Disentangling Suboptimal Updating: Complexity, Structure, and Sequencing.[*]

Marina Agranov[†]   Pëllumb Reshidi[‡]

November 8, 2023

Abstract

We study underlying reasons for the failure of individuals to adhere to Bayes' rule and decompose this departure into three elements: (i) task complexity, (ii) information structure, and (iii) timing of information release. In a series of controlled experiments, we systematically alter all three elements and quantify their magnitude. We link task complexity with the degree of non-linearity embedded in Bayesian updating. We experimentally explore this link and find empirical support for it.

1   Introduction

Across various contexts, updating beliefs is critical in the decision-making process. In the last few decades, economists and psychologists made significant progress in understanding how people engage with new information. In many situations, the standard Bayesian updating framework accurately describes the evolution of beliefs (Grether, 1978; Camerer, 1987; Charness and Levine, 2005).[1] At the same time, there are frequent and consistent deviations from Bayesian updating. These deviations persist even when people have ample opportunities to learn (Esponda et al., 2023) and when stakes are very high (Gneezy et al., 2023). Moreover, these deviations are prevalent even among professionals who frequently deal with such issues (Benjamin, 2019). What makes some belief-updating tasks more difficult and more mistake-prone than others? This is the focus of our paper.

We present findings from a series of lab experiments involving belief updating tasks, where we manipulate three essential factors: task complexity, information structure, and information sequencing. We show that adjusting these three factors can alter observed behavior from resembling Bayesian reasoning to exhibiting sizable deviations. We introduce a novel link between task complexity (difficulty) and the degree of non-linear calculations required for Bayesian updating, capturing the challenge of adapting to such nonlinearities. Information structure pertains to the idea that the same content of information can be presented in different ways. In particular, a scenario in which an individual has an informative prior can alternatively be re-formulated as one in which they start with an uninformative prior but receive the equivalent amount of information through a signal. Though mathematically equivalent, these expositions differ conceptually. Our experiment offers the first empirical evidence of this equivalence. Finally, information sequencing requires consideration since any task involving belief updating inherently involves a particular order of presenting information. Our experiment is designed to enable the manipulation of these three factors and measure the impact of each separately. While we study the influence of all three components, our primary emphasis and contribution centers around task complexity.

To see the idea behind our connection of complexity with non-linear calculations, consider the following scenario. A person tries to calculate the probability of having a particular disease. Initially, they believe there's an equal chance they have or don't have the disease. This person has access to two diagnostic tests; the first test has accuracy

---

[1]Not only humans employ Bayesian updating! Valone (2006) reviews experiments involving animals and concludes that a variety of them, across different ecological contexts, behave in a manner consistent with Bayesian updating.

85% and turns out to be negative, while the second has accuracy 80% and turns out to be positive.[2] Taking into account these outcomes, they do their best to form a correct posterior, incorporating all information at their disposal. If both tests had equal accuracy—say, 80%—it would be natural to conclude that there's a 50% chance of having the disease. However, since the first test is more precise by five percentage points, it's logical to give it slightly more weight and compute a posterior that leans more toward negative than positive. The Bayesian posterior, in this case, the probability that the individual has the disease, is 41.3%, seemingly a plausible value near which individuals' average posteriors may lie. Alternatively, consider a more extreme case, with the same setup and test outcomes but with test accuracies equal to 97% and 92%, respectively. As before, since the first test is more accurate than the second by five percentage points, people may reason they should weigh the first slightly more and, thus, compute a posterior that leans more toward negative than positive. However, in this case, the Bayesian posterior is 26.2%, drastically influenced by the first signal, although the gap between the signal accuracies is still only five percentage points. Examples become even more extreme for higher accuracy signals and/or as the gap between signal accuracies increases.

Intuitively, although the setup in both cases is identical up to parameters, it seems likely that people would make bigger mistakes in the second case. Our notion of complexity formalizes this intuition. We do so in different ways. First, we take a parsimonious approach, arguing that the difficulty in the examples above stems from the non-triviality of incorporating information from signals with different accuracies. The challenge is knowing how much more one needs to react to the higher accuracy signal than the lower accuracy one. This is not an easy question to answer. In the example above, the Bayesian posterior is linear—and thus, the marginal increase is constant—only when facing a single signal. For all other cases, as the accuracy of either signal changes, so does the marginal reaction towards a signal, i.e., the Bayesian posterior is no longer linear. It is these nonlinearities in Bayesian updating that we believe contribute to the varying levels of difficulty. We classify regions where these nonlinearities are more pronounced as more difficult. We further offer two models we delegate to the appendix, one viewed as a stylized model and one as a behavioral one. In the first, we revisit Grether (1980), which has been extensively used in empirical work studying belief updating, and make minor changes incorporating the aforementioned concepts. In the second, we offer an alternative model resulting from an agent's inability to fully follow changes in the second derivative and thus only partially react to nonlinearities.

All three approaches lead to the same testable implications: an increase in the level of

---

[2]By accuracy, we refer to the probability that the test correctly identifies the disease status.

and the difference between signal accuracies increases deviations from Bayesian posteriors. It is these testable implications that we bring to the data.

We conducted a series of preregistered experiments using a diverse online sample of participants through the Prolific online platform. Our data comprises decisions of approximately 1,200 participants across twelve treatments. These treatments differ in the information structure, information timing, and parameter values. With regard to information structure, we vary whether participants start with an informative prior and receive one signal or whether they start with an uninformative prior and receive two signals. We do so while keeping the information content unchanged. This allows us not only to study the impact of these different information structures but also to vary the timing of information. We utilize this possibility and thus speak to information sequencing by varying whether the information is delivered simultaneously or sequentially.

To capture the link between task complexity and non-linear calculations, we vary the parameter values we utilize. For our low-nonlinearity variant, we employ parameters originally used in the seminal paper by Kahneman and Tversky (1973), and in many papers studying base-rate neglect thereafter (Benjamin et al., 2016; Benjamin, 2019; Esponda et al., 2023). This constitutes our low-nonlinearity setup, in which both the level of and difference between signal accuracies are relatively low. Conversely, in our high-nonlinearity variant, we employ both a higher level of and a higher difference between signal accuracies. To further test the link between complexity and environment nonlinearities and to isolate the effects of the level and difference between signal accuracies, we run four additional treatments further varying in parameter values. In summary, we designed our experimental setup in order to uniquely identify and assess the impact of task complexity, information structure, and information sequencing on belief updating.

We have four main results.

First, our data strongly supports the link between complexity and environmental nonlinearities: as the level and difference between signal accuracies increase, so does the scale of participants' errors. Both these features play a distinct role in exacerbating errors. As a robustness check of our first result, we utilize a different dataset collected and analyzed by Enke and Graeber (2023), in which the authors link mistakes in belief-updating tasks to participants' cognitive uncertainty. Even when accounting for cognitive uncertainty, our results remain unchanged. We find it reassuring that our complexity metric proves versatile and replicates in a completely distinct sample.

Second, when faced with the standard base-rate neglect setup, i.e., an informative prior and a signal, participants make sizable mistakes and under-utilize the prior. This is consistent with the extensive base-rate neglect literature. However, when starting with

an uninformative prior and an equivalent amount of information delivered through two simultaneous signals, participants make fewer mistakes.[3] In fact, in the low-complexity environment, this later setup eradicates base-rate neglect from the onset of the treatment; observed posteriors are statistically indistinguishable from Bayesian predictions. We consider this to be an important finding, especially given the comprehensive body of work in both psychology and economics that establishes base-rate neglect as one of the most persistent deviations from Bayesian updating, showing minimal responsiveness to higher incentives Gneezy et al. (2023), ample feedback Esponda et al. (2023), or the use of more relatable contexts over mathematical constructs Gigerenzer and Hoffrage (1995).

Third, we investigate what drives the immediate correction of base-rate neglect in the low-complexity environment and its partial mitigation in the high-complexity environment. When transitioning from the standard setup with an informative prior and a single signal to an arrangement with an uninformative prior and two signals, two elements change. First, the structure of information changes; initially, information comes from a prior and a signal, whereas subsequently, all information is sourced solely from signals. Second, the timing of this information also differs; in the former scenario, information is sequential as the prior necessarily precedes the signal, while in the latter, information arrives simultaneously. To disentangle these effects, we contrast outcomes from treatments with simultaneous and sequential arrival of information. We find that changes in the structure of information affect individual-level behavior but leave aggregate observed posteriors unaltered. In contrast, the sequence in which information is received significantly impacts aggregate observed posteriors. Specifically, participants tend to overreact to more recent signals, a pattern which echoes the recency effect documented in the literature (Grether, 1992; Benjamin, 2019).

Fourth, we show that altering the sequence in which information is presented—either simultaneously or sequentially, and in the latter case, by varying the order of high and low accuracy signals—can reduce errors made by participants. However, the most effective approach for disseminating information varies based on task complexity. When complexity is low, simultaneous information release is most effective. However, when complexity is high, we leverage the recency bias to counteract the bias induced by task complexity, making a sequential release of information more advantageous.

---

[3]We do not suggest that re-framing the standard belief-updating task as one with an uninformative prior and two simultaneous signals is a practical solution. Rather, we employ this re-formulation to understand what drives sub-optimal updating. This exemplifies one of the notable strengths of experimental economics: the ability to create highly controlled, synthetic scenarios that may not exist outside the lab for the purpose of understanding behavior. See, for instance, Kagel and Levin (1993) who study behavior in the third-price auction in order to test the validity of Nash bidding theory.

Lastly, we carry out individual-level analysis and categorize participants into various groups using K-means clustering, a popular machine-learning technique. This method allows us to take a hands-off approach to identifying the optimal number of clusters in our dataset, the clusters themselves, and the sorting of participants into these clusters. The outcomes at the individual level, particularly the shifts in cluster assignments across treatments, reinforce the findings we observe in our aggregate-level analysis.

Taken together, our experiments provide evidence that nonlinearities are influential in belief updating tasks. We believe that people's limited capacity to fully internalize nonlinearities extends beyond the realm of belief updating and presents an intriguing avenue for future research.

The remainder of the paper is structured as follows. We survey the literature in Section 1.1. In Section 2, we lay out the conceptual framework. We dedicate Section 3 to the experimental design and procedures. In Section 4, we present the aggregate results of our experiment. In Section 5, we dive into individual-level analysis. Finally, we draw conclusions in Section 6.

1.1 Literature Review

Our paper relates to and builds upon several strands of literature. We discuss these strands below and highlight our contribution in comparison to the prior findings.

Complexity literature. There is a rapidly developing and intriguing literature on decision complexity. The debate surrounding the definition of complexity and the empirical methods for identifying it are still in their early stages and largely dependent on context. In the domain of rules, complexity is found to be influenced by the number of states and transitions required to implement such a rule (Oprea, 2020; Banovetz and Oprea, 2022; Camara, 2021). In the domain of lotteries, complexity has been linked to the number of distinct outcomes in the lottery support (Bernheim and Sprenger, 2020; Puri, 2022; Fudenberg and Puri, 2022), the cognitive difficulty of aggregating outcomes and objective probabilities into a single value (Oprea, 2022), and excess similarity between lotteries in the choice set (Enke and Shubatt, 2023). In the inter-temporal choice problems, complexity has been linked to the difficulty of evaluating streams of future payments (Enke et al., 2023).[4]

---

[4]Regardless of the domain, much of this literature has focused on determining whether known behavioral anomalies, such as probability weighting, loss aversion, present bias, hyperbolic discounting, and so on, are *preference-based* or, instead, are an empirical fingerprint of the *complexity* of the decision being made.

Closer to our setting is a recent paper by Ba et al. (2023), suggesting that the difficulty of incorporating new information depends on the size of the state space. The authors develop a behavioral model that incorporates two known psychological frictions, noisy cognition and representativeness, and show that underreaction to new information is predicted when the state space is simple (consists two states only), while the opposite is true when the state space is more complex.[5] This notion of complexity, however, does not have bite in the binary state case, which is the focus of the vast majority of papers in the literature, including ours. In contrast, our measure of complexity does not rely on the cardinality of the state space or the number of signals, for that matter, and can vary even in binary settings. We, therefore, view our work as complementary to the work of Ba et al. (2023) and proceed to formalize the connection between task complexity and the non-linearities embedded in Bayesian updating.

A related but distinct approach is proposed by Enke and Graeber (2023), who suggests a link between decisions and cognitive uncertainty. In Section 4.3, we show that these are two distinct phenomena that coexist and jointly determine the rate of mistakes in belief-updating tasks.

Finally, we note that our complexity notion relates to the literature that documents challenges and sub-optimal decisions apparent in environments with nonlinear features. Among the more prevalent patterns are the exponential growth bias, which is the tendency to underestimate compound growth processes prevalent in financial decisions (Wagenaar and Sagaria, 1975; Stango and Zinman, 2009; Levy and Tasoff, 2016, 2017), the schmeduling heuristics, which suggest simplified ways to construct mental representations of nonlinear incentive schemes (Rees-Jones and Taubinsky, 2020), and the difficulty in discounting atemporal payments that feature a large number of steps given the induced parameters (Enke et al., 2023).

Sub-optimal belief updating. There is a vast body of research in psychology and economics documenting errors in probabilistic reasoning that result in sub-optimal beliefs, i.e., beliefs that do not align with Bayesian predictions. The most relevant to our study are papers that document the effect information sequencing has on belief updating, e.g., primacy and recency effects (Pitz and Reinhold, 1968; Edenborough, 1975; Grether, 1992). Benjamin (2019) provides the most recent and comprehensive survey focusing on biases related to random samples and belief updating in general.

---

[5]See also Augenblick et al. (2023) who study the over- and under-inferences in belief-updating tasks with binary states and find that over-reaction is a robust finding for weak signals, while under-reaction is common for strong signals. The authors develop a theory of cognitive imprecision about signal informativeness to accommodate both patterns.

Base-rate neglect holds a prominent place in this literature as it is one of the more persistent phenomena and, therefore, one of the most frequently studied. First introduced by Kahneman and Tversky (1972) and Bar-Hillel (1980) and followed by many empirical papers scrutinizing this bias and theoretical models contemplating its origins (Benjamin, 2019; Benjamin et al., 2019). Of special interest is a recent paper by Esponda et al. (2023), which shows that even after ample opportunities to learn, base-rate neglect persists. The authors explore the main forces that hinder learning from feedback and find that mistakes that stem from misrepresentation of primitives of the environment are likely to be persistent. In this paper, we use base-rate neglect as our case study. However, we note that the insights we present are general and applicable to any setting involving probabilistic reasoning in the presence of new information.

What mitigates biased beliefs? Given the prevalence of biased beliefs and their importance in determining decisions, great efforts have been made to understand how responsive these biases are to various features of the environment and how to mitigate them.[6] Gneezy et al. (2023) study the role of financial incentives and find that base-rate neglect is largely unresponsive to stakes. Several papers document that the extent of base-rate neglect depends on whether the task is presented in terms of frequencies as opposed to probabilities (Koehler, 1996; Barbey and Sloman, 2007), whether the task is framed in an intuitive and contextual manner as opposed to an abstract way (Cheng and Holyoak, 1985; Gigerenzer and Hoffrage, 1995; Gneezy et al., 2023; Ganguly et al., 2000), and whether the task is presented as forecasting the future outcomes as opposed to updating beliefs about the state Fan et al. (2022). Esponda et al. (2023) find that information about primitives of the environment might hinder learning. Contrasting a treatment in which no primitives were provided with standard treatment, they find that in the former, elicited beliefs are eventually closer to the Bayesian posterior.[7]

We share with this literature the goal of understanding drivers of base-rate neglect and ways to mitigate it. While several manipulations from the papers discussed above reduce the bias to some extent, none come close to eliminating it fully. We examine a different manipulation and demonstrate that, for the typical parameters used in the literature, it eliminates base-rate neglect from the onset of the treatment.

---

[6]See also the 'nudge' literature, which explores how to steer people into making better choices (Thaler and Sunstein, 2008; Thaler and Benartzi, 2004; Madrian and Shea, 2001).

[7]See also a stream of recent papers that discuss the relationship between incorrect mental models and task complexity (Enke and Zimmermann, 2019; Enke, 2020; Graeber, 2023). Worth highlighting is also Esponda et al. (2023) who, in one of their treatments, provide subjects with empirical frequencies of the joint distribution of signals and outcomes. Their findings suggest this approach can significantly mitigate the issue of base-rate neglect. However, it is important to note that the effectiveness of this solution is contingent upon gathering data over multiple rounds.

2 Conceptual Framework

2.1 Setup

Consider a standard belief-updating task. The state is binary $\omega \in \{F, S\}$, denoting, for example, whether a project is a Failure or a Success. A decision-maker does not know the state but holds prior belief $P(F) = p_0$. They observe the realization of a signal s, which can take either a negative ($s = n$) or positive ($s = p$) value. The signal has accuracy θ_s, which summarizes the probability that it correctly reveals the state, $P(s = n|F) = P(s = p|S) = \theta_s$.

Upon observing a signal, Bayes' rule dictates that the updated beliefs in the form of a posterior-odds ratio can be written as

$$\frac{P(F|s)}{P(S|s)} = \frac{P(s|F)}{P(s|S)} \frac{P(F)}{P(S)} = \frac{P(s|F)}{P(s|S)} \frac{p_0}{1 - p_0}, \tag{1}$$

where $\frac{P(s|F)}{P(s|S)}$ is the base factor of the signal, and $\frac{p_0}{1-p_0}$ is the prior probability ratio. The posterior odds ratio underscores that both the prior and the signal contain valuable information, which the decision-maker uses to update their beliefs.

2.2 Information Structure

For our analysis, it is useful to distinguish between an informative and uninformative prior, where the latter is one that has minimal impact on the posterior. While identifying an uninformative prior is nontrivial in general, for our binary case, an uninformative prior is one that assigns equal probability to both states $p_0 = 1/2$.[8] With this prior, once an agent receives a signal, their posterior is fully determined by the value and accuracy of this signal. Specifically, if the accuracy of the signal is θ_s and the signal realization is negative (positive), the agent's posterior beliefs, as calculated by Bayes' rule, are that the project is a Failure (Success) with probability θ_s and a Success (Failure) with complimentary probability $1 - \theta_s$. Thus, the new information fully determines their distribution of beliefs. In this spirit, we intend to convey that the prior is uninformative: it does not interfere with the information from the signal, it is optimal to fully follow the signal.[9]

---

[8]The typical issues that make identifying an uninformative prior challenging in more general settings involve the prior not being a conjugate prior, being an improper prior, or conveying information after re-parametrization. To tackle these issues, different approaches have been taken, such as defining the uninformative prior as one that maximizes entropy, as one that maximizes the Kullback distance between the prior and posterior, as a limit in conjugate families, as priors that are robust to transformations, such as Jeffreys prior, and so on. In our binary setup, however, these issues are not present.

[9]In practice, these priors bridge the gap between Bayesian and frequentist econometricians, bringing them as close as possible. An uninformative prior, in essence, emphasizes that the data (information) is all

Henceforth, we call a prior *uninformative* if it assigns a probability of ½ to each state and call it *informative* otherwise.

Consider an agent with an informative prior $P(F) = p_0 > ½$. If this agent receives no other information, their posterior-odds ratio will be $\frac{P(F)}{P(S)} = \frac{p_0}{1-p_0}$. We can reinterpret this prior as a posterior originating from an initial uninformative prior $P(F) = \tilde{p}_0 = ½$, and a signal with accuracy $\theta_s = p_0$ whose realized value is negative.[10] In this case, the posterior-odds ratio will be

$$\frac{P(F|s = n)}{P(P|s = n)} = \frac{P(s = n|F)}{P(s = n|S)}\frac{P(F)}{P(P)} = \frac{\theta_s}{1 - \theta_s}\frac{½}{½} = \frac{p_0}{1 - p_0}$$

Thus, having a prior $p_0 > ½$ is equivalent to having an uninformative prior $\tilde{p} = ½$ and receiving a signal with accuracy p_0, which turns out to be negative. The latter case, an agent with an uninformative prior and a signal, we believe, emphasizes the information content compared to the former, an agent with an informative prior. Importantly, this is the only difference between the two setups. The mathematical requirements to calculate posteriors remain unchanged. We refer to this equivalence as *information structure*, pertaining to the idea that the same content of information can be presented in different ways.

This equivalence holds more generally, for instance, in cases in which an agent has an informative prior and receives additional information from the outset. Consider an agent with an informative prior $P(F) = p_0 > ½$ who receives a signal with accuracy θ_s. Once more, we can express the updated beliefs in the form of a posterior-odds ratio as in equation (1). Now, consider an alternative case where an agent receives two signals with accuracy θ_1 and $\theta_2 = \theta_s$, and has an initial prior $\tilde{p}_0$. Conditional on $s_1 = n$, the posterior-odds ratio will be

$$\frac{P(F|s_2, s_1 = n)}{P(S|s_2, s_1 = n)} = \frac{P(s_2|F)}{P(s_2|S)}\frac{P(s_1 = n|F)}{P(s_1 = n|S)}\frac{P(F)}{P(S)} = \frac{P(s|F)}{P(s|S)}\frac{\theta_1}{1 - \theta_1}\frac{\tilde{p}_0}{1 - \tilde{p}_0}.$$

In the special case in which the accuracy of the first signal is $\theta_1 = p_0$, and the prior is uninformative, $\tilde{p}_0 = ½$, the above reduces to

$$\frac{P(F|s_2, s_1 = n)}{P(S|s_2, s_1 = n)} = \frac{P(s|F)}{P(s|S)}\frac{p_0}{1 - p_0}\frac{½}{½}. \tag{2}$$

that matters.

[10]Having $p_0 > ½$ is without loss of generality. If $p_0 < ½$, let the realized signal be $s = p$ and signal accuracy be $\theta_s = 1 - p_0$.

Note that equation (2) is equal to the posterior-odds ratio under the initial problem, equation (1).[11] Albeit mathematically equivalent, we consider this to be a crucial distinction because, in practice, individuals may or may not process information from their priors and signals in the same way.

Consider, for example, an individual who knows that the prevalence of a disease in the population is only 1% and tests positive on a test with 90% accuracy. A Bayesian agent would recognize the significance of both the prior and the signal, forming a posterior that integrates all available information (resulting in P(have the disease|positive test) = 8.3%). However, in practice, an individual might downplay the relevance of the population distribution after receiving a personalized test, emphasizing the signal's information instead. Our exercise allows us to keep the information content unchanged while shifting where the information comes from. Another benefit of delivering information via two signals, instead of an informative prior and one signal, is that it creates an environment in which we can manipulate the timing of information, with signals arriving simultaneously or sequentially.

To sum up, the two information structures, one with an informative prior and a signal and another with an uninformative prior and two signals, are mathematically equivalent. Importantly, this equivalence holds regardless of whether the signals arrive simultaneously or sequentially. This is critical for our setup because, as we move from one treatment to another, we alter the structure and sequencing of information, yet the core mathematical problem remains unchanged.

Experiment Related Features. The discussion above may give the impression that the particular realization of the additional signal (prior-replacing signal) is important. However, for our purposes, this is not the case, and any realized signal allows us to directly compare settings with informative and uninformative priors. Naturally, there are situations where we may care about the specific direction towards which the prior leans, especially if there is a relationship between this direction and belief updating, e.g., leaning Democrat or Republican might influence how certain information is processed. However, in numerous settings—our study being one of them—the particular direction is an abstraction, and we have no reason to expect belief updating to be influenced by it. As an example, if the prior leans 85% towards F and an agent receives a positive signal with accuracy 80%, we expect the change in beliefs to be comparable with a situation in which the prior leans 85% towards S but the signal realization is negative.

---

[11] $\frac{P(s|F)}{P(s|S)}$ will be equal to $\frac{\theta_s}{1-\theta_s}$ or $\frac{1-\theta_s}{\theta_s}$ depending on the realized value of the signal.

In other words, what we care about is how beliefs are updated and not the particular initial direction. This leaves two possible outcomes of relevance: information is *aligned* (the realized signal is in the direction in which the prior leans; both signals have the same realization) or *misaligned* (the realized signal is opposite to the direction in which the prior leans; signals have different realizations). As a result, ex-ante, having a prior that leans one way or another with probability p_0 and a signal with accuracy θ_s will generate the same distribution of aligned and misaligned information as having an uninformative prior and receiving two signals, one with accuracy p_0 and one with accuracy θ_s.[12]

2.3 Complexity

Motivating Example. To see the general idea behind our complexity measure, we revisit the example briefly outlined in the introduction. Consider two information structures, both with an uninformative prior and two signals, the accuracy of which differs by five percentage points. In the first case, signal accuracies are $\theta_1 = 0.85$ and $\theta_2 = 0.80$, while in the second, they are $\theta_1 = 0.97$ and $\theta_2 = 0.92$. Consider a decision-maker who observes two signals, $s_1 = n$ and $s_2 = p$, and updates posterior beliefs. If both signals had identical accuracy, say both equal to 0.80, or both equal to 0.92, then it seems natural to place equal weight on both signals and form a posterior equal to the original prior of 50%. However, because the first signal is more accurate than the second, individuals may argue they should weigh the first signal slightly more and, thus, compute a posterior that leans more toward Failure than Success.

Although this is true in both cases, the extent to which the first signal is more informative than the second is quite different between them. The Bayesian posterior is 41.3% in the first case, seemingly a plausible value near which peoples' average posteriors may lie. In contrast, the Bayesian posterior is 26.2% in the second case, despite the difference in signal accuracies remaining only five percentage points. This pattern becomes even more pronounced with higher accuracy signals or as the disparity between signal accuracies increases.

Complexity and Nonlinearities. The difficulty illustrated above stems from the non-triviality of incorporating information from signals with different accuracies. The challenge is to know how much more one needs to react to the higher accuracy signal compared to the lower accuracy one. We argue individuals may struggle to fully internalize

---

[12]Specifically, in both cases, the probability that information is aligned is $p_0\theta_s + (1 - p_0)(1 - \theta_s)$ and misaligned with complementary probability.

the extent to which Bayesian updating requires non-linear thinking. While there are several ways to capture this idea, in this section, we offer a parsimonious method relying only on Bayesian posteriors.

Figure 1 depicts the derivative of the posterior with respect to the high-accuracy signal.[13] The marginal increase is unchanged only when facing a single signal ($\theta_2 = \frac{1}{2}$ reduces the low-accuracy signal to an uninformative one). In all other cases, the marginal change in the posterior depends not only on the signal's own accuracy but on the accuracy of the other signal as well, i.e., the Bayesian posterior is no longer linear. It is these nonlinearities in Bayesian updating that we believe contribute to the varying levels of difficulty.

Figure 1: Bayesian Posterior Derivative w.r.t high accuracy signal, θ_1

Notes: The left graph depicts a heatmap (warmer colors represent higher values), and the right graph is 3D plot. In both graphs, we focus on $\theta_1 \geq \theta_2$, since θ_1 denotes the high accuracy signal.

When comparing two signals with different accuracies, the region of interest becomes the region between the juncture where both signals have the same accuracy, all the way to the point characterized by the accuracies of the signals. A Bayesian agent would be able to fully follow these changes and figure out how much more they need to weigh the high-precision signal compared to the lower-precision one. However, an agent struggling to follow such nonlinearities might make errors proportional to these cumulative nonlinearities. To link the complexity of a task with these cumulative nonlinearities, we integrate the absolute value of the second derivative of the Bayesian posterior, starting from the juncture where two signals have equal accuracy up to the point of interest

---

[13]Predictions are unchanged if alternatively, we focused on the low accuracy signal.

$$C(\theta_2, \theta_1) = \int_{\theta_2}^{\theta_1} \left| \frac{d^2 P(\tilde{\theta}_1, \theta_2 | s_1, s_2)}{d\theta_1^2} \right| d\tilde{\theta}_1. \tag{3}$$

Above, $P(\tilde{\theta}_1, \theta_2 | s_1, s_2)$ represents the Bayesian posterior given signal accuracies $\tilde{\theta}_1$, θ_2, and signal realizations s_1 and s_2.[14] Low-complexity environments will be those in which the Bayesian posterior is rather linear, and thus, neglecting nonlinearities will not matter much, resulting in low $C(\theta_2, \theta_1)$ values. In such environments, a somewhat linear approximated understanding of the environment proves effective. High-complexity environments will be those in which the Bayesian posterior is rather non-linear. In these environments, aggregate nonlinearities are large, leading to large $C(\theta_2, \theta_1)$ values. In these environments, relying on a linear approximated understanding of the environment leads to sizable discrepancies.

Testable Implications. For ease of exposition, we express the accuracy of a more accurate signal θ_1 in terms of the less accurate signal θ_2, i.e., $\theta_1 = \theta_2 + \eta$, where η is a positive constant.[15] The above definition of complexity leads to two main testable implications.

1. Complexity increases as the level of signal accuracies increases: $\frac{\partial C(\theta_2, \theta_2 + \eta)}{\partial \theta_2} > 0$.

2. Complexity increases as the gap between signal accuracies increases: $\frac{\partial C(\theta_2, \theta_2 + \eta)}{\partial \eta} > 0$.

As the level and/or the gap between signal accuracies increases, the updating task takes place on a more nonlinear region, which, according to our predictions, may lead participants to make larger mistakes.[16] It is these testable implications that we take to the data. As we describe in detail in Section 3, we utilize six parameterizations, across which we systematically vary both the absolute level of signal accuracies and the gaps between them. The empirical footprint of these testable implications is the wedge between posteriors reported by participants and Bayesian predictions. If our notion of complexity holds water, we expect to see greater disparities in belief-updating problems with higher complexity—higher levels of and gaps between signal accuracies.

---

[14]This is one of many ways to capture the nonlinearities of the environment. Another measure that gives almost identical predictions in this setup is the Gini coefficient (Lorenz curve), which looks at the deviation of a graph of interest (the distribution of wealth) from the 45-degree line (a fully linear function).

[15]Naturally, the value of η is restricted to be $\eta \in (0, 1 - \theta_2)$.

[16]When both signals have the same realizations, there are parameter values that make the effect of the level non-monotonic. However, our study focuses on cases with misaligned information, as declared in the preregistration and described further in Section 3. For misaligned signals, the above predictions hold true for all parameter values.

Alternative Formulations. In the Appendix, we offer two alternative models that yield qualitatively similar predictions to those described here. First, in Section A.1, we explore a modification of the Grether (1980) model, which has been extensively used in empirical work studying belief updating, and make minor restrictions incorporating the aforementioned concepts. This modified model is more inert than the Bayesian model, leading to large differences, especially in regions where large changes in posteriors are expected for small changes in signal accuracies. Second, in Section A.2, we explicitly model an agent's inability to fully follow changes in the second derivative and, thus, only partially react to nonlinearities. The agent performs well when his linear representation is a good approximation of the Bayesian posterior, which happens when both signals have low accuracies or when the gap between signal accuracies is relatively low and makes large mistakes otherwise.

The approach covered in this section, alongside the two alternative models in the Appendix, represents some of possibly many methods aiming at highlighting the idea that nonlinearities involved in Bayesian updating may influence an individual's ability to carry out belief updating properly. We suspect—and hope—that future work will identify and test even more specifications that speak to this phenomenon.

3 Experiment Design

We designed our experiment to manipulate three key aspects: task complexity, information structure, and information sequencing.

To capture the impact of complexity, we vary parameter values. To capture the impact of information structure, we vary whether participants receive information through an informative prior and one signal or through an uninformative prior and two signals. Finally, to capture the impact of the sequencing of information delivery, we vary whether the two signals arrive simultaneously or sequentially. To decompose any possible interaction between the timing of a signal's arrival and its accuracy, we have two sequential information arrival treatments in which we vary the order of signal accuracies: high accuracy followed by low accuracy and vice versa. Below, we first describe the main task participants encountered in each treatment and then provide details about the structure and parameters of each treatment.

Main Task. In each treatment, participants face a standard belief-updating task with a binary state and binary signal(s). At the outset of each round, the state is represented by a project selected from a pool of projects; each project has a p chance of being a Failure

and a $1 - p$ chance of being Successful. Participants know p but do not know whether the project is successful or not. Depending on the treatment, participants receive either one or two conditionally independent signals with known accuracies.[17] In treatments with two signals, participants receive signals with different accuracies. Let θ_1 (θ_2) denote the accuracy of the higher (lower) accuracy signal. We use the strategy method and elicit participants' posterior beliefs for each possible signal realization.[18] Relying on the strategy method ensures we collect a balanced dataset. Each participant participates in only one treatment and plays a total of 20 rounds.[19]

Feedback. At the end of each round, participants are informed about the realized value of the signal(s) (positive or negative) and state (Success or Failure). Realized signal and state values from all previous rounds are stored at the bottom of the screen in an easy-to-read table. We include detailed feedback to mitigate potential memory issues that may disrupt learning and confound results; see screenshots in the Online Appendix.

Treatments. In the **Baseline** treatment, participants have an informative prior p and receive one signal with accuracy θ_2.[20] This treatment mimics the classic experiment of Kahneman and Tversky (1972) and the follow-up literature on base-rate neglect (Benjamin, 2019; Esponda et al., 2023).

In the **Simultaneous** treatment, participants have an uninformative prior and receive two signals simultaneously. The signals are conditionally independent and have accuracies θ_1 and θ_2. With two binary signals, there are four possible combinations of signal realizations. However, since the prior is uninformative, the case in which both signals are positive is the mirror equivalent of that in which both signals are negative. Similarly, the first signal being negative and the second being positive is the mirror of the case in which the first signal is positive and the second is negative. In other words, asking four questions would have been redundant. To keep treatments comparable, we randomly draw a value for the first signal and rely on the strategy method to allow for both a positive and a negative realization of the second signal. For more details, see the Online Appendix.[21]

---

[17]Signal accuracy is the probability that the signal correctly reveals the state: $P(s = p|S) = P(s = n|F) = \theta_s$.

[18]The strategy method is a common practice in many experiments, including beliefs-updating experiments (Gneezy et al., 2013; Esponda et al., 2023). For the comparison between the strategy method and the direct response method, see Brandts and Charness (2011).

[19]Repetition of a task is a standard technique in experiments, which allows participants to adjust to the interface and further arrive at their optimal response. We address learning in the data analysis section.

[20]We vary whether p represents the probability of Success or Failure. This assignment is determined randomly, with equal probability, for each participant at the beginning of the experiment and remains unchanged throughout.

[21]In addition to guaranteeing a balanced dataset, this design ensures that participants encounter all

In the **Sequential High-Low** treatment, participants have an uninformative prior and receive two signals. Unlike the Simultaneous treatment, these signals arrive sequentially. Participants receive the higher accuracy signal first. After observing the first signal realization, participants submit their updated beliefs about the state. Afterward, relying on the strategy method, participants submit their beliefs for two possible realizations of the second signal.[22] Thus, beliefs are elicited twice, once after the first signal realization and once via the strategy method for both possible values of the second signal.

Finally, the **Sequential Low-High** treatment is identical to the Sequential High-Low treatment except that participants receive the lower accuracy signal first.

In summary, for each treatment in each round, we elicit two beliefs: one when the signal aligns with the prior (or the first and second signals align) and another when the signal is misaligned with the prior (or the first and second signals are misaligned).[23]

Parameters. Table 1 summarizes our treatments and parameters. We use two main sets of parameters, denoted below by A and B. Under parametrization A, in treatments in which participants receive two signals, signal accuracies are set to $(\theta_2, \theta_1) = (0.80, 0.85)$, whereas in the Baseline treatment (denoted by $\tilde{A}$), the prior accuracy is set to $p = 0.85$ and signal accuracy to $\theta_2 = 0.80$. This is the classic set of parameters used in the base-rate neglect literature (Kahneman and Tversky, 1972; Benjamin, 2019; Esponda et al., 2023). Under parametrization B, in treatments in which participants receive two signals, signal accuracies are set to $(\theta_2, \theta_1) = (0.85, 0.95)$, whereas in the Baseline treatment (denoted by $\tilde{B}$), the prior accuracy is set to $p = 0.95$ and signal accuracy to $\theta_2 = 0.85$. As described in Section 2.2, within each parameterization, the Bayesian predictions are identical for all treatments. The primary purpose of parametrization B is to assess the validity of the notion of complexity described in Section 2.3, and to act as a test of robustness for the findings from the initial parametrization A.

Four additional parameterizations, C, D, E, and F, allow us to decompose how the level and gap between signal accuracies are linked to complexity. Based on our conceptual framework presented in Section 2.3, we can rank all six sets of parameters in terms of complexity, with parameterization A being the least complex and parameterization B being the most complex one. The remaining sets can be ranked based on the two organizing principles, where C_j denotes the complexity of parameterization j:

---

possible combinations of positive and negative signal realizations throughout the rounds.

[22]This makes the design comparable with the Simultaneous treatment in which the first signal value is drawn, whereas beliefs are elicited for both possible values of the second signal.

[23]We say the signal is aligned with the prior if the prior leans towards Failure (Success) and the realized signal value is negative (positive). In treatments with two signals, signals are aligned (misaligned) if both have the same realized value (one is positive and the other is negative).

Table 1: Sessions, Treatments, and Parameter Values

Session	# Participants	Treatment	Parameter	Prior	Low Accuracy Signal	High Accuracy Signal
1	101	Baseline	$\tilde{A}$	p=0.85		—
2	101	Simultaneous			$\theta_2 = 0.80$	
3	101	Sequential High-Low	A	p=0.50		$\theta_1 = 0.85$
4	100	Sequential Low-High				
5	99	Baseline	$\tilde{B}$	p=0.95		—
6	99	Simultaneous			$\theta_2 = 0.85$	
7	102	Sequential High-Low	B	p=0.50		$\theta_1 = 0.95$
8	100	Sequential Low-High				
9	99		C		$\theta_2 = 0.75$	$\theta_1 = 0.85$
10	100	Simultaneous	D	p=0.50	$\theta_2 = 0.80$	$\theta_1 = 0.90$
11	100		E		$\theta_2 = 0.85$	$\theta_1 = 0.90$
12	100		F		$\theta_2 = 0.90$	$\theta_1 = 0.95$

1. Complexity, as defined in equation (3), increases as the level of signal accuracies increases, holding constant the gap between them:

   - $C_A(0.80, 0.85) < C_E(0.85, 0.90) < C_F(0.90, 0.95)$

   - $C_C(0.75, 0.85) < C_D(0.80, 0.90) < C_B(0.85, 0.95)$

2. Complexity increases as the gap between signal accuracies increases, holding constant the level:

   - $C_A(0.80, 0.85) < C_D(0.80, 0.90)$

   - $C_E(0.85, 0.90) < C_B(0.85, 0.95)$

Subject Pool. We conducted our experiment on the Prolific platform with roughly 100 participants in each of the 12 treatments, for a total of 1202 participants. We recruited participants between the ages of 18 and 70, who were living in the United States, were fluent in English, and had a high approval rating on Prolific. For each treatment, an equal number of men and women were recruited. The main experiment was carried out in October - December 2022, while two additional treatments (D and E) were conducted in July 2023.

Participants' Payments. In all treatments, participants received a $5 payment upon completion. In addition, each participant had a 20% chance to be selected into a bonus group. For the selected participants, one of the experimental rounds was randomly chosen for payment. The answers submitted in the chosen round determined whether the

selected participant received an additional bonus of $20. We used the standard BDM method to incentivize subjects to truthfully state their beliefs.[24] The experiment lasted 20 minutes on average, and participants earned, on average $7.97.

Implementation. The experiment was approved by Caltech (IR22-1237) and Duke University IRB (2023-0033) and preregistered on aspredicted.org.[25] The experimental software was programmed in oTree Chen et al. (2016). Instructions and screenshots of the interface are presented in the Online Appendix.

4 Results

We begin the analysis with parameterization A ($\tilde{A}$). The first results we present in Section 4.1 and Section 4.2 gradually peel off layers that contribute to sub-optimal updating, including the structure and sequencing of information delivery. Having done so, we reach the Simultaneous treatment—an environment free of the effects of structure and sequencing. It is in this environment that we study our notion of complexity. The main analysis of complexity is presented in Section 4.3, where we utilize all six parameterizations, supplemented by data from Enke and Graeber (2023) for an additional assessment. In Section 4.4 we explore the robustness of our findings across parameterizations A ($\tilde{A}$) and B ($\tilde{B}$). In Section 4.5, we use the discrepancy arising from signal sequencing to counter the discrepancy arising from complexity. Finally, in Section 4.6, we quantify the extent to which information structure, sequencing, and complexity affect belief updating across the least and most complex parameterizations A ($\tilde{A}$) and B ($\tilde{B}$), respectively. The individual-level analysis is presented in Section 5.

For the remainder of the paper, with a slight abuse of notation, we will refer to parameters $A(\tilde{A})$ and $B(\tilde{B})$ as simply parametrization A and B, respectively.

---

[24]The BDM is theoretically an incentive-compatible method for eliciting truthful responses regardless of participants' risk attitudes Becker et al. (1964). In addition, to help participants understand this method, we told them that they had no incentive to report beliefs falsely if they wanted to maximize the expected payoff in the experiment. Danz et al. (2021) shows that announcing that truth-telling is optimal is an effective way to elicit true beliefs.

[25]The experiment was conducted in three waves: the initial wave with parametrization A and B in October 2022; complexity treatments, parametrization C and F, in December 2022; additional complexity treatments, parametrization D and E, in June 2023. Each wave was separately preregistered on aspredicted.org; see preregistration 1, 2, and 3.

Approach to data analysis. We focus our main analysis on cases in which participants receive misaligned information.[26] We do so because when participants receive aligned information, the Bayesian posterior probabilities are very close to zero.[27] With predicted values so close to the 0 border, implementation errors participants may have are unlikely to be mean zero. We regard utilizing this data for our main analysis as less than ideal. Thus, as stated in our preregistration, our focus is on elicitations from misaligned signals.[28] Nonetheless, we utilize all elicitations when conducting individual-level analysis in Section 5 and illustrate them in the Appendix.

To simplify the presentation and eliminate redundancies, in our data analysis, we normalize the prior and the high-accuracy signal to be negative. Since we focus on the elicited beliefs from misaligned signals, this normalization implies a positive value for the low-accuracy signal.

For the majority of results, we present round-by-round data. In cases when we pool data across rounds, we cluster errors on the individual level.

4.1   Baseline vs Simultaneous

We begin our analysis by comparing the Baseline and the Simultaneous treatments under parametrization *A*, which exhibits low complexity according to our measure (Section 2.3). In the Simultaneous treatment, both signals are received at the same time, eliminating any effects of information sequencing. Moreover, the effect of information structure is minimized as all available information is conveyed through signals alone, without the presence of an informative prior. Therefore, the Simultaneous treatment under parametrization A is the natural initial comparison to the Baseline treatment as it minimizes the effect of information structure, information sequencing, and complexity.

In Figure 2, we present the round-by-round average posteriors under misaligned information. The two dashed lines represent the Bayesian posterior (41.38) and the posterior for an agent exhibiting perfect base-rate neglect (80). The latter refers to an agent that completely disregards the information from the prior and relies solely on the signal to form a posterior.

---

[26]Recall, in the Baseline treatment, information is aligned (misaligned) if the signal's realization agrees (disagrees) with the direction in which the prior leans. In all other treatments, information is aligned (misaligned) if the realized values of the signals are the same (different).

[27]These probabilities are 0.042 for parametrization *A*, 0.009 for *B*, 0.055 for *C*, 0.027 for *D*, 0.019 for *E*, and 0.006 for parametrization *F*.

[28]An alternative approach would be to treat these elicitations as truncated. However, doing so requires making assumptions about the nature of the truncation and the distribution of implementation errors. These assumptions would naturally not be without loss of generality.

Figure 2: Posteriors in Baseline and Simultaneous treatments, parameterization A

Notes: We report the round-by-round average posteriors under misaligned information, alongside the 95% confidence intervals, clustered at the individual level. The lower horizontal dashed line depicts the Bayesian posterior, while the top dashed line depicts the perfect base-rate neglect posterior.

Throughout all rounds, we see minimal learning in the Baseline treatment, maintaining an average posterior of 63.69. These features are consistent with the existing literature.[29] This characteristic of minimal learning holds for all treatments in our experiment. In the Simultaneous treatment, however, the average belief is 41.65 and not statistically distinguishable from the optimal Bayesian level (see Table 5); the p-value of the difference is 0.786. We find this observation striking, considering that, as emphasized in Section 2, the mathematical problem underlying both treatments is identical. Moreover, it is worth noting that this complete correction occurs right from the outset.

Result 1 (Belief Updating Immediate Correction). *In low-complexity environments, delivering information through two simultaneous signals leads to an estimated mean statistically indistinguishable from the Bayesian posterior.*

We consider this to be an important finding, especially given the comprehensive body of work in both Psychology and Economics that establishes base-rate neglect as one of the most persistent deviations from Bayesian updating, showing minimal responsiveness to higher incentives Gneezy et al. (2023), ample feedback Esponda et al. (2023), or the use of more relatable contexts over mathematical constructs Gigerenzer and Hoffrage (1995).

---

[29]See, for example, Esponda et al. (2023) who have the same baseline as us with the same parameter values but run the experiment in a laboratory with university students as their sample pool. There, the average beliefs in round one are around 64. These levels drop with learning but at a very slow rate.

Certain approaches that have managed to mitigate deviations from optimal updating involve the utilization of aggregate statistics computed over several rounds. However, in real-life situations, individuals may not encounter the same problem repeatedly, and aggregating private information can be challenging. Thus, we consider it noteworthy that appropriate structuring and timing of information alone can effectively aid in, or in this case, completely rectify belief updating.

4.2   Information Structure and Sequencing

Having established the differences between the Baseline and Simultaneous treatments, we turn to examine factors that drive this difference. There are two main features that distinguish the Simultaneous treatment from the Baseline. First, the Baseline treatment features an information structure with an informative prior and one signal, while the Simultaneous treatment has an uninformative prior and delivers an equivalent amount of information via two signals. Second, in the Simultaneous treatment, all information is delivered at the same time, whereas the Baseline presents information sequentially; the prior precedes the signal. Either or both of these features may be responsible for the differential behavior in the two treatments.

We use our sequential information treatments to decompose the role of these features in correcting sub-optimal belief updating. Figure 3 depicts the round-by-round average posterior beliefs for all four treatments in parameterization *A*. Table 5 presents average posteriors when pooling data across all rounds.

The Effect of Information Structure per se.   The informative prior in the Baseline treatment conveys an equivalent amount of information as the first signal in the Sequential HL treatment. The accuracy of the second signal is the same in both treatments, as is the sequencing of information. Thus, the difference between the two comes exclusively from the way information is communicated—the information structure.

We evaluate this comparison in two steps. First, we compare beliefs elicited after the first signal in the Sequential HL treatment with the induced prior in the Baseline treatment. While these are predicted to be identical theoretically, their empirical equivalence needs to be established. Our data directly speaks to this equivalence. Barring implementation errors, we find that participants correctly update their beliefs when faced with an uninformative prior and one signal only (see Figure 7 and the discussion in Section 4.3).

The second step is to compare the final beliefs elicited in the Simultaneous HL treatment with those in the Baseline. Here, as well, we find no difference in elicited beliefs at

Figure 3: Posteriors in all *A* parametrization treatments
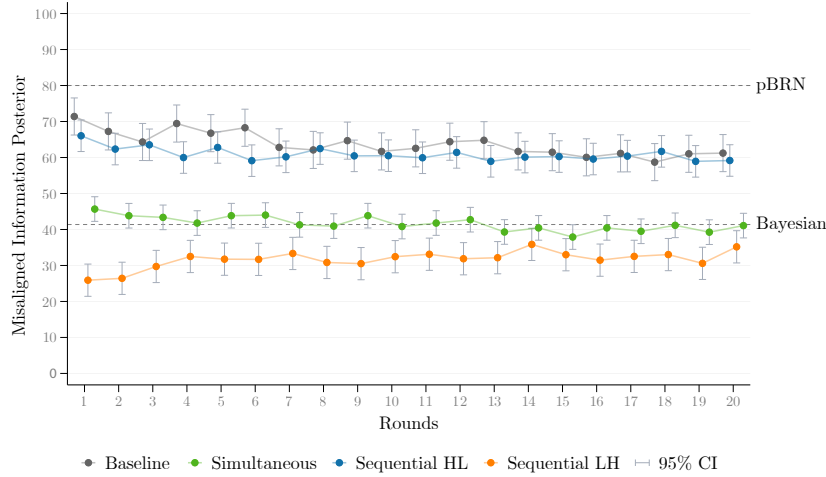
Notes: We report the round-by-round average posteriors under misaligned information, alongside the 95% confidence intervals, clustered at the individual level. The lower horizontal dashed line depicts the Bayesian posterior, while the top dashed line depicts the perfect base-rate neglect posterior.

the aggregate level, see Figure 3 and Table 5, the p-value of the difference is 0.28. This aggregate data, however, hides some individual differences, which we explore in Section 5.[30]

Result 2 (Information Structure Effect). *We document no aggregate effect on participants' reported beliefs when altering the information structure.*

The Effect of Information Sequencing. The Sequential HL and Sequential LH treatments differ only in the order in which the high- and low-accuracy signals are received. These two treatments are necessary to disentangle the effect of the order in which signals are received from the effect of, say, conditioning more heavily than theoretically optimal on the high-accuracy signal regardless of its timing.

Figure 3 shows that the sequencing of information has a substantial impact on belief updating. In all Sequential treatments, participants put a higher-than-optimal weight on the most recent signal they received. In the Sequential HL treatment, this means participants overweight the second (positive) signal by forming posteriors that hover around 60.89 across the twenty rounds, while in the Sequential LH treatment, they overweight the second (negative) signal and arrive at the lower-than-Bayesian posterior of 31.70.[31]

---

[30]To preview these results, in Section 5, we show that a change in the information structure mostly affects beliefs of participants who rely on all information they receive: their beliefs in the Sequential HL treatment are less extreme compared to reported beliefs in the Baseline treatment.

[31]Recall from Section 4 that in our data analysis, we normalize the high-accuracy signal to be negative and the low-accuracy signal to be positive.

On the contrary, when both signals are received at the same time, as happens in the Simultaneous treatment, participants pay attention to both signals while being influenced more by the higher accuracy signal, leading to posteriors that hover around 41.65 across the twenty rounds.

Result 3 (Recency Bias). *We document a sizable recency bias independent of signal accuracy.*

Our results share similarities with previous literature documenting a recency bias. However, we capture this effect in a framework in which we can decompose and quantify its influence apart from other factors. Additionally, by introducing variations in the environment, we assess its robustness across parameterizations; see Section 4.4 for details. It is worth emphasizing that we document a sizable recency bias, despite the fact that the timing between signal arrival and belief elicitation is minimal. Participants receive their first signal, their beliefs are elicited, and immediately after, posteriors incorporating the second signal are elicited. Furthermore, during this second elicitation, the interface reminds participants of the realized value of the first signal. Despite the minimal time elapsed between observing the first signal and being reminded of its realized value, we still observe a significant recency bias. If there was a substantial time gap between signal delivery, one would expect this bias to be even stronger.

4.3 Complexity

The Simultaneous treatment presents an ideal environment for testing the link between environment nonlinearities and complexity. This treatment reduces the impact of information structure and sequencing by utilizing an uninformative prior and two signals delivered simultaneously. In this section, we primarily focus on Simultaneous treatments and investigate the empirical relevance of our complexity measure in organizing the data across all six parameterizations.

We start with contrasting the lowest and the highest complexity parameterizations *A* and *B*, given our measure described in Section 2.3. Recall that under parametrization *A*, participants' choices in the Simultaneous treatment match the Bayesian posterior exactly (see Figure 2). Figure 4 replicates Figure 2 for parameterization *B* and demonstrates that the Simultaneous treatment no longer achieves the optimal Bayesian level. The average of elicited posteriors across the 20 rounds is 33.39, whereas the Bayesian posterior is 22.97. With both information structure and sequencing being controlled, the only difference between the Simultaneous parametrization A and Simultaneous B treatments is the complexity of the problem. Thus, in line with our predictions, an increase in complexity drives a wedge between observed and optimal behavior.

23

Figure 4: Posteriors in Baseline and Simultaneous treatments,
parameterizations B
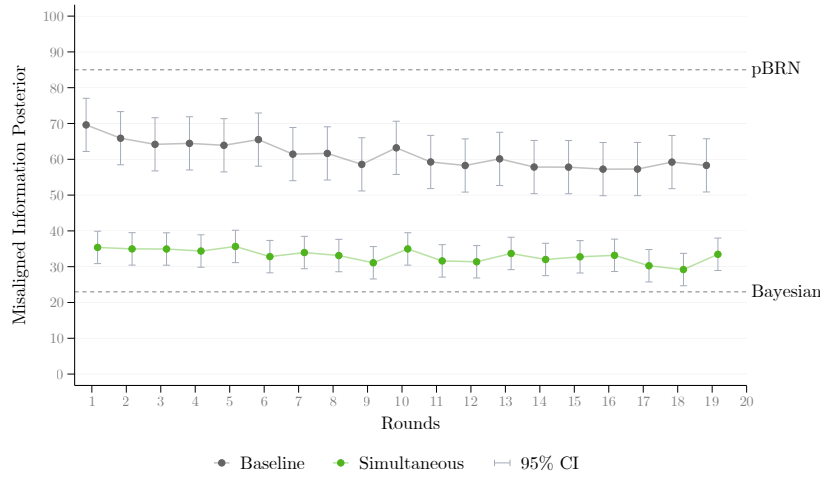
Notes: We report the round-by-round average posteriors under misaligned information, alongside the 95% confidence intervals, clustered at the individual level. The lower horizontal dashed line depicts the Bayesian posterior, while the top dashed line depicts the perfect base-rate neglect posterior.

**Result 4** (Complexity). *Environment complexity drives a wedge between observed and Bayesian posteriors.*

Note that parameterization B features both higher levels of and higher differences between signal accuracies compared to parameterization A. Our next step is to investigate separately how these changes affect belief updating. To do so, we utilize four additional parameterizations C $(75, 85)$, D $(80, 90)$, E $(85, 90)$, and F $(90, 95)$. With these additional treatments, we can separately analyze the impact the level and difference between signal accuracies have.[32]

Empirical Ranking by Observed Complexity. In Figure 5, we display the disparity between the observed and Bayesian posteriors as we increase the level of signal accuracies, keeping the difference fixed. In the left panel, the difference between signal accuracies is 5, whereas in the right panel, the difference is 10. Our data confirms that regardless of the difference between signal accuracies, an increase in the level leads to a larger gap between reported posteriors and Bayesian predictions.

In Figure 6, we display the disparity between the observed and Bayesian posteriors as we increase the difference between signal accuracies, keeping the level fixed. In the

---

[32]Technically speaking, we could have done so with a single additional treatment that differs from parametrization A only in the level or gap between signal accuracies. Nevertheless, through the variation of parameter values, we ensure that the observed effect holds true across different regions.

Figure 5: The Impact of Signal Accuracy Levels on Task Complexity

Accuracy Difference 5

Accuracy Difference 10

Notes: We report the difference between the observed and Bayesian posteriors, averaged across participants in all 20 rounds, alongside the 95% confidence intervals, clustered at the individual level. The horizontal axis depicts the accuracy of the low-accuracy signal and lists the label of each parameterization below.

left panel, the signal accuracy level is fixed at 80, whereas in the right panel, the level is 85. Our data confirms that regardless of the level of signal accuracies, an increase in the difference results in an increase in the observed gap between participants' choices and Bayesian posteriors.

Figure 6: The Impact of Signal Accuracy Difference on Task Complexity

Accuracy Level 80

Accuracy Level 85

Notes: We report the difference between the observed and Bayesian posteriors, averaged across participants in all 20 rounds, alongside the 95% confidence intervals, clustered at the individual level. The horizontal axis depicts the difference in signal accuracies and lists the label of each parameterization below.

We collect these findings in Table 2, where we present a regression of the gap between the reported and the Bayesian posterior on a constant, the level, and the difference between signal accuracies.

Table 2: Accuracy level and Difference impact on observed Gap

| | Gap | |
	All Rounds	Last 5 Rounds
Difference	1.566***	1.555***
	(0.240)	(0.289)
Level	0.464***	0.421***
	(0.124)	(0.152)
N (observations)	11980	2995
K (individuals)	599	599

Individual-level clustered errors in parentheses

$^*\ p < 0.10,\ ^{**}\ p < 0.05,\ ^{***}\ p < 0.01$

The first column of Table 2 utilizes the whole dataset from the six treatments, whereas the second column relies on data from the last five rounds. The regression confirms the observation from the graphs, where we see that both the difference and the level of signal accuracies have a sizable and statistically significant effect.

Result 5 (Complexity: Level and Difference). *The level of and difference between signal accuracies increase the gap between observed and Bayesian posteriors.*

Additional Evidence from Sequential Treatments. Recall that in the sequential treatments, participants' beliefs are elicited twice, once after receiving the first signal and once more via the strategy method conditional on the second signal. Further, recall from our definition of complexity in Section 2.3 that when faced with an uninformative prior and a single signal, the Bayesian posterior is linear. Such tasks are, according to our measure, not difficult. We next test this prediction.

Our data shows that barring implementation noise, participants' reported beliefs after receiving the first signal are largely in line with Bayesian posteriors, see Figure 7. The upper (lower) graphs represent parametrization *A* (*B*), with dashed lines indicating the Bayesian posterior centered at 15, 20, 5, and 15, respectively. The vast majority of participants' choices correspond to these levels. Because noise to the left is bounded by 0, whereas the boundary for noise to the right is much further, the estimated averages tend to be a few percentage points higher than the accuracy of the signal, whereas the median values of the posteriors match the Bayesian posteriors in all four cases. Therefore, when confronted with an uninformative prior and a single informative signal, participants accurately estimate posterior beliefs.[33] This aligns with our notion of complexity and the linear nature of this particular case.

---

[33]Two deviations worth mentioning are a small share of individuals choosing a posterior of 50 and a small share of individuals making the inverse of the correct choice: 85 instead of 15, 80 instead of 20, 95 instead of 5, and 85 instead of 15, respectively. A deeper dive into the data reveals that these are occasional mistakes a few participants make rather than systematic deviations.

It is important to highlight that increasing the cardinality of information sources does not inherently indicate an increase in complexity. Take, for instance, our *A* parametrization, where the Bayesian posterior region is relatively linear, resulting in a low complexity measure. There, the addition of an information source did not affect errors made by participants. In contrast, the *B* parametrization positions the updating task in a very non-linear region, yielding a high complexity measure and, as observed, leads to substantial deviations. In other words, the presence of more than one channel of information does not imply but simply allows for belief-updating tasks to become more difficult.

Figure 7: Sequential Treatments - Posteriors after the First Signal
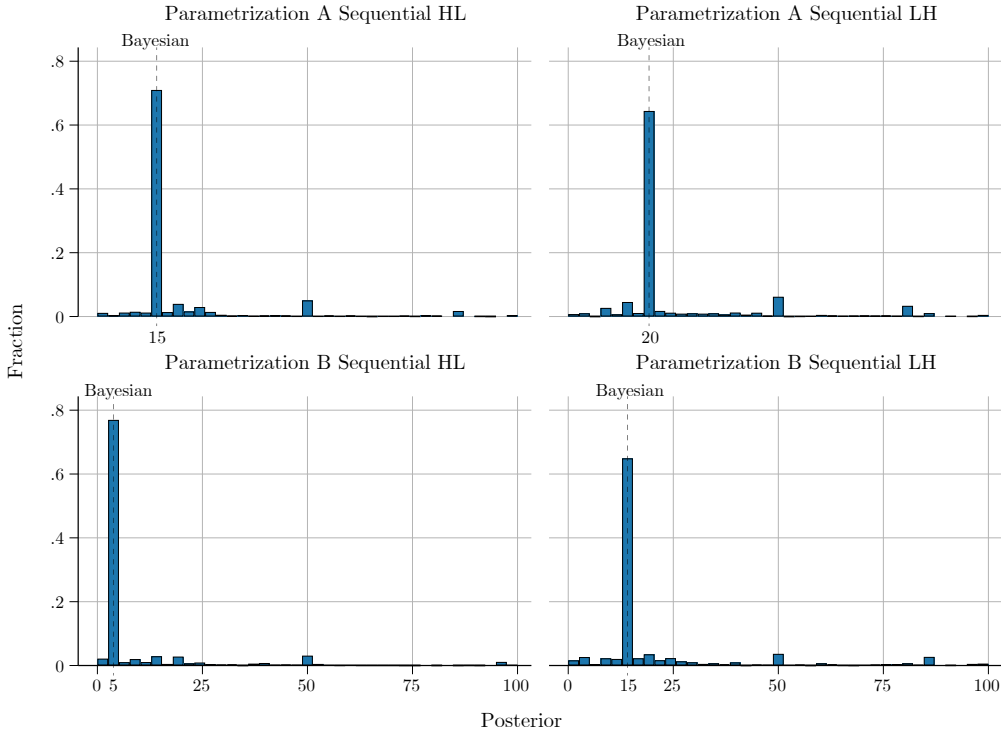
Notes: Above, we report the histogram of participants' posteriors. The vertical axis represents the fraction of choices, whereas the horizontal axis corresponds to the particular posterior.

Importantly, Figure 7 also reveals that our participants are capable of making correct choices far from the middle (50%) point, e.g, in the most extreme case, when the correct choice is 5 in the Sequential B HL treatment, almost 80% of participants make this choice.

Complexity and Cognitive Uncertainty: A Robustness Check. We believe it is useful to compare the implications of our complexity measure with those predicted by the cognitive uncertainty notion studied in a recent paper by Enke and Graeber (2023). Cognitive uncertainty has been linked to a compression effect, according to which individuals tend

to report beliefs closer to the middle value of 50 when they are more uncertain of their answers. Enke and Graeber (2023) provides empirical support for this idea by gathering a new dataset consisting of belief elicitations as well as measures of cognitive uncertainty at the individual level.

We utilize data from Enke and Graeber (2023).[34] Aiming for comparability with our dataset, we examine cases where participants have an informative prior and receive a single signal.[35] Regression analysis, presented in Table 3, demonstrates that the level of and difference between the accuracies of information sources, which represent our measure of complexity, significantly impact the gap between reported and Bayesian posteriors.

This effect remains substantial and statistically significant even after controlling for participants' age, education, Raven scores, college education, gender, and, most importantly, their elicited cognitive uncertainty (second column).[36]

We interpret these results to mean that our notion of complexity and cognitive uncertainty are distinct phenomena that can coexist and jointly influence the accuracy of belief-updating tasks. In other words, our measure is not necessarily linked to perceived uncertainty: a participant may be very confident and very wrong at the same time.[37]

Table 3: Robustness Check

	Gap	
Difference	0.419***	0.444***
	(0.0385)	(0.0391)
Level	0.260***	0.251***
	(0.0550)	(0.0545)
Cogn Unc		0.0529***
		(0.0176)
Other Controls	No	Yes
N	2866	2866

Individual-level clustered errors in parentheses

$^{*}\ p < 0.10,\ ^{**}\ p < 0.05,\ ^{***}\ p < 0.01$

Linear Thinking Estimation. In the Appendix, we analyze two models whose predicted behavior is in line with our notion of complexity. Specifically, in Section A.2, we analyze the case in which an agent can only partially incorporate nonlinearities. After imposing restrictions on the derivatives of the posterior, solving a partial differential equation leads

---

[34]We are grateful to the authors for providing us with their data.

[35]The parameters utilized have prior values of $(50, 70, 90, 95, 99)$ and signal accuracies of $(65, 70, 75, 90)$.

[36]Our findings from Section 4.2 indicate that beliefs observed in the case with an informative prior and one signal should closely resemble those obtained from a treatment involving sequential information arrival with two signals and an uninformative prior. The sequential arrival of information is less than ideal for the study of complexity, as sequencing also impacts elicited beliefs. However, in Section A.5 in the Appendix, utilizing symmetric parameter cases, we perform a back-of-the-envelope calculation and find that, given these parameter values, sequencing can account for only a small part of the increase in the observed gap, with the majority being attributed to the increase in complexity.

[37]Interestingly, running a regression of cognitive uncertainty on the difference and level of signal precisions leads to an insignificant estimated value on the level and a negative and statistically significant estimated value on the difference. That is, participants reported uncertainty does not appear to rise with increased signal precision. On the contrary, reported uncertainty tends to *decrease* as the difference between signal precisions increases.

to the following posterior

$$\tilde{\pi}(S|s_2 = p, s_1 = n) = \alpha \underbrace{\frac{\theta_2(1-\theta_1)}{\theta_2 + \theta_1 - 2\theta_2\theta_1}}_{\text{Bayesian Posterior}} + (1-\alpha)\underbrace{\left(\frac{1}{2} - \theta_1 + \theta_2\right)}_{\text{Fully Linear}}.$$

The α parameter quantifies the agent's ability to incorporate nonlinearities, where a value of $\alpha = 0$ indicates a complete failure to incorporate nonlinearities in belief updating, while $\alpha = 1$ signifies behavior indistinguishable from Bayesian updating. After rearranging the above equation, we can estimate α via a linear regression. The estimation results are presented in Table 4. In the first column, we analyze data pooled from all 20 rounds, while the second column focuses on data from the last five rounds. The estimation reveals that participants are capable of some nonlinear thinking but not to the extent required for Bayesian updating. The graph of the estimated model is presented in Figure 8.

Figure 8: Estimated Model

Table 4: Estimated α

	All Rounds	Last 5 Rounds
$\hat{\alpha}$	0.417***	0.564***
	(0.0555)	(0.0663)
N (observations)	11980	2995
K (individuals)	599	599

Individual-level clustered errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The figure illustrates the Bayesian ($\alpha = 1$) and the fully linear ($\alpha = 0$) models through transparent graphs, along with the estimated model ($\alpha = 0.56$) via the yellow graph.

The estimated value of α is approximately 0.42 (0.56) when considering the entire dataset (the last five rounds).[38] Therefore, the model that best fits the data suggests that participants lie between Bayesian and linear updaters, indicating that participants face challenges in properly incorporating nonlinearities while still exhibiting some degree of nonlinear thinking.

---

[38] By conducting a similar exercise and employing a nonlinear regression on the alternative modified Grether model, analyzed in Section A.2, we obtain an estimated value of approximately 0.46 (0.58) when utilizing the entire dataset (the last five rounds).

Figure 9: Posteriors in all *B* parametrization treatments
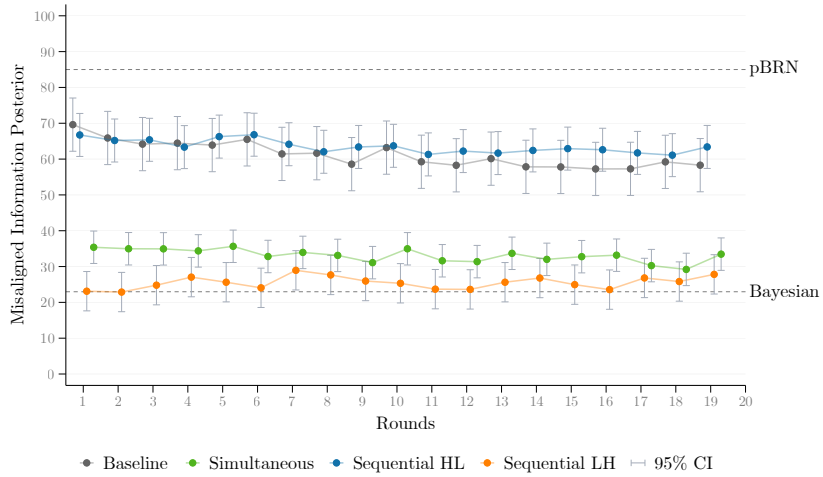
Notes: We report the round-by-round average posteriors under misaligned information, alongside the 95% confidence intervals, clustered at the individual level. The lower horizontal dashed line depicts the Bayesian posterior, while the top dashed line depicts the perfect base-rate neglect posterior.

Result 6 (Complexity: Bayesian vs Linear Updating). *Participants are capable of incorporating nonlinearities only partially. Behavior is best described by a model that roughly lies between Bayesian and linear updating.*

4.4 Robustness Across Parametrizations

We next examine elicited posteriors across all treatments in parametrization B, which, according to our measure, displays the highest degree of complexity. We use this analysis as a robustness check of our findings reported for parametrization A, the least complex treatment. The round-by-round average posteriors are reported in Figure 9, and the average posterior across all rounds can be found in Table 5.

A quick comparison of Figure 9 with Figure 3 reveals that the ranking of estimated means across treatments remains unchanged under both parameterization A and B. The average posterior for the Baseline and the Sequential HL treatments are not statistically distinguishable, followed by a significantly lower posterior for the Simultaneous treatment and an even lower one for the Sequential LH treatment. Under both parametrizations, we find that the structure of information, i.e., the difference between delivering information via an informative prior and one signal versus an uninformative prior and two signals, has no effect on posteriors in aggregate. On the contrary, sequencing plays an important role in determining posterior beliefs. Specifically, we find a large recency bias, consistent with findings reported in Section 4.2.

Result 7 (Robustness). *The ranking between all four treatments is preserved across different parameterizations, as are results regarding information structure and sequencing.*

4.5 Countering Biases

Despite the findings in Section 4.4 showing that the general ranking of the estimated means remains consistent between the two parameterizations, there is an important distinction. In the least complex parametrization A, the simultaneous release of information leads to elicited beliefs not statistically distinguishable from the Bayesian posterior. This is not the case in the most complex parameterization B, where the simultaneous release of information does not completely align posterior beliefs with Bayesian ones, though it does bring them closer.

In the preceding sections, we highlighted two significant biases that impact belief updating. In complex environments, when information is presented simultaneously, by failing to adequately react to the high-accuracy signal, participants tend to under-follow it. Conversely, we have also observed a pronounced recency bias, whereby participants tend to over-follow the most recent signal. With these biases in mind, we explore the possibility of mitigating the complexity bias by releasing information sequentially, with the high-accuracy signal delivered last to utilize the sequencing bias.

An examination of Figure 9 confirms that in high-complexity environments, Sequential LH treatment yields results closer to the Bayesian posterior compared to the Simultaneous treatment. This observation is further supported by posteriors averaged across all rounds presented in Table 5: posteriors in the Sequential LH treatment are not statistically different from the Bayesian level (p-value=0.312).

This analysis suggests that the optimal information release strategy depends on the particular environment. In low-complexity settings where signals are generally weighted correctly, simultaneous information release proves to be optimal. Conversely, in high-complexity environments where there's a risk of incorrectly weighting signals, leveraging the recency bias can help mitigate this issue.

Result 8 (Countering Biases). *The bias arising from sequential information arrival (recency bias) can help mitigate the bias arising from environment complexity (complexity bias).*

4.6 Drivers of Base-Rate Neglect

In Table 5, we present average reported beliefs for all treatments under parametrizations A ($\tilde{A}$) and B ($\tilde{B}$). We provide these estimations for both the entire dataset and for the last five rounds separately. Notably, the observed changes in means are relatively small.

31

Table 5: Estimated Means

	Parameters A ($\tilde{A}$)		Parameters B ($\tilde{B}$)	
	All Rounds	Last 5 Rounds	All Rounds	Last 5 Rounds
Baseline	63.79	60.43	61.93	57.97
	(1.967)	(2.423)	(2.854)	(3.447)
Simultaneous	41.65	40.29	33.39	31.77
	(0.985)	(1.293)	(1.435)	(1.678)
Sequential HL	60.89	59.95	63.70	62.35
	(1.785)	(1.966)	(2.205)	(2.538)
Sequential LH	31.70	32.56	25.04	25.79
	(1.633)	(1.849)	(2.093)	(2.464)

Individual-level clustered errors in parentheses

In both parameterizations, the Baseline treatment exhibits comparable relative levels of base-rate neglect

$$\frac{\mu^A_{Bench} - \mu^A_{Bayes}}{\mu^A_{pBRN} - \mu^A_{Bayes}} = \frac{63.79 - 41.37}{80 - 41.37} \approx 0.58, \qquad \frac{\mu^B_{Bench} - \mu^B_{Bayes}}{\mu^B_{pBRN} - \mu^B_{Bayes}} = \frac{61.93 - 22.97}{85 - 22.97} \approx 0.63.$$

In the calculations above, a score of 0 implies that, on average, participants choose the Bayesian posterior, while a score of 1 implies that, on average, participants choose the perfect base-rate neglect (pBRN) posterior. Recall that pBRN agents disregard the initial information and solely follow the signal.

Next, we delve into a decomposition of the observed level of base-rate neglect, attributing it to information structure, sequencing, and complexity. As discussed before, information structure has no effect on beliefs at the aggregate level, while the complexity of the environment and the sequencing do. We first quantify the extent to which base-rate neglect is influenced by sequencing.

$$Sequencing^A = \frac{\mu^A_{SeqHL} - \mu^A_{Sim}}{\mu^A_{Bench} - \mu^A_{Bayes}} = \frac{60.89 - 41.65}{63.79 - 41.37} \approx 0.86.$$

The remaining portion of base-rate neglect, although not statistically significant, is due to information structure. We visually represent these effects in Figure 10, which essentially illustrates the y-axis of Figure 3 after aggregating data across rounds.

Figure 10: Average posteriors in all *A* parametrization treatments

We also decompose base-rate neglect observed under the second parametrization B as follows

$$Sequencing^B = \frac{\mu^B_{SeqHL} - \mu^B_{Sim}}{\mu^B_{Bench} - \mu^B_{Bayes}} = \frac{63.70 - 33.39}{61.93 - 22.97} \approx 0.78,$$

$$Complexity^B = \frac{\mu^B_{Sim} - \mu^B_{Bayes}}{\mu^B_{Bench} - \mu^B_{Bayes}} = \frac{33.39 - 22.97}{61.93 - 22.97} \approx 0.27.$$

Thus, sequencing continues to be the main driver of base-rate neglect under parametrization B, while complexity also becomes of first-order importance.[39] We depict these effects graphically in Figure 11, which illustrates the y-axis of Figure 9 after aggregating data across rounds.

Figure 11: Average posteriors in all *B* parametrization treatments

Result 9 (Drivers of Base-rate Neglect)**.** *Information sequencing is the main catalyst of base-rate neglect, with environment complexity also playing a crucial role.*

---

[39]Note the calculations add up to 1.05 since information structure, although statistically insignificant, increases base-rate neglect in the second parametrization.

Figure 12: Average Individual Choices
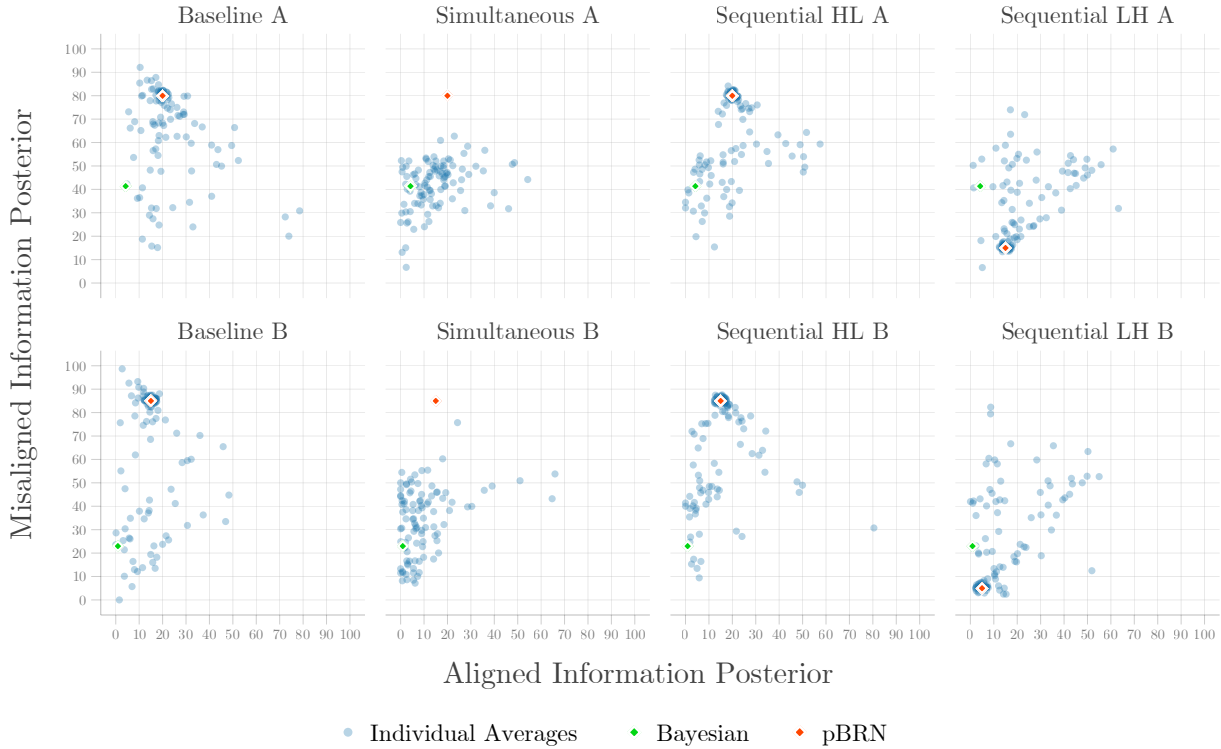
Notes: To help distinguish the large amount of data bundled on the pBRN level, we apply a jitter of 1.5 magnitude. This jittering perturbs the datapoint no further than a distance of 1.5 from the initial value. The top(bottom) row displays data across treatments under parametrization A(B).

5 Individual Level Analysis

5.1 Primary Data Patterns

We now shift our attention to individual-level behavior. For each participant, we calculate their average elicited beliefs across all rounds for both aligned and misaligned information and present these averages in Figure 12.[40] Therefore, each datapoint in the figure represents the average behavior of a single participant.

As can be seen, whenever information is released sequentially, the individual-level average posteriors are heavily bunched around the pBRN level.[41] This bunching phenomenon persists regardless of whether the sequential information delivery stems from an informative prior and a single signal (Baseline treatment) or an uninformative prior

---

[40]In the Online Appendix, we show the counterpart of Figure 12 utilizing only the last five rounds. All main features remain unchanged.

[41]Note that, given our normalization, in Sequential LH, the pBRN level is (15,15) under the first parametrization and (5,5) under the second.

and two signals (Sequential treatments). Only when information is released simultaneously do we observe beliefs that are not heavily concentrated around the pBRN levels. These findings align with Result 9, demonstrating the substantial impact of recency bias.

Result 10 (De-Bundling). *The Simultaneous treatment is the only treatment leading to individual-level beliefs that are not strongly concentrated around the pBRN level.*

5.2 Individual-Level Effect of Information Structure

Recall that we observe no significant differences between average beliefs reported in the Baseline and the Sequential HL treatments in both parameterizations A and B (Section 4.2 and Section 4.4). Although this holds true on average, in this section, we explore whether the information structure plays a role on an individual level.

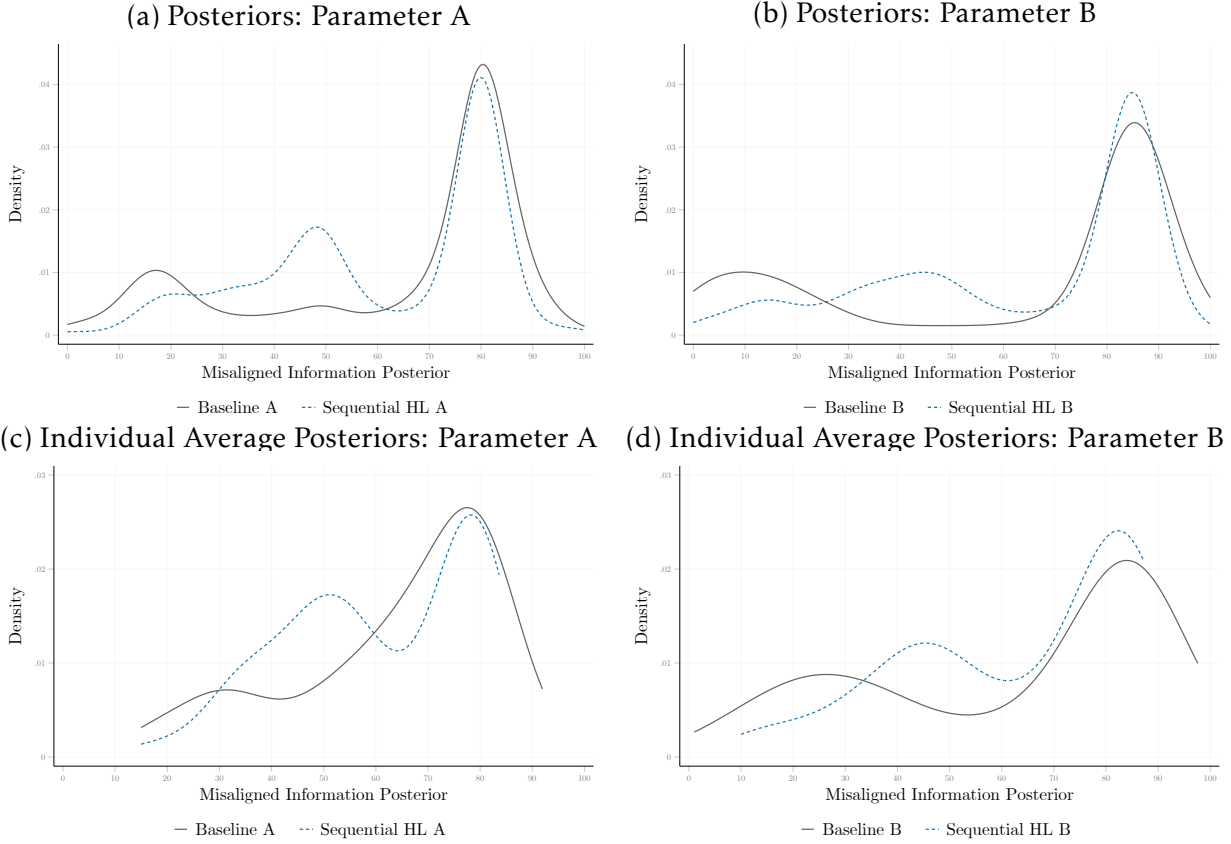In Figure 13a and Figure 13b, we show estimated kernel densities of posteriors from the Baseline and Sequential HL treatments, under parametrization A and B respectively. Within a parametrization, the estimation pools posteriors across participants and rounds.[42] In both parameterizations, despite the similar means, the distributions exhibit notable differences. In the sequential HL treatments, there is a greater concentration towards intermediate values, which are close to the Bayesian level. While the fractions of participants choosing posteriors around the pBRN levels (80 for A and 85 for B) are comparable, in the sequential HL treatment we see fewer values above these levels and fewer values for low posteriors. In both parameterizations, this mass is redistributed from the more extreme values towards the center in such a way that keeps the mean roughly unchanged. However, we run a Kolmogorov-Smirnov test between the Baseline and Sequential HL distributions and reject the null that they are the same ($p < 0.01$) under both parametrizations A and B.

The change between the distributions can be due to small changes in the behavior of many participants, drastic changes in the behavior of some participants, or both. To further explore this, we compute the average posterior for each participant across the 20 rounds and estimate kernel densities based on these average posteriors. Doing so allows us to focus on the variation across participants. In Figure 13c and Figure 13d, we present estimated kernel densities of the average individual-level posteriors in the baseline and sequential HL treatments, under parametrization A and B respectively. As can be seen, a considerable portion of participants, on average, choose levels near the pBRN levels (80 for A and 85 for B). These participants disregard the initial information and solely follow

---

[42]Due to the consistent behavior exhibited by participants, conditioning on any round results in a qualitatively indistinguishable graph.

Figure 13: Distribution of Posteriors and Individual-Level Averaged Posteriors

(a) Posteriors: Parameter A   (b) Posteriors: Parameter B

(c) Individual Average Posteriors: Parameter A   (d) Individual Average Posteriors: Parameter B

the second signal. For the remainder of the distributions, we see that in the Sequential HL treatment, fewer individuals choose extreme values. Our interpretation of these additional estimated kernel densities is that changing the information structure has no effect on individuals who solely follow the second signal, while for others, it steers their choices towards less extreme values—closer to the Bayesian level.

In Section A.4 in the Appendix, we estimate the kernel densities under both parameterizations after removing participants that seem to only follow the second signal.[43] Results reveal that the difference between the estimated kernel densities becomes even starker. This is in line with our interpretation of the differential effect of information structure. In other words, for participants who only focus on the most recent information, the particular information structure does not play a substantial role—they ignore the initial information regardless. On the other hand, for participants who somewhat incorporate both the initial and the more recent information, the specific information structure can influence belief updating.

---

[43]We remove participants from the analysis if their average misaligned and aligned posteriors are within 10 points from the pBRN posterior level.

Result 11 (Effect of Information Structure). *Elicited beliefs of participants who exclusively rely on recent information are unaffected by information structure. For other participants, a change in the information structure results in less extreme reported beliefs.*

5.3   Classifying Types: K-means Clustering

We next proceed by classifying participants into different types. To determine the number of types and the types themselves, we utilize K-means clustering, which, simply put, is a method that partitions *n* observations into *k* clusters/groups. Each observation is associated with the cluster with the nearest mean (centroid). This results in a partitioning of the data space into Voronoi cells. Specifically, K-means clustering minimizes within-cluster variances (squared Euclidean distances). This is one of the most commonly used unsupervised classifiers.[44] By employing this procedure, we bypass the need to determine types subjectively. Instead, we rely on the unsupervised classification procedure to determine both the number of types and their characteristics. To determine the number of clusters, we employ two commonly used approaches, the *elbow method* and the *silhouette score*. Details of these approaches are presented in the Online Appendix. Based on this initial analysis, the suggested number of clusters is three.

Since our aim is to evaluate how the share of different types changes across treatments, we separate the typical K-means clustering into two parts. The first part, clustering, involves determining the centroids for each cluster. We do this by pooling the data across treatments within a parametrization. Having determined the centroids, we then proceed with the second part, classification, which simply associates each observation to the cluster with the nearest mean.[45] To identify the centroids, we use a standard iterative refinement technique. To summarize, we do the following: (i) determine the number of clusters, (ii) determine the centroids, and (iii) classify participants.

We follow the exercise described above for treatments one through three. We exclude the sequential LH treatment for technical reasons.[46] The clustered data is shown in Figure 14, along with the three centroids and the corresponding Voronoi sets they generate.

---

[44]An unsupervised classifier is a machine learning algorithm that automatically identifies patterns and groups data without prior labeled training examples.

[45]Had we not followed the procedure described above, and instead, had we estimated centroids for each treatment, there would be no natural way to compare shares of participants belonging to different groups across treatments since what a group is would differ from treatment to treatment.

[46]In the sequential LH treatment, if we do not normalize the data, as we have done in the main analysis, the Bayesian posterior will have a different position compared to the three other treatments. If we normalize the data, the pBRN posterior will have a different position compared to the three other treatments. This, in turn, hampers our ability to have a natural interpretation of the clusters. Hence, we proceed with the clustering exercise for the first three treatments only.

Figure 14: Parametrization A Clustering

Parametrization A Baseline | Parametrization A Simultaneous | Parametrization A Sequential

Misaligned Information Posterior

Aligned Information Posterior

<u>Notes:</u> Participants are categorized into three separate clusters. Dark gray dots mark the centroids of these clusters, and dashed lines represent the Voronoi cells corresponding to these centroids.

We see the emergence of three distinct clusters, with their centroids exhibiting close proximity to the Bayesian posterior (4.2,41.38), the 50-50 posterior, and the pBRN posterior (20,80). Consequently, we interpret the first group as roughly Bayesian, or closest to Bayesian, the second as potentially exhibiting confusion, and the third as roughly pBRN, or closest to pBRN. We label the second cluster as potentially confused due to the fact that, regardless of the prior and signal value, whether positive or negative, participants consistently opt for values close to 50. We display the fraction of each type across treatments in Figure 15.

Figure 15: Parametrization A Cluster Histogram

Parametrization A Benchmark | Parametrization A Simultaneous | Parametrization A Sequential HL

Fraction

Closest to Bayesian — Confused — Closest to pBRN

As can be seen, a simultaneous release of information under the first parametrization leads to the largest share of participants classified as closest to Bayesian, with a minuscule

share of agents closest to the pBRN level. Importantly, although in the previous section, we saw that the estimated mean under the Baseline and Sequential HL treatments was not statistically different, we see that the composition of the type of participants differs. The Sequential HL treatment is characterized by a higher share of closest-to-Bayesian agents and a lower share of closest-to-pBRN agents. Thus, in line with the evidence presented in Section 5.2, the information structure does seem to have an effect on individual-level behavior.

Result 12 (Information Structure and Type Classification). *Information structure affects participant categorization.*

We next turn our attention to parametrization B. We once again follow the procedure described above and show the clustered data in Figure 16, along with the three centroids and the corresponding Voronoi sets they generate.

Figure 16: Parametrization B Clustering

Notes: Participants are categorized into three separate clusters. Dark gray dots mark the centroids of these clusters, and dashed lines represent the Voronoi cells corresponding to these centroids.

We see the emergence of three distinct clusters, with their centroids exhibiting close proximity to the Bayesian posterior (0.9,22.97), an in-between posterior, and the pBRN posterior (15,85). Consequently, we interpret the first group as roughly Bayesian, or closest to Bayesian, the second as in between the two extremes, and the third as roughly pBRN, or closest to pBRN. We display the fraction of each type across treatments in Figure 15.

Figure 17: Parametrization B Cluster Histogram

Once more, a simultaneous release of information leads to the largest share of participants classified as closest to Bayesian, with a minuscule share of agents classified as closest to the pBRN level. We once again see that the classification of participants differs between the baseline and sequential HL treatments.

Result 13 (Types Across Parameters). *We observe variation in both clusters as well as the distribution of participants among these clusters across different parameters.*

6 Conclusions

Through a series of lab experiments, we examined how task complexity, information structure, and information sequencing influence belief updating. Our findings highlighted that adjusting these factors can alter observed behavior from closely resembling Bayesian reasoning to exhibiting sizable deviations. Our analysis revealed that each of these elements, at the aggregate or individual level, exerts a distinct impact. We quantified these effects and explored strategies to leverage one factor against another, aiming to minimize deviations from Bayesian updating.

Through a range of treatments, as well as supplementary data, we conducted a comprehensive test of a notion of complexity rooted in the nonlinearities of the underlying problem. Taken together, this evidence suggests that nonlinearities are influential in belief updating. We believe peoples' limited ability to fully internalize nonlinearities extends beyond the realm of belief updating and presents an intriguing avenue for future research.

A   Alternative Models of Complexity

A.1   Modified Grether Model

Grether (1980) proposed a generalization of Bayesian updating that accommodated a variety of deviations commonly observed in experiments. This generalization is extensively used in the literature on beliefs Benjamin (2019). In this section, we explore a modification of the Grether (1980) model that yields comparable qualitative predictions to our complexity concept discussed in Section 2.3.

Grether (1980) writes the posterior-odds ratio in the following form

$$\frac{\pi(S|s_2, s_1)}{\pi(F|s_2, s_1)} = \left(\frac{P(s_2|S)}{P(s_2|F)} \frac{P(s_1|S)}{P(s_1|F)}\right)^{\alpha} \left(\frac{P(S)}{P(F)}\right)^{\beta}$$

where $\pi(\cdot)$ captures an agent's possibly biased beliefs. If $\alpha = \beta = 1$, the model reduces to Bayesian updating, whereas $\alpha < 1$ ($\beta < 1$) implies underinference, extracting less information from the signal (prior) than prescribed by Bayes' rule. For comparability purposes, we focus on the uninformative prior case with $\frac{P(S)}{P(F)} = 1$.

We make two modifications to this formulation, given an uninformative prior: (i) agents properly update beliefs when receiving only one signal, and (ii) in the presence of more than one signal, agents under-follow signals more, the more accurate signals are. The first modification relates to the idea that complexity, or high nonlinearities, only appear when there is more than one source of information. The second modification captures the idea that under Bayesian updating, agents are expected to react more and more to small changes in signals' accuracies as these accuracies grow larger. It is in such cases that, we believe, the aforementioned sluggishness is emphasized, and individuals fail to properly Bayesian update.

One way to accommodate these two modifications in the Grether (1980) model is to replace parameter α with a function $\gamma(\theta_1, \theta_2)$ that depends on the signals' accuracies:

$$\gamma(\theta_1, \theta_2) = \alpha \cdot \underbrace{(2min(\theta_1, \theta_2) - 1)}_{\in [0,1], \text{ weight on } \alpha} + 1 \cdot \underbrace{(2 - 2min(\theta_1, \theta_2))}_{\in [0,1], \text{ weight on } 1} \qquad 0 \leq \alpha \leq 1 \qquad (4)$$

Note that the above functional form reflects the two proposed modifications. Indeed, if $min(\theta_1, \theta_2) = 1/2$ then $\gamma(\theta_1, \theta_2) = 1$ in line with the first modification and as $min(\theta_1, \theta_2)$ increases $\gamma(\theta_1, \theta_2)$ decreases, in line with the second modification.

To illustrate, Figure 18 focuses on the case of two misaligned signals and compares the Bayesian posterior with the one induced by the modified Grether model across dif-

ferent θ_1 and θ_2 values.[47] By design, the functions are in agreement when one signal is uninformative (θ_1 or θ_2 is 0.5). The functions are also in agreement if both signals have equal accuracy ($\theta_1 = \theta_2$) because, in such cases, both models assign equal weight to each signal.[48]

Figure 18: Gap Between Bayesian Posterior and Modified Greather

Notes: The figures above show the Bayesian posterior (left), a transparent Bayesian posterior, and the posterior of the modified Grether model (middle), as well as their difference (right). The graphs are plotted for values θ_1 and $\theta_2 \in [0.5, 0.99]$. The value of α is set to 0.

In all other cases, the two functions differ. If signal accuracies are low or if the gap between the signal accuracies is small, the modified Grether model results in posteriors close to Bayesian ones. However, when signals have very different accuracies, or when both signals are very accurate, the modified Grether model produces noticeably more inert updating than the Bayesian one. This is the region in which the sluggishness of our model plays a large role, leading to large differences. These predictions mimic those described in Section 2.3.

A.2   Linear Thinking Model

In this section, we provide an alternative model that captures the difficulty of incorporating nonlinearities in belief-updating tasks. We assume that agents struggle to fully

---

[47]The posterior belief implied by the modified Grether model for two misaligned signals is

$$\pi(S|s_2 = p, s_1 = n) = \frac{\left(\frac{\theta_2}{1-\theta_2}\frac{1-\theta_1}{\theta_1}\right)^{\gamma(\theta_1,\theta_2)}}{\left(\frac{\theta_2}{1-\theta_2}\frac{1-\theta_1}{\theta_1}\right)^{\gamma(\theta_1,\theta_2)} + 1}$$

[48]The two functions are also in agreement if one of the signals is fully informative (θ_1 or θ_2 is 1). To make displaying the graphs clearer, this region is clipped off. Figure 18 only displays posteriors for θ_1 and $\theta_2 \in [0.5, 0.99]$.

appreciate rapid changes required to form correct posterior beliefs in response to small changes in fundamentals and, instead, are only able to adapt to these changes partially. Focusing on the case of misaligned signals, let $\tilde{\pi}(S|s_2 = p, s_1 = n)$ be the posterior of an agent who struggles to incorporate nonlinearities. We define the derivative of this posterior to be a weighted average with weight $\alpha \in [0,1]$ on the derivative of the Bayesian posterior and a constant. A constant is chosen to respect the idea that agents update beliefs correctly when there is only one signal, i.e.,

$$\frac{d\tilde{\pi}(S|s_2 = p, s_1 = n)}{d\theta_1} = \alpha \frac{dP(S|s_2 = p, s_1 = n)}{d\theta_1} + (1 - \alpha) \overbrace{\frac{dP(S|s_2 = p, s_1 = n)}{d\theta_1}}^{\text{a constant}} \Bigg|_{\theta_2 = \frac{1}{2}}$$

The derivative with respect to the other signal accuracy is similar. Solving the above partial differential equation and using the fact that when $\alpha = 1$ the equation reduces to the Bayesian posterior, yields

$$\tilde{\pi}(S|s_2 = p, s_1 = n) = \alpha \underbrace{\frac{\theta_2(1 - \theta_1)}{\theta_2 + \theta_1 - 2\theta_2\theta_1}}_{\text{Bayesian Posterior}} + (1 - \alpha) \underbrace{\left(\frac{1}{2} - \theta_1 + \theta_2\right)}_{\text{Fully Linear}}$$

Note that since the first derivative is a convex combination of the Bayesian first derivative and a constant, the second derivative will always be lower in magnitude than the Bayesian second derivative.

As a final step, we compare the Bayesian posterior $P(S|s_2 = p, s_1 = n)$ with the posterior from the modified Grether model $\pi(S|s_2 = p, s_1 = n)$ discussed in Section A.1, and the linear thinking posterior derived in this section $\tilde{\pi}(S|s_2 = p, s_1 = n)$. We show these posteriors for various levels of θ_1 with $\alpha = 0$ in the first row and $\alpha = 1/2$ in the second row of Figure 19. The modified Grether and the linear thinking model produce similar predictions, both in the direction of departure from Bayesian updating as well as in magnitude. The main difference between the two models emerges when one of the signals becomes fully informative ($\theta_1 = 1$). Near this region, the modified Grether model quickly converges to 0, while the linear thinking model does not. Again, since the fully informative case is not the focus of this study, for relevant regions (away from $\theta_1 = 1$), for any value of α, the two models produce almost identical predictions. Importantly, under both models, when the level of signal accuracies is low, the departure from Bayesian updating is small, even for sizable differences in signal accuracies. In contrast, when the level of signal accuracies is high, even small differences in signal accuracies lead to large gaps from the predicted

Figure 19: Model Prediction Comparison

Notes: Above, from left to right, we utilize $\theta_2 = \{0.80, 0.90, 0.95\}$ and $\alpha = 0$. The blue, green, and orange graphs represent the Bayesian, modified Grether, and Linear Thinking posteriors, respectively.

Notes: Above, from left to right, we utilize $\theta_2 = \{0.80, 0.90, 0.95\}$ and $\alpha = 1/2$.

Bayesian behavior.

A.3 Aligned Information Posteriors

In Figure 20a and Figure 20b, we graph the average round-by-round beliefs of participants when information is aligned. For the Baseline treatment, these are cases when the realized signal is in the direction in which the prior leans. For all other treatments, these are cases in which both signals have the same realized value.

A.4 The Effect of Information Structure on Individual-Level Beliefs

Below, we estimate the kernel densities under both parameterizations after removing participants that seem to behave in a pBRN manner. We remove participants from the analysis if their average aligned and misaligned posteriors are within 10 points from the pBRN posterior level.[49]

---

[49]For parametrization A, this implies we drop participants whose misaligned posteriors are between 70 and 90 and whose aligned posteriors are between 10 and 30. For parametrization B, this implies we drop
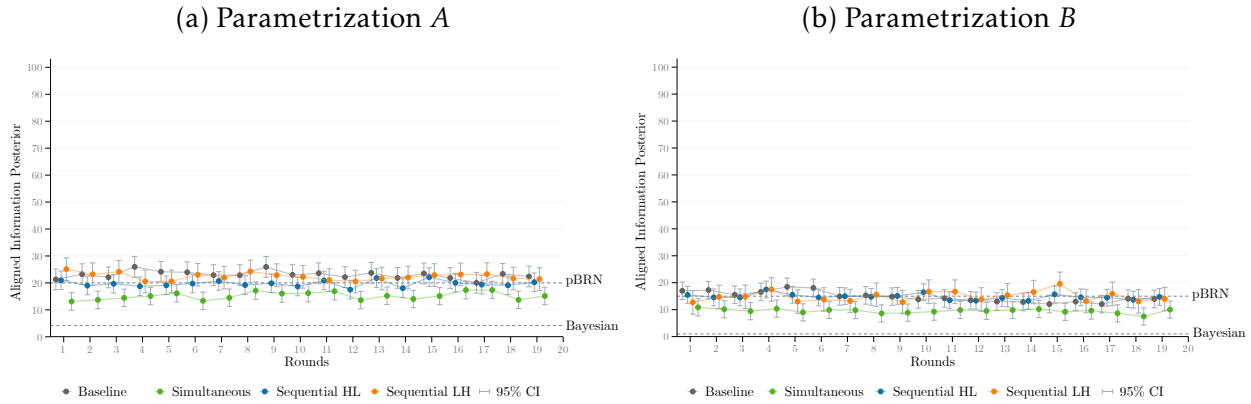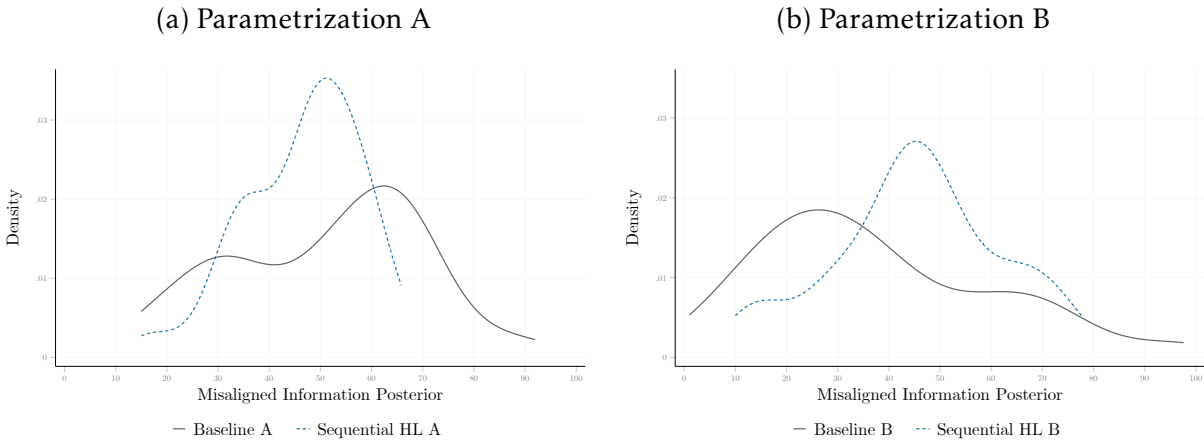
Figure 20: Aligned Information Posteriors

(a) Parametrization *A*

(b) Parametrization *B*

Figure 21: Individual-Level Averaged Posteriors Excluding pBNE

(a) Parametrization A

(b) Parametrization B

Compared to Figure 13a and Figure 13b, the difference between the estimated kernel densities becomes starker. This is in line with our interpretation that information structure seems to have minimal to no effect on participants who show pBRN behavior; however, for other non-pBRN participants, the effect is sizable.

A.5 Complexity and Cognitive Uncertainty: Additional Calculations

The sequential arrival of information is less than ideal for the study of complexity, as sequencing also impacts elicited beliefs. However, utilizing symmetric parameter cases, we can perform a back-of-the-envelope calculation to account for the effect of sequencing. Consider the (70,70) and (90,90) treatments, where the first (second) number corresponds

participants whose misaligned posteriors are between 75 and 95 and whose aligned posteriors are between 5 and 25.

to the accuracy of the prior (signal). Given our measure of complexity, these are both low-complexity treatments. The Bayesian posterior is 50 for both. Data reveals that the mean posteriors are 56 and 53, respectively. Thus, $gap_{low} \in [3, 6]$, which is rather small despite the fact that belief elicitation is hindered by sequencing. Now, consider two examples with relatively high complexity (90,70) and (70,90). Because we, once again, normalize the high accuracy signal/prior to be negative, the Bayesian posterior in both cases is 20.6. The data reveals that in the former case, the posterior is 43, whereas in the latter case, the posterior is 37. This gap (between 43 and 37) is due to sequencing; in the former case, the negative and high-accuracy information arrives first, whereas, in the second case, it arrives later. Naturally, a treatment in which all information arrives simultaneously must lie between these two values. Consequently, the gap, in this case, must be $gap_{high} \in [17.6, 23.6]$. Thus, the gap has risen from $[3, 6]$ to $[17, 23.6]$, implying that the gap has increased by at least 11.23 or at most 20.6 percentage points. In the former (later) case, complexity would account for $^{11.23}/17 = 66\%$ ($^{20.6}/23.6 = 87\%$) of the gap. Hence, while indirectly accounting for the effect of sequencing, we see that the bulk of the increase in the gap can not be related to sequencing and, therefore, must be complexity related.

References

Augenblick, N., Lazarus, E., and Thaler, M. (2023). Overinference from weak signals and underinference from strong signals. *working paper*.

Ba, C., Bohren, A., and Imas, A. (2023). Over- and underreaction to information. *working paper*.

Banovetz, J. and Oprea, R. (2022). Complexity and procedural choice. *manuscript*.

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44:211–233.

Barbey, A. and Sloman, S. (2007). Base-rate respect: From ecological rationality to dual processes. *Behavioral and Brain Sciences*, 30.

Becker, G., DeGroot, M., and Marschak, J. (1964). Measuring utility by a singleresponse sequential method. *Behavioral Science*, 9:226–232.

Benjamin, D., Bodoh-Creed†and, A., and Rabin, M. (2019). Base-rate neglect: Foundations and implications. *working paper*.

Benjamin, D., Rabin, M., and Raymond, C. (2016). A model of nonbelief in the law of large numbers. *Journal of European Economic Association*, 14(2):515–544.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:69–186.

Bernheim, B. D. and Sprenger, C. (2020). On the empirical validity of cumulative prospect theory: experimental evidence of rank-independent probability weighting. *Econometrica*, 88(4):1363–1409.

Brandts, J. and Charness, G. (2011). The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14:375–398.

Camara, M. (2021). Computationally tractable choice. *manuscript*.

Camerer, C. (1987). Do biases in probability judgment matter in markets? experimental evidence. *The American Economic Review*, 77(5):981–997.

Charness, G. and Levine, D. (2005). When optimal choices feel wrong: A laboratory study of bayesian updating, complexity, and affect. *American Economic Review*, 95(4).

Chen, D. L., Schonger, M., and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Cheng, P. and Holyoak, K. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, 17(4):391–416.

Danz, D., Vesterlund, L., and Wilson, A. (2021). Belief elicitation and behavioral incentive compatibility. *American Economic Review*, 112(9):2851–2883.

Edenborough, R. (1975). Order effects and display persistence in probabilistic opinion revision. *Bulletin of the Psychonomic Society*, 5 (1):39–40.

Enke, B. (2020). What you see is all there is. *Quarterly Journal of Economics*, 135(3):1363–1398.

Enke, B. and Graeber, T. (2023). Cognitive uncertainty. *Quarterly Journal of Economics*.

Enke, B., Graeber, T., and Oprea, R. (2023). Complexity and time. *manuscript*.

Enke, B. and Shubatt, C. (2023). Quantifying lottery choice complexity. *working paper*.

Enke, B. and Zimmermann, F. (2019). Correlation neglect in belief formation. *The Review of Economic Studies*, 86(1):313–332.

Esponda, I., Vespa, E., and Yuksel, S. (2023). Mental models and learning: The case of base-rate neglect. *American Economic Review*.

Fan, T., Liang, Y., and Peng, C. (2022). The inference-forecast gap in belief updating. *manuscript*.

Fudenberg, D. and Puri, I. (2022). Evaluating and extending theories of choice under risk. *manuscript*.

Ganguly, A., Kagel, J., and Moser, D. (2000). Do asset market prices reflect traders' judgment biases? *Journal of Risk and Uncertainty*, 20(3):219–245.

Gigerenzer, G. and Hoffrage, U. (1995). How to improve bayesian reasoning without instruction: frequency formats. *Psychological review*, 102 (4).

Gneezy, U., Enke, B., Hall, B., Martin, D., Nelidov, V., Offerman, T., and van de Ven, J. (2023). Cognitive biases: Mistakes or missing stakes? *Review of Economics Studies*.

Gneezy, U., Rockenbach, B., and Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior and Organization*, 93:293–300.

Graeber, T. (2023). Inattentive inference. *Journal of the European Economic Association*, 21(2):560–592.

Grether, D. (1992). Testing bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization*, 17(1).

Grether, D. M. (1978). Recent psychological studies of behavior under uncertainty. *American Economic Review*.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly journal of economics*, 95(3):537–557.

Kagel, J. and Levin, D. (1993). Independent private value auctions: Bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. *The Economic Journal*, 419:868–879.

Kahneman, D. and Tversky, A. (1972). On prediction and judgement. *ORI Research Monograph*, 12(4).

Kahneman, D. and Tversky, A. (1973). On the psychology of prediction. *Psychological review*, 80(4).

Koehler, J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenge. *Behavioral and Brain Sciences*, 19.

Levy, M. and Tasoff, J. (2016). Exponential-growth bias and lifecycle consumption get access arrow. *Journal of the European Economic Association*, 14(3):545–583.

Levy, M. and Tasoff, J. (2017). Exponential growth bias and overconfidence. *Journal of Economic Psychology*, 58:1–14.

Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics*, 116(4):1149–1187.

Oprea, R. (2022). Simplicity equivalents. *manuscript*.

Oprea, R. D. (2020). What makes a rule complex. *American Economic Review*, 110(12):3913–3951.

Pitz, G. and Reinhold, H. (1968). Payoff effects in sequential decision-making. *Journal of Experimental Psychology*, 77 (2):249–257.

Puri, I. (2022). Preference for simplicity. *manuscript*.

Rees-Jones, A. and Taubinsky, D. (2020). Measuring schmeduling. *Review of Economic Studies*, 87:2399–2438.

Stango, V. and Zinman, J. (2009). Exponential growth bias and household finance. *Journal of Finance*, 64(6):2807–2849.

Thaler, R. H. and Benartzi, S. (2004). Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112(51):5164–5187.

Thaler, R. H. and Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness.* Yale University Press: New Haven & London.

Valone, T. J. (2006). Are animals capable of bayesian updating? an empirical review. *Oikos*, 112:252–259.

Wagenaar, W. and Sagaria, S. (1975). Misperception of exponential growth. *Perception and Psychophysics*, 18(6):416–422.