beyond bias

Sensible as it sounds, the Bayesian brain theory has been a topic of heated debate in recent years, leading to a growing consensus that belief formation in both healthy and atypical individuals is far from Bayesian (Kube & Rozenkrantz, 2021; Williams, 2018). Despite some dissenting voices (Tappin & Gadsby, 2019), it seems that the human brain’s belief updating process often falls short of the optimally rational Bayesian ideal. Our pursuit of truth and coherence seems to often clash with other motivational forces. This suboptimal behaviour can manifest in two distinct directions, depending on a host of social and cognitive factors.

On the one hand, some individuals exhibit “stronger-than-rational” updating in different scenarios, whereby new ev-

idence is met with a disproportionate response that underweighs prior beliefs. This phenomenon, also known as “base-rate neglect” (Benjamin, Bodoh-Creed, & Rabin, 2019; Kahneman & Tversky, 1973), coexists with a “recency bias” (Ashinoff, Buck, Woodford, & Horga, 2022). Base-rate neglect is more likely to manifest in cases where prior beliefs are weakly connected to self-concept and are generally valence-independent.

On the other hand, in some cases, people may be more resistant to revising their beliefs, clinging stubbornly to prior convictions even in the face of overwhelming contrary evidence. A large body of behavioural and neural data has demonstrated this effect of conservatism in belief revision (Powell, 2022; Edwards, 1968). People are especially resistant to updating beliefs that challenge their self-concept (Dunning, Meyerowitz, & Holzberg, 1989; Sanitioso, Kunda, & Fong, 1990; Cohen, Aronson, & Steele, 2000). The well-established phenomenon of “confirmation bias”, (Nickerson, 1998) is also connected to this tendency to maintain the original content of beliefs and structure of cognitive models.

Consistent with this view is the account of asymmetric valence-dependence updating, manifested in optimism bias. People tend to update beliefs in a valence-dependent manner: they are more likely to integrate positive news into their beliefs while disregarding negative news (Sharot, 2011; Sharot, Korn, & Dolan, 2011; but see Shah, Harris, Bird, Catmur, & Hahn, 2016). These patterns of belief revision have important implications for understanding the cognitive processes underlying human decision-making and belief formation. However, a significant limitation in much of the existing literature is the lack of a comprehensive account for the revision of cognitive models beyond isolated beliefs with weak priors in single or small evidence samples.

Learning through others

Belief updating in the context of learning from others represents a distinct aspect of information processing. The influence of social norms on belief updating is well-documented, as evidence that aligns with prevailing norms or is presented as normative is more likely to be integrated into one’s beliefs (Vlasceanu & Coman, 2022; Orticio, Martí, & Kidd, 2022). When it comes to learning from social partners, the perception of their competence, reliability, and the level of trust placed in them are crucial factors that predict the extent to which individuals update their beliefs based on social information (Pilditch, Madsen, & Custers, 2020). This is especially relevant to the study of advice-taking behaviour. Hertz, Bell, and Raihani (2021) find that people exhibit a stronger inclination to follow advice rather than imitate observed choices, and this preference was influenced by their level of trust in the advisor. These findings, along with the study paradigm employed, lay the groundwork for further investigation into trust, perceptions of social partners, and the patterns of belief updating in synchronous, sequential learning contexts.

It is worth noting that while belief updating mechanisms

have been extensively studied in isolation from social learning contexts, there is a paucity of research directly exploring individuals’ perceptions of updating patterns in others. This represents an important gap in the current literature. By examining how individuals perceive and interpret updating patterns exhibited by social partners, we can gain a deeper understanding of the social dynamics and cognitive mechanisms that shape belief updating processes.

Methods

Preregistration

This study’s data collection procedure, experimental design, materials and measures, as well as the main hypotheses were registered on the Open Science Framework (https://osf.io/4zjvn/?view_only=9aec98dd9b874fdb99e333b863f4845a).

Design

The study used a 2×4 between subjects design. The two factors examined were “environment” and “advisor”. The context of the experiment was a game, where participants went on a “fishing vacation” at a site with two lakes. There, they were introduced to Taylor, a local fisherman who gave them the opportunity to choose a lake to fish from, and they fish from the other. The aim was to catch as many fish as possible in 35 rounds of choosing between two lakes to fish from.

The “environment” factor had two levels: “stable” and “changeable”. In the stable conditions the environment consisted of two lakes represented by a Poisson distribution of “fish” with λ of 5 and 6 respectively, that remain stable throughout the duration of the task. In the changeable conditions the environment consisted of the same two lakes, but there was a change after the 20th round, where the distributions shifted to λ of 5 and 2 respectively.

The “advisor” factor had four levels: “Bayesian”, “Volatile” (underweight priors), “Rigid” (overweighs priors), and a “no-advisor” learning condition. In the no-advisor learning conditions, participants were called to choose a lake from which to fish and subsequently received feedback from both lakes. All conditions started with five “learning” trials where no advice was given, identical to the no-advisor learning conditions. The advisors, however, were learning behind the scenes since round one. In the Bayesian advisor conditions, after completing the first five rounds, participants were introduced to Jamie, an expert player who would provide advice on which lake to fish from for the remaining rounds of the task. Jamie in this case was a learning model that adjusts its advice output each round by updating its estimate of each lake based on the feedback from the previous round using Bayes’ rule. In the volatile advisor conditions, the model applies the same strategy, but weighs its priors 0.2×, making each new datapoint heavily influential in their estimations. Conversely, the rigid advisor model weighs its priors 5×, resulting in the first few datapoints heavily influencing their advice in all subsequent rounds. All advisors started the game

with uniform priors that were equal for both lakes. The advice provided consisted of a binary choice between the two lakes with no additional information shown to participants.

The experiment consisted of three phases: learning, testing, and survey. The experimental task was inspired by Hertz et al. (2021), and the materials and procedure were adapted to address the research questions at hand.

Procedure

The study was administered through oTree, an open-source platform for web-based behavioural experiments (Chen, Schonger, & Wickens, 2016). Prior to the beginning of the study, participants were informed about the nature and requirements of their participation and provided their informed consent. Participants were randomly split into one of eight experimental conditions. They are informed that one lake is more profitable than the other and that their, or anyone else's, fishing does not affect the amount of fish in each lake. A performance bonus of £0.01 per five fish caught was paid to all participants.

In the learning phase of the study, all participants were introduced to the two lakes and did five rounds of fishing to establish some expectations about the populations of fish in each lake. In this phase, participants were called to choose a lake from which to fish and subsequently received information about the number of fish each lake yielded in that round, delivered in the form of their and the fisherman's catch. This process was repeated for 5 rounds.

During the testing phase of the study, participants in the "no-advisor" learning conditions performed the same task as in the learning phase for all subsequent rounds. Participants in the "advisor" conditions were introduced to their advisor at this stage. The task for subsequent rounds had three elements: advice, lake choice, feedback. The advice was visible in both the advice and choice page to ensure participants were aware of it while making the choice. This process was repeated for 30 rounds. After rounds 10, 20 and 35 during this phase, participants were asked to indicate which lake they are more likely to get the most fish from on a sliding scale ranging from 0 to 100 (0 being Bagel, 100 being Pacman). They were also asked to indicate the lake they would like to recommend to future players and provide a rationale for their recommendation in free text.

Finally, in the survey phase, participants were asked questions about their final estimates of the average number of fish in each lake and their perceptions of the advisor. Study data is publicly available on OSF (https://osf.io/4zjvn/?view_only=9aec98dd9b874fdb99e333b863f4845a). The study falls within the remit of the approval given by the UCL Research Ethics Committee to the Causal Cognition Laboratory (study number: EP/2018/005).

Materials and measures

Experimental materials. The experimental materials, including the basic task structure, prose, and graphics, were

adapted from Hertz et al. (2021).

Environment. In all conditions, the environment consisted of two lakes, lake "Bagel" and lake "Pacman", each represented by a Poisson distribution of "fish". The Poisson distribution has a single parameter, λ , which represents the expected number of events that occur in a given time interval or space. The expected value and variance of a Poisson-distributed random variable are both equal to this parameter. For a random discrete variable x that follows the Poisson distribution, the probability of x is given by:

$$f(x) = \frac{\lambda^x e^{-\lambda}}{x!},$$

This distribution was chosen as it is suitable for representing natural count data, such as fish in lakes. In each round, a random process generated a number for each lake's distribution as that lake's "catch". In all rounds for the stable environment conditions and in rounds 1-20 for the changeable conditions the λ for lake Bagel was set to 5 and for lake Pacman to 6. For the changeable conditions, the λ shifted to 5 for lake Bagel and 2 for lake Pacman after round 20, and remained stable throughout rounds 21-35. The chosen parameters were determined based on two main criteria. Firstly, they needed to be believable and ecologically valid. Secondly, they needed to be appropriate as a stimulus in relation to the hypotheses that were being tested. They were finally tested in a pilot (N=98) to ensure the difference between the original lake payoffs was subtle, but eventually apparent after several rounds, and that the environmental shift was noticeable to participants.

Learning models. The learning models were utilised as stimulus for the participants to invoke judgements of different belief updating patterns, presenting as expert players providing advice. To specify the Bayesian model, we used the Poisson-Gamma conjugate pair, for sake of simplicity and computational tractability. Where observations x_1, \dots, x_i are a random sample from Poisson(Λ), and prior distribution for Λ is Gamma(α, β), the posterior distribution for Λ is Gamma(α_i, β_i), where:

$$\alpha_i = \alpha + \sum_{i=1}^n x_i, \beta_i = \beta + n$$

For the Bayesian model, the updating followed this exact formula. For the volatile model, a weighing factor was added, multiplying $\sum_{i=1}^n x_i$ by 0.2. For the rigid model, the weighing factor was 5. The expected value is then given by the following formula:

$$s_i = \frac{\alpha_i}{\beta_i^2}$$

Before the game started, all models started with uniform prior parameters $\alpha_0 = 4$ and $\beta_0 = 2$. n remained $n_i = 1$ for all conditions in all rounds. Each round and the α

and β parameters for each lake were updated according to the Gamma posterior hyperparameters formula (or its two variations) based on the number of fish caught from each lake in the previous round. Then, the expected values were calculated using the specified formula. Finally, the models “decided” on the lake recommendation by choosing the lake with the higher expected value so far. If the expected values happened to be equal the recommendation was chosen randomly.

Participants

A power analysis determined the minimum sample size for detecting two-way ANOVA effects as 560. Parameters included $\alpha = 0.05$, $\text{power} = 0.95$, $\text{effect size} = 0.25$, and 8 conditions (2×4). A final sample of 572 recruited via Prolific. Participant mean age was 40 years old (range = 18, 77) and 52% of participants were female. All participants were paid a base rate of £9/hour, and bonus of £0.01 per five fish caught was paid to all participants.

Results

Performance trust A 2×3 between-subjects ANCOVA, including advisor accuracy as a covariate, found a small but significant main effect of environment change ($F(1,419) = 4.15$, $p = .042$, $\eta^2 = .01$) on advice-taking, indicating that the mean proportion of rounds in which participants followed the advisor’s recommendation was higher for the stable condition ($M = 74.4\%$, $SD = 17.7$), than the change condition ($M = 71\%$, $SD = 17.7$). Contrary to our prediction, there was no significant main effect of advisor pattern. The main effect of advisor accuracy was large and significant; the more accurate the advisor, the higher the likelihood of following the advice ($F(1,419) = 70.76$, $p < .001$, $\eta^2 = .14$). Advisor accuracy was calculated as the proportion of rounds where the advisor recommended the lake with higher (or equal) payoff in a given round for each participant.

There was a significant but small interaction effect between environment change and advisor pattern ($F(2,419) = 4.22$, $p = .015$, $\eta^2 = .02$). Post hoc comparisons using the t-test with Bonferroni correction indicated that advice taking was significantly lower ($p < .05$) for the Rigid advisor in the environment change condition than for any other combination of the two factors, except compared to the Volatile advisor in the stable environment condition. There were no other significant pairwise comparisons. Figure 1 illustrates these effects.

The results differ for the secondary outcome measures of performance trust: perceptions of competence and trustworthiness. Neither environment condition or advisor pattern had significant effects on ratings of the advisor’s competence and trustworthiness. Advisor accuracy had a significant and large positive effect on both outcome measures; more accurate advisors were rated as more competent ($\chi^2(1,419) = 61.07$, $p < .001$) and more trustworthy ($\chi^2(1,419) = 65.86$, $p < .001$).

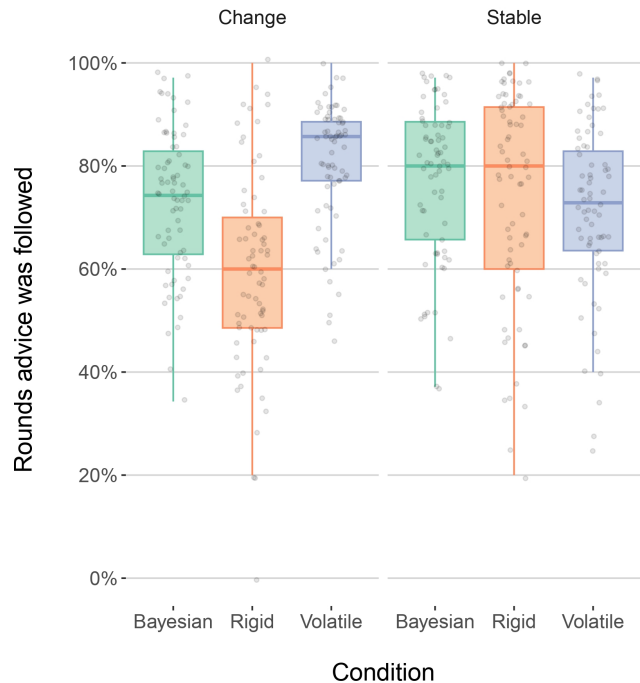


Figure 1: Advice-taking by condition.

Identification of “good” lake Results of a 2×4 ANOVA indicate that neither the advisor type nor environment change had a significant effect on the overall task performance of participants, as measured by the proportion of rounds in which they selected the “good” lake. A 2×4 ANOVA found a large and significant main effect of environment change ($F(1,559) = 158.02$, $p < .001$, $\eta^2 = .22$) on post-hoc identification of the “good” lake, as measured by a fixed-pie scale indicating the likelihood of getting the higher yield from each of the two lakes. Likelihood attributed to the “good” lake was overall higher in the stable condition $M = 67.7\%$, $SD = 25.1$ compared to the change condition $M = 38.4\%$, $SD = 29.9$. There was no significant effect of advisor condition or the interaction term.

Choice stochasticity A 2×4 ANOVA found a moderate and significant main effect of environment change ($F(1,538) = 84.28$, $p < .001$, $\eta^2 = .14$) on choice stochasticity, with stochasticity being higher in the environment change condition $M = 0.90$, $SD = 0.16$ compared to the stable condition $M = 0.74$, $SD = 0.25$). There was no significant effect of advisor condition or the interaction term on choice stochasticity.

Learning pattern A preliminary analysis was performed to explore the differences in learning patterns among participants and formal models. In Figure 2, the learning patterns of each formal model are compared to the actual learning pattern of participants who were not given any advice (No Advice conditions). Learning is approximated by calculating the proportion of individual agents that chose the “good” lake in

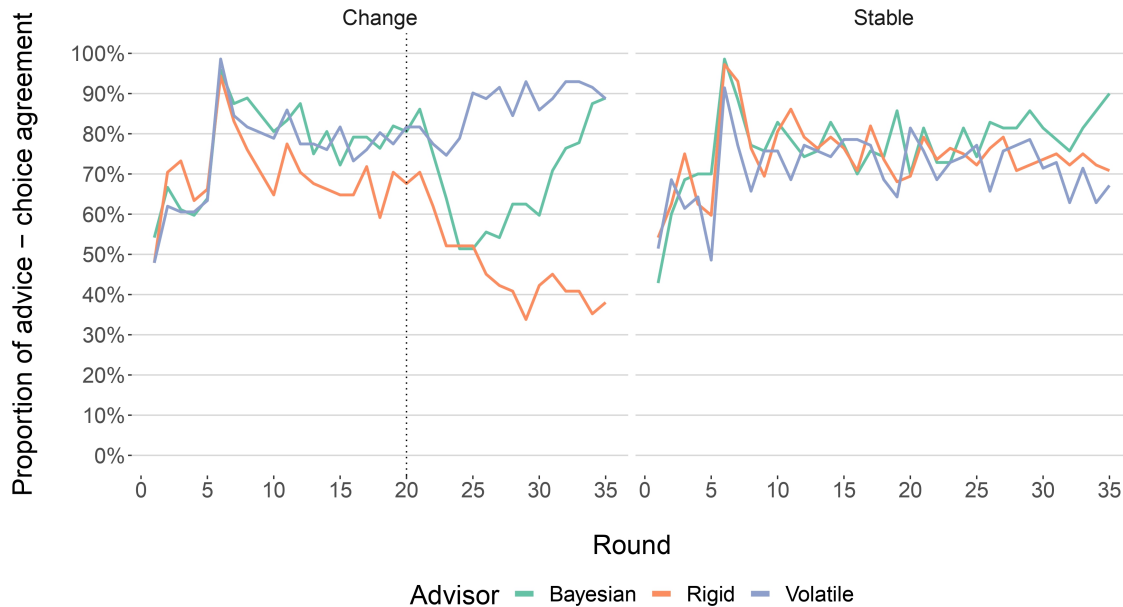


Figure 2: Advice-taking across rounds by condition.

a given round.

In a 2×4 ANOVA, significant main effects were found for environment change and agent type, as well as their interaction. The main effect of agent type was large and significant ($F(3, 272) = 19.38, p < .001, \eta^2 = .18$). Post-hoc comparisons reveal significant differences between all agent types on learning pattern (with Bonferroni-adjusted $t(272)$ ranging from -3.87 to $7.57, p < .05$), but not between the Volatile model and Participants ($t(272) = 0.77, p = 1.000$). This finding contradicts our initial hypothesis, which predicted that participants' learning would be most similar to the Bayesian model. Based on the results, we cannot reject the null hypothesis that there is a difference between the learning of participants and the Volatile model.

Discussion

In the present research, we compared different strategies of information integration to develop a better understanding of the factors influencing belief updating and perceptions of others' learning. Results suggested that, overall, people tend to rely more on social partners in stable environments rather than dynamic ones, albeit to a small extent. The interaction between environment change and advisor pattern indicated that adaptability was important in uncertain, changing environments. Participants exhibited higher trust in the Bayesian and Volatile advisors, though there was no effect on perception of competence and trustworthiness when controlling for advisor accuracy. The study did not find significant effects of advisor condition on task performance, but environment change affected preference for the "good" option and increased choice stochasticity.

The results of the present study suggest that participants'

advice-taking behaviour is influenced by the combination of environment and advisor pattern. The finding that participants were more likely to follow the advisor's recommendation in the stable environment is consistent with previous research showing that people are more likely to rely on heuristics and cognitive shortcuts in familiar, stable environments (Gigerenzer, 2020). The significant interaction between environment change and advisor pattern suggests that the effect of advisor pattern on advice-taking is modulated by the environment. The finding that advice-taking for the Rigid advisor was significantly lower in the environment change condition seems to suggest that there is a penalty for being rigid in dynamic environments rather than a benefit in being flexible or optimally Bayesian in any type of environment. This is consistent with previous research showing that people are more likely to update their beliefs in response to new information when they perceive the environment to be unpredictable or volatile (see Nassar, Wilson, Heasley, & Gold, 2010; Fenneman & Frankenhuis, 2020).

Overall, these findings have important implications for understanding how change and the introduction of contradicting evidence play into preferences for updating patterns. This study contributes to understanding the mechanisms of belief updating, particularly in sequential tasks, and expands on previous research by exploring social learning aspects and comparing participant behaviour to formal models. We illustrate that people can differentiate between Bayesian updating, and its "stronger-than" and "weaker-than" variations, and exhibit preferences for these updating patterns, in different environment structures. This study differs from previous studies on updating in three main ways. Firstly, we indirectly study the mechanisms involved in belief updating, by examining par-

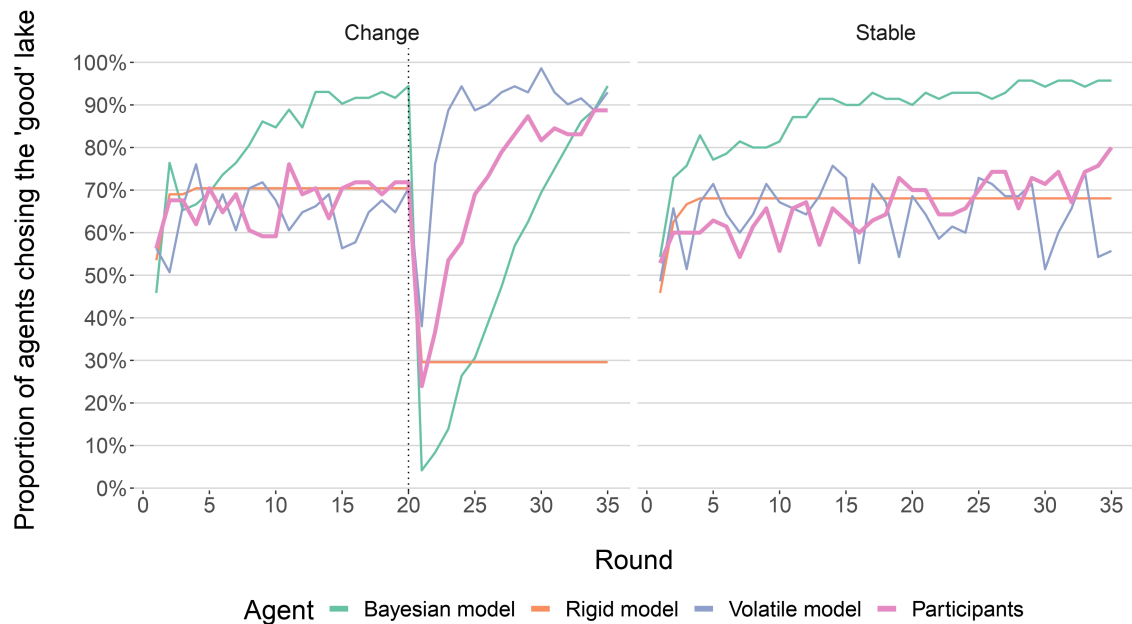


Figure 3: Learning pattern for each model type compared to the participants by environment condition. Participant learning is calculated based on participants in the “No Advice” independent learning condition.

participants’ reactions to the updating of social partners. This social learning element of updating has not been explored previously. Secondly, we directly compare different patterns of updating in a task involving sequential information, rather than single instances. Finally, we also collect data to compare participants own updating behaviour to that of formal models in the same task.

We show preliminary evidence that humans seem to importantly diverge from optimal Bayesian updating and adopt a more volatile strategy, underweighing priors and putting excessive weight on new evidence, regardless of whether the environment is stable or dynamic. This finding sheds light on the mechanisms of belief updating in sequential tasks, and is in line with Ashinoff et al. (2022). It contradicts findings that suggest humans prefer volatile updating only in dynamic or noisy environments (Nassar et al., 2010; Piray & Daw, 2020; but see Findling, Skvortsova, Dromnelle, Palminteri, & Wyart, 2018). This finding however may be largely explained by the fact that participants have no priors or expectations about the environment of the task.

It must be noted that the study’s conclusions regarding optimal updating and any deviations from it are constrained by the fact that Bayesian learning might not be appropriately characterised as “optimal” or “rational” in dynamic environments. Bayesian models assume an environment that is static over time. That means that in the changeable environment condition, we are not actually testing for the optimal model. The Volatile Bayesian model (one that assumes a constant drift in the payoffs of the two lakes) would be much closer to the correct assumptions about the environment, and as such represent optimal adaptability, as compared to the other two.

The results both on trust and learning reinforce that suggestion. Future research should explore the use of models such as a change-point detection model, which would more faithfully correspond to optimal updating in a dynamic environment, in the way that Bayesian updating does in a stable one. Thus, we might conclude that individuals tend to engage in optimal Bayesian updating, but are more likely to adapt their mental models to accommodate the dynamic nature of their environment, always assuming the possibility of change.

There are a number of additional limitations to consider when interpreting the findings of this study. One of the main limitations of the study is related to the measures used. Although advice-taking provides a robust behavioural measure of trust, in this paradigm, it is not clear whether participants always followed the advice or made choices based on their pre-existing beliefs. Therefore, some instances of advice-taking may not accurately reflect actual trust in the advisor. This ambiguity could potentially be addressed in future studies by including additional measures to assess the degree to which participants rely on the advice provided or seek it. Another limitation relevant to the experimental paradigm lies in the limited exploration of parameter values for the environment and the advisor models. Future studies may benefit from exploring a wider range of parameter values for the environment and using a more diverse set of advisor models to better capture the complexity of belief updating patterns in real-world settings.

These findings contribute to our understanding of belief updating and preferences for updating patterns in different environments. Further research is needed to explore the underlying mechanisms and reasons behind these effects.

References

- Ashinoff, B. K., Buck, J., Woodford, M., & Horga, G. (2022, 12). The effects of base rate neglect on sequential belief updating and real-world beliefs. *PLOS Computational Biology*, *18*(12), 1-37. Retrieved from <https://doi.org/10.1371/journal.pcbi.1010796> doi: 10.1371/journal.pcbi.1010796
- Benjamin, D. J., Bodoh-Creed, A. L., & Rabin, M. (2019). Base-rate neglect: Foundations and implications..
- Bennett, D. (2015). The neural mechanisms of bayesian belief updating. *Journal of Neuroscience*, *35*(50), 16300–16302. Retrieved from <https://www.jneurosci.org/content/35/50/16300> doi: 10.1523/JNEUROSCI.3742-15.2015
- Chen, D. L., Schonger, M., & Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88-97. Retrieved from <https://www.sciencedirect.com/science/article/pii/S2214635016000101> doi: <https://doi.org/10.1016/j.jbef.2015.12.001>
- Cohen, G. L., Aronson, J., & Steele, C. M. (2000). When beliefs yield to evidence: Reducing biased evaluation by affirming the self. *Personality and Social Psychology Bulletin*, *26*, 1151 - 1164.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: the role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, *57*, 1082-1090.
- Ecker, U., Lewandowsky, S., Cook, J., Schmid, P., Fazio, L., Brashier, N., ... Amazeen, M. (2022, 01). The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, *1*, 13-29. doi: 10.1038/s44159-021-00006-y
- Edwards, W. (1968). *Conservatism in human information processing*.
- Fenneman, J., & Frankenhuys, W. (2020, 03). Is impulsive behavior adaptive in harsh and unpredictable environments? a formal model. *Evolution and Human Behavior*, *41*. doi: 10.1016/j.evolhumbehav.2020.02.005
- Findling, C., Skvortsova, V., Dromnelle, R., Palminteri, S., & Wyart, V. (2018). Computational noise in reward-guided learning drives behavioral variability in volatile environments. *bioRxiv*. Retrieved from <https://www.biorxiv.org/content/early/2018/10/11/439885> doi: 10.1101/439885
- Friston, K. (2012). The history of the future of the bayesian brain. *NeuroImage*, *62*(2), 1230-1233. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1053811911011657> (20 YEARS OF fMRI) doi: <https://doi.org/10.1016/j.neuroimage.2011.10.004>
- Gerstenberg, T., & Tenenbaum, J. B. (2020). Intuitive theories. *Encyclopedia of Creativity, Invention, Innovation and Entrepreneurship*.
- Gigerenzer, G. (2020). How to explain behavior? *Topics in Cognitive Science*, *12*(4), 1363-1381. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1111/tops.12480> doi: <https://doi.org/10.1111/tops.12480>
- Goodman, N. D., Tenenbaum, J. B., & Contributors, T. P. (2016). *Probabilistic Models of Cognition* (Second ed.). <http://probmods.org/v2>. (Accessed: 2023-5-6)
- Hertz, U., Bell, V., & Raihani, N. (2021, 10). Trusting and learning from others: immediate and long-term effects of learning from observation and advice. *Proceedings of the Royal Society B: Biological Sciences*, *288*. doi: 10.1098/rspb.2021.1414
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*, 237-251.
- Knill, D. C., & Pouget, A. (2004). The bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*, 712-719.
- Kube, T., & Rozenkrantz, L. (2021). When beliefs face reality: An integrative review of belief updating in mental health and illness. *Perspectives on Psychological Science*, *16*(2), 247-274. Retrieved from <https://doi.org/10.1177/1745691620931496> (PMID: 32818386) doi: 10.1177/1745691620931496
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, *30*(37), 12366–12378. Retrieved from <https://www.jneurosci.org/content/30/37/12366> doi: 10.1523/JNEUROSCI.0822-10.2010
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*(2), 175-220. Retrieved from <https://doi.org/10.1037/1089-2680.2.2.175> doi: 10.1037/1089-2680.2.2.175
- Orticio, E., Martí, L., & Kidd, C. (2022, 07). Social Prevalence Is Rationally Integrated in Belief Updating. *Open Mind*, *6*, 77-87. Retrieved from https://doi.org/10.1162/opmi_a.00056 doi: 10.1162/opmi_a.00056
- Pilditch, T. D., Madsen, J. K., & Custers, R. (2020). False prophets and cassandra's curse: The role of credibility in belief updating. *Acta Psychologica*, *202*, 102956. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0001691819300952> doi: <https://doi.org/10.1016/j.actpsy.2019.102956>
- Piray, P., & Daw, N. D. (2020). A simple model for learning in volatile environments. *PLoS Computational Biology*, *16*.
- Powell, D. (2022). A descriptive bayesian account of optimism in belief revision. In H. R. . V. R. J. Culbertson A. Perfors (Ed.), *Proceedings of the 44th annual conference of the cognitive science society*.
- Sanitioso, R. B., Kunda, Z., & Fong, G. T. (1990). Motivated recruitment of autobiographical memories. *Journal*

- of personality and social psychology*, 59 2, 229-41.
- Shah, P., Harris, A. J., Bird, G., Catmur, C., & Hahn, U. (2016). A pessimistic view of optimistic belief updating. *Cognitive Psychology*, 90, 71-127. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0010028516301177> doi: <https://doi.org/10.1016/j.cogpsych.2016.05.004>
- Sharot, T. (2011). The optimism bias. *Current Biology*, 21(23), R941-R945. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0960982211011912> doi: <https://doi.org/10.1016/j.cub.2011.10.030>
- Sharot, T., Korn, C., & Dolan, R. (2011, 11). How unrealistic optimism is maintained in the face of reality. *Nature neuroscience*, 14, 1475-9. doi: 10.1038/nn.2949
- Tappin, B. M., & Gadsby, S. (2019). Biased belief in the bayesian brain: A deeper look at the evidence. *Consciousness and Cognition*, 68, 107-114. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1053810018305075> doi: <https://doi.org/10.1016/j.concog.2019.01.006>
- Vilares, I., & Kording, K. (2011, April). Bayesian models: The structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences*, 1224(1), 22-39. doi: 10.1111/j.1749-6632.2011.05965.x
- Vlasceanu, M., & Coman, A. (2022). The impact of social norms on health-related belief update. *Applied Psychology: Health and Well-Being*, 14(2), 453-464. Retrieved from <https://iaap-journals.onlinelibrary.wiley.com/doi/abs/10.1111/aphw.12313> doi: <https://doi.org/10.1111/aphw.12313>
- Williams, D. (2018). Hierarchical bayesian models of delusion. *Consciousness and Cognition*, 61, 129-147. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1053810017306219> doi: <https://doi.org/10.1016/j.concog.2018.03.003>