

Belief Decay or Persistence? A Mixed-method Study on Belief Movement Over Time

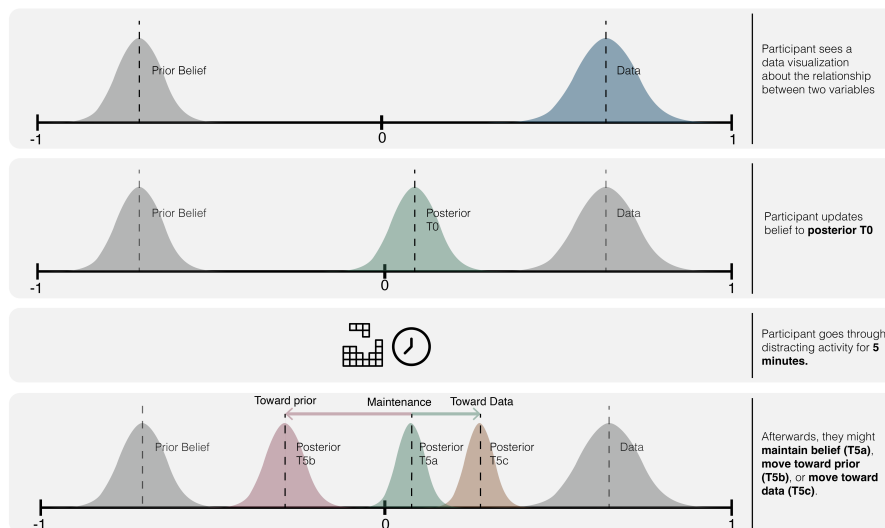
Shrey Gupta¹ , Alireza Karduni² & Emily Wall¹ ¹ Emory University, Atlanta, GA, USA² IDEO, Chicago, IL, USA

Figure 1: This paper compares posterior belief elicitation immediately after seeing new data (T_0) v. after a time delay (e.g., 5 minutes T_5). We conduct quantitative analysis to understand the incidence of belief maintenance (T_5a), belief movement towards the prior (T_5b), and belief movement towards the data (T_5c) after a temporal delay, and conduct qualitative analyses to understand the respective reasons.

Abstract

When individuals encounter new information (data), that information is incorporated with their existing beliefs (prior) to form a new belief (posterior) in a process referred to as belief updating. While most studies on rational belief updating in visual data analysis elicit beliefs immediately after data is shown, we posit that there may be critical movement in an individual's beliefs when elicited immediately after data is shown v. after a temporal delay (e.g., due to forgetfulness or weak incorporation of the data). Our paper investigates the hypothesis that posterior beliefs elicited after a time interval will “decay” back towards the prior beliefs compared to the posterior beliefs elicited immediately after new data is presented. In this study, we recruit 101 participants to complete three tasks where beliefs are elicited immediately after seeing new data and again after a brief distractor task. We conduct (1) a quantitative analysis of the results to understand if there are any systematic differences in beliefs elicited immediately after seeing new data or after a distractor task and (2) a qualitative analysis of participants' reflections on the reasons for their belief update. While we find no statistically significant global trends across the participants beliefs elicited immediately v. after the delay, the qualitative analysis provides rich insight into the reasons for an individual's belief movement across 9 prototypical scenarios, which includes (i) decay of beliefs as a result of either forgetting the information shown or strongly held prior beliefs, (ii) strengthening of confidence in updated beliefs by positively integrating the new data and (iii) maintaining a consistently updated belief over time, among others. These results can guide subsequent experiments to disambiguate when and by what mechanism new data is truly incorporated into one's belief system.

CCS Concepts

• **Human-centered computing** → Empirical studies in visualization; Visualization theory, concepts and paradigms;

1. Introduction

In this data-affluent century, people are experiencing new information including data on topics such as climate, politics, or business analytics via sources like news media outlets, blogs, and social media. This information is often presented using textual or visual representations of data in static and interactive formats and plays an important role in the impactful deliverance of information [MOC21]. Therefore, cogently presented information affects people's perception of the data and helps shape their beliefs: it can convincingly influence an audience due to its impactful messaging [KWMD21] or provoke a biased reaction by using vague, misleading, or unrelated visuals, which is often titled as misinformation [LGS*22, Fox83]. This process of belief updating has been systematically studied whereby, former studies have investigated the patterns that influence the way new information (data) is incorporated with existing beliefs (prior) to form a new belief (posterior) [KKGMD20, KMWD20, WXCW17].

Prior studies have gauged this process by capturing user beliefs through a Bayesian model [WXCW17] and found that some visualizations can improve people's Bayesian reasoning [KKGMD20]. However, these studies share a commonality in that they measure the posterior beliefs immediately after showing new data. Hence, given this trend where only immediate posteriors are examined, we hypothesize that **beliefs elicited after a delay would be systematically different from beliefs elicited immediately after seeing new data**, such that people's beliefs would "decay" back towards their prior over time. This can be due to reasons such as:

1. Beliefs may differ because those elicited immediately after seeing new data may be the result of individuals mimicking the trend shown in the data rather than actually incorporating the data into their beliefs [KOSL12]. Hence, beliefs elicited after a delay may trend back towards the prior. This may be especially true for strongly held beliefs, where sufficient time to mentally process the data may lead to weak or no incorporation of the data into one's belief system.
2. Beliefs may also differ because of forgetfulness [BKIS*19]. While individuals may have incorporated the data into their beliefs upon seeing new information, over time they may forget this information. This may mimic more realistic data analysis practices in that one does not typically see new data and immediately halt their analysis. They continue to interact with data and interact in the world.

We sought to analyze this hypothesis in a controlled experiment. Specifically, **to what extent do people retain updated beliefs over a period of time?** In addition to this hypothesis, we sought to explore the *reasons* for the movement of beliefs, i.e., change in posterior beliefs (updated beliefs) over time [AR21]. In this paper, we present the results of an empirical study with 101 participants who completed three trials where correlation beliefs were elicited immediately after seeing new data (presented using static scatterplots) and again after a brief distractor task (Tetris). While we find that greater than 50% of participants do display some form of belief movement, we find no statistically significant evidence of systematic movement of beliefs toward the prior after the distractor task. However, our qualitative investigation of the reasons for belief movement based on individuals' reflections on their belief up-

dates yielded 10 prototypical reasons for belief movement within the categories of (i) belief movement towards the prior, (ii) belief movement towards the data, and (iii) belief maintenance. This qualitative analysis yielded insights such as "decay" (belief movement towards the prior) in response to forgetfulness or skepticism of new information, and belief movement towards the data in response to rationalizing with oneself to explain the displayed trends. We discuss these trends at length in our qualitative and exploratory analysis. We conclude with a discussion of the potential underlying mechanisms for such belief movement based on our findings.

2. Related Work

To lend context to the forthcoming experiment, we first describe several particularly important bodies of related work including belief updating, Bayesian belief modeling, and belief elicitation.

Belief Updating: The research in belief updating stems from a desire to understand the decision-making process for an individual [And13]. Numerous studies in cognitive science (e.g., [Gri06, GCK*10, GTFG08, GT06, SPG16]) involve understanding human inference and belief updating via experiments where participants are (1) presented some information and (2) asked generalized questions based on that information. For instance, participants learn that horses, cows, and dolphins possess a certain property such as being warm-blooded, and must decide if all mammals possess this property [GCK*10]. Hence, belief updating can be defined as *the change in an individual's belief as an effect of newly presented evidence about a subject* [HE92].

Bayesian Belief Modeling: A Bayesian framework is often used to model the process of belief updating. Studies have compared probabilistic and connectionist models [RM*04] for decision-making and concluded that the Bayesian belief updating framework is intuitively and formally apt for measuring inference and rationality in human thinking. One particularly important property is that Bayesian statistics provides a notion of a normative value against which observations can be compared to assess rationality [BC16]. At its core, a Bayesian framework involves assessing the degree to which an individual will believe a hypothesis h after seeing new data d (their posterior belief), which is determined as a function of the certainty of information in d and their belief in the hypothesis before (prior). Hence, the 3 primary components of belief update in the Bayesian framework are, (i) prior (elicited belief) (ii) data shown (visual representation of information) (iii) posterior (elicited belief). We adopt a similar Bayesian framework in this paper for our study described in the next section.

Bayesian Belief Modeling in Visualization Research: Recent studies in the visualization community have also explored belief updating and decision-making using Bayesian statistics (e.g., [KMWD20, KKGMD20, WXCW17, MDF12, OPH*15, MOC21]). These studies have focused on finding ways to optimally capture participant beliefs using interactive visualizations [KMWD20, KKGMD20], measuring the impact of different visual representations of information [OMHC12, MOC21], capturing "irrational" behaviors, e.g., cognitive biases, in how individuals perceive information [WXCW17, Nov79, PMN*14], studying individuals' exploration patterns (low-level user interactions

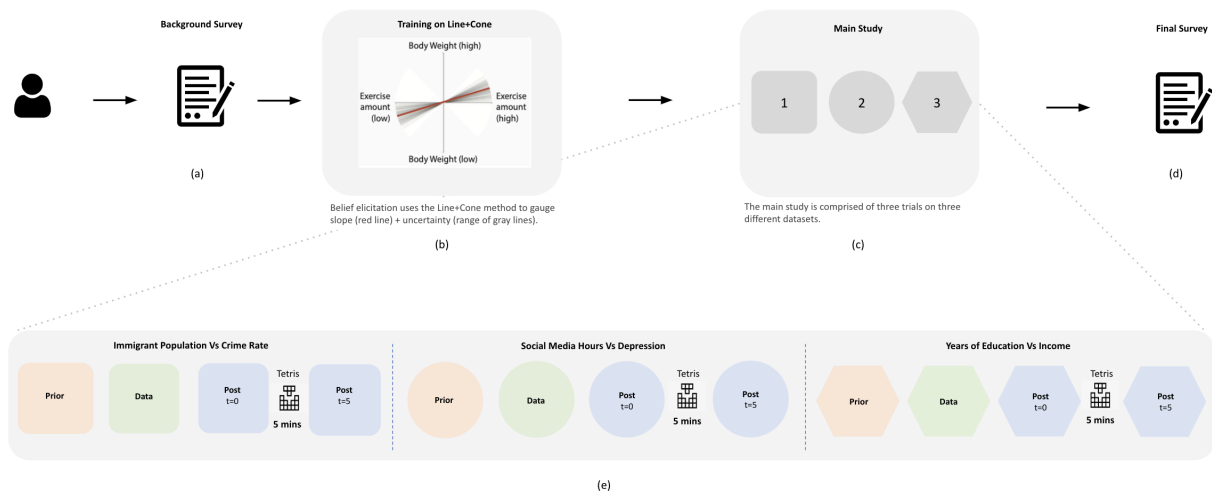


Figure 2: Study Procedure

with visualizations) [MGO20] as well as studying the importance of static [BLE19, GRH13, KCF07] and interactive visualization [KBH18, Shn03, TMK11, KBGH15, MDF12, OPH*15] for Bayesian reasoning [MOC21].

Belief Movement: Augenblick et al. [AR21], discuss the concept of belief movement and how an individual's belief is continuously updated with new revelations (evidence). The authors lead with the hypothesis that a rational belief update should reduce uncertainty with more proof. They formalize this intuition by providing measures for calculating movement and uncertainty reduction when beliefs are constantly updated with new evidence. They successfully show the relationship between cognitive biases, belief movement, and uncertainty reduction. Enke et al. [ESZ20] similarly study how association (utilizing associative memory to restructure prior beliefs) leads to the formation of constantly changing beliefs. Jarrett et al. [JHVDS21] modeled a framework for sequential decision behavior based on participants' changing beliefs (belief movement). Similar to belief updating, we use the related term "belief movement" to focus our analysis on changes in the *direction, magnitude, and uncertainty* of a belief update, which can be attributed to rationalizing and formalizing the information shown over time. We utilize the concept of belief movement as a critical component of the analysis in our experiment.

Belief Elicitation: Prior efforts toward capturing participants' beliefs have primarily been laboratory-based experiments involving participants performing tasks, whereby their actions are associated with their inherent beliefs i.e. captured as their elicited beliefs for the task [ST14, STWB03]. This grounded process helps in validating the meaning behind the captured elicited beliefs. However, such experiments also suffer from *hedging* and *risk-aversion* problems where payoffs are attached to actions and subjects try to coordinate their beliefs with the expected actions. More recently, tasks inspired by visual information inference and reasoning have been used for belief elicitation [Tve05]. This has led visualization re-

searchers to use similar visual interactive belief-capturing methodologies rooted in human-cognition models, which account for how people process and represent information [GRF09, PBG*14, LS10]. Kim et al. [KKGMH20] used an inference-assisted uncertainty framework for Bayesian belief modeling, wherein individuals' beliefs for a point estimate were captured along a slider with an adjustable window to represent uncertainty around those beliefs. Such visualizations also capture uncertainty and allow for finer-grained representation of participants' beliefs by providing room for error [WSM08, HRA15]. Mahajan et al. [MCK*22] study belief-driven-visualizations which are effective in eliciting an individual's beliefs. They provide a framework for designing an impactful belief-driven visualization using the datasets from narrative media.

Line+Cone: One particularly important belief elicitation technique is the 'Line+Cone' methodology introduced by Karduni et al [KMWD20]. 'Line+Cone' is an interactive visual elicitation method that allows an individual to express their belief in a correlation as a line and their uncertainty about the correlation as a cone around the line. The uncertainty margins lie in the interval $[-1, 1]$ and represent a normal distribution around the mean correlation. In our study, we capture individuals' beliefs about correlations using the 'Line+Cone' belief elicitation methodology.

3. Methodology

We conducted three sequential pilot studies to arrive at the final study design shown in Figure 2. In the *first pilot study*, we investigated the comprehension and knowledge of participants for the various datasets. In the *second pilot study*, we investigated the optimal time interval for the study by analyzing a range of time intervals from 1 to 30 minutes using a between-subjects design. In the *third pilot study*, we investigated a within-subjects design where beliefs were elicited immediately after the data was shown and again after a distractor task. We finally adopted the within-subjects design to control for variability at the individual level (see

supplementary and git repository: <https://github.com/shrey-gupta/belief-persistence-evis23> for more information). We ultimately arrived at the pre-registered (https://aspredicted.org/HJW_M4K) study design (Figure 2) as discussed below.

3.1. Procedure

The sequence of the study is depicted in Figure 2. After providing informed consent (Figure 2(a)), participants completed a brief training session on how to use the 'Line+Cone' elicitation method (Figure 2(b)). We ensured a common baseline comprehension of statistics and the 'Line+Cone' elicitation technique by requiring that 4 comprehension questions be answered correctly to advance in the study. Next, the participants completed the main study (Figure 2(c)) where we elicited their beliefs for 3 consecutive trials, and finally, the study ended with participants providing qualitative feedback (Figure 2(d)) in a final questionnaire.

The *main study* phase includes three trials corresponding to three datasets presented in a randomized order. For each trial, we (i) elicited participants' beliefs about the correlation between two variables, (ii) presented new information on the data using a scatterplot, (iii) elicited their updated posterior belief at $t = 0$ min, (iv) intervened using a distractor task (Tetris) for 5 minutes, and finally (v) elicited their updated posterior belief again at $t = 5$ min. The prior and posterior beliefs were elicited using the 'Line+Cone' elicitation method [KMWD20], where the belief about the correlation is represented by the slope of the line, and the uncertainty around the belief is represented by the range of possible values (Figure 2(b)). We also included an attention check after the second trial to make sure participants remained vigilant.

After completing all three trials of the main study, we collect qualitative responses from participants regarding the rationale for the prior and posterior beliefs they expressed. Specifically, for each trial (data domain), participants are asked to provide the reasoning behind their prior and posterior choices in free-form text responses. We also include a Likert scale [JKCP15] asking participants about the likelihood of the supporting data shown being manipulated.

3.2. Data

We used three datasets in this study which were presented to participants in a randomized order (Figure 2(d)). The datasets were bivariate, representing relationships between (1) *Immigrant Population x Crime Rate*, (2) *Social Media Hours x Depression Rate*, and (3) *Years of Education x Income*. These datasets were chosen with the following goals in mind: (i) datasets should represent areas of general social and political knowledge (i.e., not requiring niche knowledge or expertise), and (ii) may likely result in opinionated responses (rather than participants having weak or no opinion about the datasets). Beliefs about these datasets were elicited as correlations between the variables using the 'Line+Cone' methodology [KMWD20], and the supporting data was shown using 50 data points in a scatterplot. We introduced the variable names only and did not provide any further contextual information (e.g., specifying a particular region), which could allow individuals to interpret

the prompts differently. However, we believe the within-subjects design mitigates this risk.

In order to maximize the opportunity to observe a shift in participants' beliefs, we followed a method wherein the data in scatterplots was generated uniquely for each participant to be incongruent i.e. trend in opposite direction to the correlation participants expressed as their prior beliefs. On the contrary, when congruent data is used, most of the participants' belief updates from prior to posterior are arbitrarily small. This is also observed in prior studies [KMWD20, MRK*22] where individuals' prior beliefs aligned closely to the original data for these datasets. Showing "real" data in these settings is unlikely to result in significant belief movement, which would confound the ability to study belief movement. Pennycook et al. [PEM*21] also observe that most participants have the tendency to neglect the "realness" of the dataset, except when explicitly questioned about its validity. The data was generated by taking 50 random samples from a multivariate normal distribution with a correlation value opposite to the prior mean i.e., $corr_u(p) - 1.0$, where $corr_u(p)$ represents the prior correlation for user u . After the conclusion of the study, we informed participants that the data was artificially generated.

3.3. Participants

101 participants were recruited from the Prolific crowd-sourcing platform. Inclusion criteria required the participants to be 18 or more years of age, based in the United States, and fluent in the English language. Participants were ineligible if they participated in any prior pilot studies. Based on power analysis from pilot studies (medium effect size, $power = 0.9$, $\alpha = 0.05$), we aimed for a target sample size of 100 participants. Among the 101 participants, we aimed to balance self-identified *gender* (51 females, 49 males, 1 other). Participants ranged in *age* from 18 to 79 ($\mu = 36.81$, $\sigma = 13.97$) and had varied highest *education* levels (high school 26, undergrad 61, masters 9, doctorate 3, other 1, prefer-not-to-say 1).

4. Results

The goal of this study is to investigate whether there is a systematic decay of beliefs over time via observing belief movement toward prior beliefs. We define *belief decay* as the comparison of the difference between prior beliefs and posterior beliefs at two different time points (T_0 and T_5) such that if the posterior belief at T_5 is closer to the prior than at T_0 , the participant experiences belief movement in the form of belief decay. The following two hypotheses (H1 and H2) along with analyses in sections 4.2, 4.5, and 4.6 were pre-registered. Section 4.1 includes descriptive statistics, and sections 4.3, 4.4, and 4.7 are exploratory (unplanned) analyses.

H1: The belief updated in response to new data will not be retained accurately over time and consequently show signs of decay in subsequent elicitation. This can be demonstrated by posterior belief showing *movement towards the prior* over time.

H2: The amount of belief decay over time will be modulated by the strength of the prior belief (measured using the *uncertainty* range), such that participants with stronger prior belief would show greater belief decay compared to those with more weakly held prior beliefs.

Notation	Description
$B_u(t) = (\mu_t, CI_t)$	The belief of user u at time t , represented as a tuple (<i>slope, uncertainty</i>).
$t \in \{T_p, T_0, T_5\}$	Timestamp for belief elicitation at prior, 1st posterior (0 min), and 2nd posterior (5 min), respectively.
d	The data shown to the user after prior elicitation.
ϵ	Error as a result of measurement i.e. noise. We define a threshold of $\epsilon = 0.05$.
$\delta_u^t(t_2)$	The <i>difference</i> between the beliefs of user u at times t_1 and t_2

Table 1: Notation used in the presentation of study results.




Figure 3: Histogram of prior beliefs and uncertainty for the three datasets. Dashed lines show the average elicited prior belief for each dataset.

Intuitively, we hypothesize that eliciting beliefs immediately after showing a visualization will result in elicited beliefs matching the presented data. In other words, immediately eliciting beliefs will likely represent what users recall from the chart, and after the passage of time, we may observe a decay such that beliefs would revert towards the prior. Therefore we elicit participants' beliefs about three datasets, once immediately after seeing data T_0 , and once at T_5 (after a 5-minute distractor task).

4.1. Descriptive Statistics on Prior Beliefs and Uncertainty

Elicited beliefs about correlations can range between -1 (strong negative correlation) to 1 (strong positive correlation), while uncertainty around beliefs can range between 0 (very certain) to 2 (very uncertain). Trends of participants' prior beliefs on correlation for the three datasets were widely diverging, as were the uncertainties around them (See Figure 3). On average, participants believed that *Years of Education x Income Rate* (Figure 3, top-right)

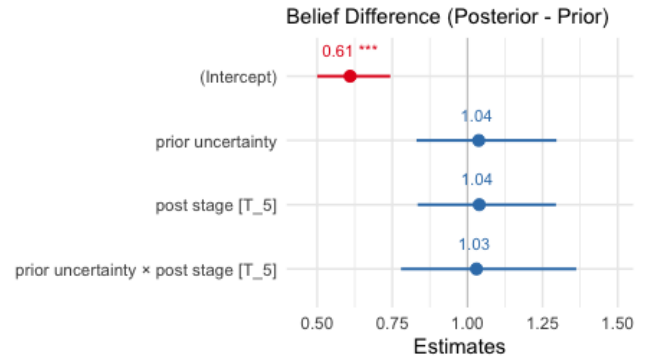


Figure 4: Fixed effects coefficients for belief difference. Error bars indicate 95% confidence intervals. Asterisks indicate statistical significance using p-values: *** 99.9%, ** 99%, * 95%. For post-stage, the reference category is T_0 .

has a strong positive correlation ($\mu = 0.511, \sigma = 0.58$); and they were fairly certain around their beliefs ($\mu = 0.59, \sigma = 0.46$). For the relationship between hours of *Social Media x Depression Rate* (Figure 3, top-center), participants' beliefs, on average, pointed towards a positive correlation while being slightly more spread out ($\mu = 0.41, \sigma = 0.60$) and they were also fairly certain about their beliefs ($\mu = 0.62, \sigma = 0.43$). In contrast, for the *Immigrant Population x Crime Rate* (Figure 3, top-left), participants believed on average that there is little to no correlation ($\mu = 0.13, \sigma = 0.44$). Moreover, Compared to the other two variable sets, participants were slightly more uncertain ($\mu = 0.70, \sigma = 0.47$).

4.2. Belief Decay After Distracting Activity

Given the baseline of prior beliefs described in the previous section, we now seek to answer whether participants' elicited posterior beliefs move towards their prior beliefs after a distracting activity (in this case, tetris).

We used a mixed-effects regression model to account for repeated trials for each participant. Since the difference between participants' elicited posterior and prior beliefs is a bounded value, we used a beta regression with a Logit link [FCN04]. The dependent variable in our model was $(Posterior(t) - Prior, \forall t \in \{T_0, T_5\})$. The fixed effects for our model were the Posterior stage (T_0 or T_5) and its interaction with the Prior Uncertainty (a continuous variable between 0 and 1). We included unique tokens for each participant as random effects. Figure 4 shows the fixed effect coefficients for the mixed-effects model.

It is important to note that in the pre-registration, we described a model with only the posterior stage as a fixed effect with no interaction terms. We also pre-registered a second hypothesis with a model that accounts for users' prior beliefs (and uncertainty) and their interaction with belief change (see pre-registration, section 5). Although this is a slight change from our pre-registration, we found that both models did not show significant effects for the posterior

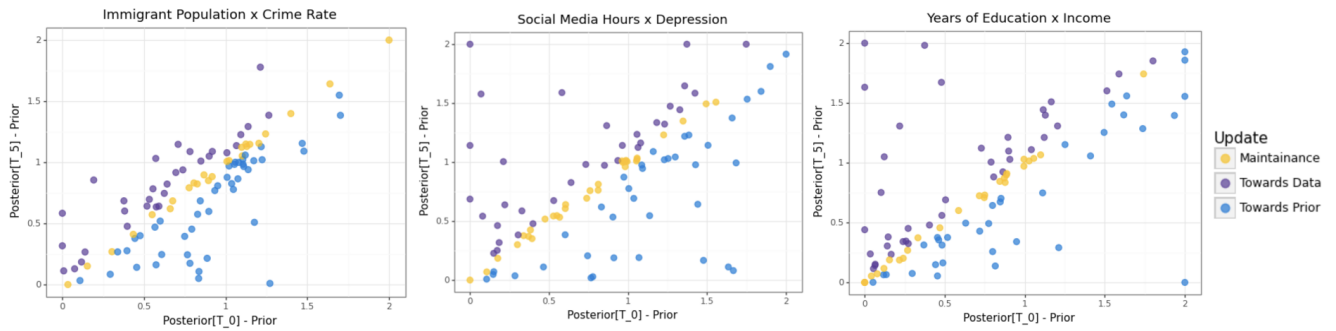


Figure 5: Scatterplot of individuals' difference of posterior T_0 beliefs and prior (x-axis) and difference of posterior T_5 beliefs and prior (y-axis). Yellow points represent those who maintained their beliefs between T_0 and T_5 , purple points represent those whose beliefs moved toward the data, and blue points are those whose beliefs moved toward the prior for the three datasets. Belief maintenance includes an error threshold $\epsilon = \pm 0.05$.

stage as fixed effects. Thus, for brevity, we elected to only report on the second, more complete model.

For $(Posterior(t) - Prior)$, we observe no significant difference between T_5 and T_0 belief elicitation ($\beta = 0.038 [-0.1810, 0.258]$, $z = 0.344$, $p = 0.731$). Thus, **we do not find evidence to support hypothesis H1** (i.e., no evidence to support that participants' posterior beliefs move towards their prior after a disrupting activity). Similarly, we do not observe an effect for uncertainty ($\beta = 0.036 [0.185, 0.259]$, $z = 0.325$, $p = 0.745$). Moreover, we do not observe a significant moderating effect for prior uncertainty ($\beta = 0.029 [-0.249, 0.308]$, $z = 0.208$, $p = 0.835$) on the difference between participants' elicited posterior (at T_0 and T_5) and prior beliefs. Hence, we find no concrete evidence that participants with strong prior beliefs have greater belief decay over time, i.e. **no evidence to support hypothesis H2**.

Finally, we explore whether participants' uncertainty around their beliefs changes after 5 minutes of a distractor task. We used a mixed effects beta regression with uncertainty size as a dependent variable to model this relationship. For fixed effects, we included post-stage (T_0 or T_5); and for random effects, we included participants' unique user IDs. Similar to the analysis for H1, we do not observe any significant effect of time interval on participants' uncertainty around their beliefs ($\beta = -0.135 [-0.3010, 0.031]$, $z = -1.594$, $p = 0.1109$).

4.3. Belief Trend for Individual Participants

The prior analysis begets the question that if we do not observe a systematic belief movement towards prior, does it indicate that participants fully interpreted the visualizations and updated their beliefs accordingly? Hence, for a more fine-grained analysis of *specific participant* behaviors, we compare participants' belief change w.r.t prior at two timestamps (T_0 and T_5) for each dataset as shown in Figure 5. Let $B_u(t)$ represents the belief of user u at time t , and let $\delta_u^t(t_2)$ represent the *difference* between the belief of user u at times t_1 and t_2 . That is, let $\delta_u^t(t_2) = B_u(t_2) - B_u(t_1)$. Therefore for this analysis, we compute the difference between elicited

posterior and prior beliefs, i.e., for each participant we calculate $\delta_u^p(0) = B_u(T_0) - B_u(p)$ (intuitively, the difference between posterior at T_0 and prior T_p) and $\delta_u^p(5) = B_u(T_5) - B_u(p)$ (intuitively, the difference between posterior at T_5 and prior (p)).

Exploring Figure 5, we observe that individuals who fall near the diagonal (yellow data points) experienced very little change in belief between T_0 and T_5 . Given noise of up to $+/- 0.05$, we consider these to be instances of belief maintenance. Individuals whose beliefs moved toward their prior (i.e., experienced "decay") would be those *below* the diagonal (where $\delta_u^p(0) > \delta_u^p(5)$) (blue data points). On the other hand, points *above* the diagonal represent individuals whose beliefs moved toward the data (where $\delta_u^p(5) > \delta_u^p(0)$) (magenta data points). Based on Figure 5, we observe greater dispersion of belief movement for the *Social Media Hours x Depression Rate* (Figure 5, center) dataset and *Years of Education x Income Rate* (Figure 5, right) dataset. Interestingly, for *Immigrant Population x Crime Rate* (Figure 5, left), participants had an overall neutral impression of correlation and, we see relatively less dispersion of belief movement away from the diagonal.

Nonetheless, when we tally the frequency of individuals who experienced belief movement towards prior v. belief movement towards data using the heatmap in Figure 6a, we observe more individuals experienced movement towards prior for datasets *Immigrant Population x Crime Rate* and *Social Media Hours x Depression Rate* (Figure 6a, row[2-3], column[1-2]), although the magnitude (dispersion) is less. Across the three datasets, there appears to be no systematic trend; however, we do observe some shifts indicative of belief movement. We examine these instances more closely throughout the Qualitative (Section 4.6) and the Exploratory (Section 4.7) analyses.

4.4. Alternative Formulations of Belief Maintenance

We also explore two alternative formulations of the concept of belief maintenance. We initially set our $\epsilon = 0.05$ to accommodate minor shifts in posterior beliefs over time. However, from the scatterplot analysis in Figure 5, we observe very close clusters of the three belief movement trends and acknowledge that the actual noise

Belief Movement	Immigrant Population x Crime Rate	Social Media Hours x Depression	Years of Education x Income
Maintenance	25	31	38
Towards Prior	45	37	28
Towards Data	31	33	35

(a) Grouped by Belief Change w.r.t Prior ($\delta_u^p(t)$) [$\epsilon = \pm 0.05$]

Belief Movement	Immigrant Population x Crime Rate	Social Media Hours x Depression	Years of Education x Income
Maintenance	45	45	51
Towards Prior	32	31	26
Towards Data	24	25	24

(b) Grouped by Belief Change w.r.t Prior ($\delta_u^p(t)$) [$\epsilon = \pm 0.1$]

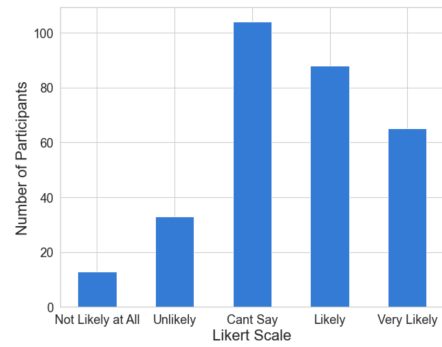
Belief Movement	Immigrant Population x Crime Rate	Social Media Hours x Depression	Years of Education x Income
Maintenance	73	72	75
Towards Prior	20	15	12
Towards Data	8	14	14

(c) When $B_u(T_5)$ falls within CI_{T_0} .**Figure 6:** Heatmap for the number of participants in each category of belief movement trend.

threshold for measuring belief change should be further explored. Hence, we consider a slightly more generous value for noise in this analysis, $\epsilon = 0.1$. By this formulation, we now observe a greater number of instances of individuals who maintained their beliefs over time, as shown in Figure 6b (compared to values shown for $\epsilon = 0.05$ as shown in Figure 6a). We observe that the number of participants who maintained their beliefs increased by approximately 50% on average across the three datasets (Figure 6b, row 1).

Consequently, for our second formulation, we consider a still more generous formulation of belief maintenance. Rather than comparing T_0 and T_5 elicitation with a noise range, we conceptualize belief maintenance as the scenario when the T_5 elicitation falls within the uncertainty range of the T_0 elicitation. That is, if $\mu_{T_5} \in CI_{T_0}$ for $B_u(T_5)$ and $B_u(T_0)$, respectively, for each user u . Intuitively, if the belief still falls within the user's initial plausible uncertainty range, then the belief has not changed. Figure 6c demonstrates this trend, where we observe that almost 70% of responses show belief maintenance (compared to 47% and 31% for $\epsilon = 0.1$ and $\epsilon = 0.05$ formulations, respectively). However, this broader encompassing formulation is not surprising, given that many participants expressed relatively wide ranges of uncertainty reflecting less confidence in the true trend. On the other hand, systematic shifts in belief movement, if such a phenomenon exists, are likely relatively small over short time intervals. We try to capture this in our Qualitative analysis (Section 4.6) next.

Given these alternative formulations of belief maintenance, we see a much wider range of possible numbers of individuals who maintain their beliefs from T_0 to T_5 . We emphasize, however, that our pre-registered analyses are contained in Section 4.2; these findings result from further exploratory analyses (Section 4.7).


**Figure 7:** Bar chart for the Likert responses of participants indicating the likelihood of data was manipulated.

4.5. Did participants trust the visualizations shown?

When faced with any data, participants must decide whether to trust its veracity when incorporating new information into their beliefs. In our study in particular, the data was dynamically generated per participant to represent a trend that was incongruent with their expressed prior belief. Thus after the main experiment, we included a 5-value Likert scale asking whether participants believed that the data shown to them was fabricated for each of the three datasets (hence $101 \times 3 = 303$ responses).

In Figure 7, we observe that when explicitly asked about trust in the provided data, most responses indicated uncertainty (104) about its veracity. From the remaining, the counts skewed toward skepticism, with 88 responses indicating that the data was 'likely' or 65 indicating 'very-likely' manipulated. Few responses indicated that the data was un-tampered, with 33 responses indicating that it was 'unlikely' and 13 responses indicating 'not-likely-at-all' that the data was manipulated. In light of the fact that most participants updated their beliefs to be closer to the provided data over a short interval, these results suggest that a longer time interval could be one factor influencing the degree of belief update and subsequent belief movement.

4.6. Qualitative Analysis

**Figure 8:** Density plot of participants' beliefs at different stages.

Although we do not observe systematic movements of participants' beliefs towards their priors, Figure 8 shows some variability in belief movement both within and across different datasets. We observe that beliefs at posterior (T_5) have a marginal shift towards peaks observed for the prior beliefs compared to the posterior (T_0)

beliefs. Moreover, we also observed that many participants do not fully trust the veracity of the data shown (Figure 7), indicating that these shifts have a probability of persisting over time.

Anticipating scenarios where variability in elicited beliefs can be elaborated by participants' inference of the data shown, we asked them to reflect on their responses, after completing the main study tasks. Specifically, we showed participants their responses and asked "This is what you indicated your belief. Can you describe why you believed this?" with respect to their expressed prior, posterior (T_0), and posterior (T_5) beliefs. Here we analyze the participants' free-form responses containing the rationale behind their beliefs. The goal of this analysis is to understand (1) if/how participants perceived movement in their belief elicitation (or in some cases, what they *wish* they did) and (2) why their beliefs moved between T_0 and T_5 elicitation stages.

Methodology: We established an initial codebook to identify characteristics of participants' rationale including whether they mentioned the direction of their belief (positive, negative, no correlation), source of belief (having read the research, news, anecdotal experience), level of confidence of elicited beliefs, and mention of any difficulties with the interface. After subsequent brainstorming, we included additional codes related to the specific mention of the distractor task, surprise reaction to the data, or skepticism about its veracity. To annotate participant responses, we used an iterative coding strategy. Each author independently coded the same 5 rows of the response data at a time until 90% agreement was reached about the richness of the codes (after approximately 4 iterations). Then, the first author proceeded to code the remaining responses. The complete codebook can be seen in Supplemental Materials.

Given this coding, we grouped responses into one of 9 possible prototypes spanning instances of belief movement towards prior (i.e., the participant perceives and intends that T_5 elicitation is closer to the prior than T_0 elicitation), belief movement towards data (i.e., the participant perceives and intends that T_5 elicitation is closer to the data than T_0 elicitation), and belief maintenance (i.e., the participant perceives and intends that T_5 elicitation is equivalent to T_0 elicitation) (as illustrated in Figure 9). For a total of 101 participants and 3 trials, the number of samples in this analysis includes a total of 303 participant responses. We note that this coding process has inherent noise, which we discuss further in Section 5. Moreover, since this analysis focuses on analyzing the change in belief from T_0 to T_5 , however, we did observe a small number of responses (30/303) that did not significantly change the belief from prior to T_0 , as the participants were either very confident in their belief (strong prior) or, skeptical of the data shown to them. We discuss this further in the Discussion.

Movement towards Prior: We observed four prototypical reasons that participants expressed a shift towards the prior at T_5 v. T_0 . About 83 responses across 46 unique participants showed signs of belief movement towards the prior.

- **Theme 1 [Strong Prior]:** Some participants resorted back to their prior beliefs after the distractor task in the T_5 elicitation as they were able to solidify their beliefs given more time ("I felt even more confident in my previous answer the more I thought

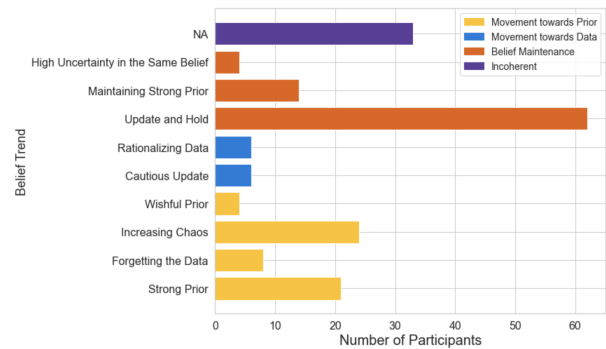


Figure 9: Barchart for participant count in each category of belief movement trend described in the qualitative analysis (Section 4.6).

about it, and felt it might even be more positively correlated than before.") or as they began to question the validity of the data ("Here, I started to believe the graphs could be made up or falsified slightly, so I went with what I believed to be true: immigrants are law abiding citizens."). There were a total of 28 responses across 21 unique participants fitting this theme.

- **Theme 2 [Forgetting the Data]:** Some participants forgot the data shown, i.e. experienced cognitive decay, due to the distractor task (tetris) and as a result had increased uncertainty after the distractor. In these instances, it's likely that the participants did not truly integrate the data into their beliefs at the T_0 posterior elicitation (see Discussion). There were a total of 12 responses across 8 unique participants fitting this theme. The posterior (T_5) responses included rationale such as, "I couldn't remember my response after viewing the data on the scatter plot so I defaulted to my original belief it seems." and "I tried to keep it as close to what the data stated. It wasn't easy to recall the data after playing tetris. So I think I got pretty close."
- **Theme 3 [Increasing Chaos]:** This includes participants who were unsure of the data shown to them and increased their uncertainty over time. There were 39 responses across 24 unique participants fitting this theme. The posterior (T_5) responses expressed uncertainty such as, "Given the data, there was an indication that there is no correlation between education and income, but I am skeptical so I made a high margin of error.", to being confused by the data shown, "That still doesn't make sense to me because why would more people earn less for more education?? so I decided to increase the chaos factor." or "I felt that I was not as sure of the answer so I made the gray area larger."
- **Theme 4 [Wishful Prior]:** In this theme, 4 responses across 4 unique participants claimed that the data shown felt manipulated but did not reflect it in their elicited responses. They expressed regret to not have resorted back to their prior beliefs. The responses were, "I went along with the data, but I shouldn't have. I think it was a mistake to change from my first graph." or, "I kept it relatively the same for the most part because I was under the assumption that the data was true. Despite me having doubts in the results, I kept it based strictly on the data before."

Movement towards Data: We observed two prototypical reasons that participants expressed a shift towards the data at T_5 v. T_0 . In

these cases, some participants updated their beliefs over time by rationalizing the data shown and integrating it with their prior beliefs. There were a total of 16 indicative responses by 11 participants.

- **Theme 5 [Cautious Update]:** Participants fitting this theme concluded to update their beliefs towards data shown but with increased uncertainty. There were a total of 8 responses across 6 unique participants in this theme. The posterior T_5 responses included, “after additional consideration, I thought maybe the headline I read could have been sensationalized and I hadn’t read the study in question, so I gave it less weight despite knowing the data could still be an unrepresentative sample of the population set” or “I think I was able to change my mind as I was not quite sure of the relationship between these events.”
- **Theme 6 [Rationalizing Data]:** Participants fitting this theme expressed more confidence in the data after the distractor task. They positively rationalized their updated beliefs with decreased uncertainty. There were a total of 8 responses across 6 unique participants in this theme. One participant expressed, “I believe that the data concluded that the more hours spent on social media, the less likely one is to have depression, and I was sure of it after the game.”

Belief Maintenance: We observed three prototypical reasons that participants expressed belief maintenance at posterior elicitations T_0 and T_5 . There were a total of 148 responses across 72 unique participants in this category.

- **Theme 7 [Update and Hold]:** In this theme, participants who updated their beliefs immediately after the data shown (T_0) held those same updated beliefs after the distractor task (T_5). There were a total of 118 responses by 62 participants fitting this theme. In these cases, participants often reiterated the same rationale for the T_0 and T_5 elicitations, such as, “I updated my belief based on the scatterplot I was shown.”
- **Theme 8 [Maintaining Strong Prior]:** In this theme, participants discarded the data shown and resorted to their prior beliefs at posterior T_0 and maintained that same updated belief after the distractor at T_5 . There were a total of 25 responses by 14 participants fitting this theme. For instance, one participant said, “I decided to stick with my initial belief because I had a feeling that I was correct, but I wasn’t 100% sure.”
- **Theme 9 [High Uncertainty in the Same Belief]:** In this theme, participants were unsure of the data shown to them but responded by maintaining the same belief but with high uncertainty. There were a total of 5 responses by 4 participants fitting this theme. The responses included, “According to the data there is a strong negative correlation but I do not believe this so I left a wide range of error.” and “It might still be true that higher income can come from more education, but with the job market getting oversaturated by overly educated employees, there probably is a lot more variation now.”

There were an additional 56 responses by 33 participants which couldn’t be successfully annotated as they were semantically incoherent. The responses were vague such as posterior responses like “Oh, I think I misunderstood it?”, or “unsure of my answers”, and did not provide any insights for successful annotations. For these responses, we had difficulty categorizing the direction of their correlation and how the time interval affected the belief update.

Movement	3-trials	2-trials	1-trial
Towards Prior	6	9	30
Towards Data	4	5	22
Maintenance	5	7	13

Figure 10: Heatmap for the Internal Consistency Analysis (Quantitative analysis for all datasets: Grouped by $\delta_u^p(t)$ [$\epsilon = \pm 0.05$])

4.7. Exploratory Analysis

In the previous sections, we reported on pre-registered quantitative and qualitative analyses and ultimately derived a set of perceived reasons for belief movement. In this section, we press further on these findings in additional exploratory analyses. Specifically, we explore the extent to which individuals exhibited similar belief movement trends across the three experimental trials. If participants exhibit movement in a consistent direction across all three trials, this may suggest a potential generalizable tendency (e.g., beliefs for some individuals will decay back towards their prior). Whereas, if the belief movement is inconsistent across the experimental trials, there may exist alternative underlying mechanisms we can explore. Hence, we analyze the belief movement for each participant across the datasets and observe the frequency of participants that show consistent belief movement (towards prior, data, or maintenance) for all 3, exactly 2, and exactly 1 dataset (Figure 10).

Based on our quantitative analysis for all 3 trials, we observe belief movement towards prior for 6 participants, belief movement towards data for 4 participants, and belief maintenance for 5 participants as shown in Figure 10 (left column). Based on our initial hypothesis pertaining to belief decay (movement towards prior) over time, we further examine the **6 participants who displayed belief movement towards the prior for all 3 trials**. We believe this could be indicative of a systematic tendency, e.g., due to weak incorporation of data and hence examine these cases next to better understand potential underlying mechanisms and discrepancies.

One of the 6 participants, who had a strong prior belief (positively correlated) for the dataset, *Years of Education x Income*, eventually became skeptical of the data and reverted back to their prior (“I went back to my original thoughts. I guess even after seeing the supporting data, I still thought my guess was closest.”) For the other two datasets, the participant indicated a weaker prior, changed their mind (“The plot point chart changed my view on the correlation”), and then perceived that they stuck with it, trying to mimic the data (“I kept it the same as the plot points”); however there was significant change in the T_0 and T_5 elicitation for the *Immigrant Population x Crime Rate* dataset ($\delta = 0.38$ back toward the prior) and a value just outside the acceptable range of $\epsilon = 0.05$ for the *Social Media Usage x Depression Rate* dataset ($\delta = 0.07$). For other participants, there was much more substantial movement back towards the prior ($\delta \geq 0.19$ for 15/18 of quantitative responses with an average shift back toward the prior of 0.47). One participant perceived that this was an error, e.g., “should have remained the same”, although it’s unclear whether they perceived that the error was theirs or the systems. Others indicated that they “had more time to think about it” or “went with [their] gut”

Similarly, we assessed the **4 individuals who displayed belief movement towards the data for all 3 trials**, to observe systematic tendencies of recalling and mimicking the data after a distractor task. For two participants, all of the belief movement towards the data was just outside the noise range $\epsilon = 0.05$, with an average belief movement across the 3 trials of 0.08. These participants perceived it as maintaining their belief, indicating “*After Tetris, I kept it relatively unchanged*” and “*I kept my graph mostly the same as the second graph but charted more uncertainty*”. The other two participants had more substantial movement toward the data, on average 0.40 and 0.46, respectively. One said “*I remembered it was a large wave pattern. I evidently made it larger.*”

Quantitative analysis revealed only **5 individuals maintained their beliefs over time for all 3 trials**. Participants indicated things like “*Tetris had no effect on my decision and closely relates to the data presented.*” However, we also note some oddities such as one participant who maintained beliefs across all 3 trials indicated it as being erroneous “*Apparently I was not thinking about the data chart. This answer does not match my original response.*”. Unfortunately due to the very small numbers, this analysis does not shed any light on potential systematic mechanisms behind belief movement. Future studies would be required to assess underlying mechanisms. However, we observe discrepancies between individuals’ actual and perceived belief movement which we discuss further in the following Discussion section.

5. Discussion

Decay or Otherwise? Our experiment was motivated by the hypothesis that individuals’ beliefs decay towards their prior over time. While we found no statistical support for a global trend, we nonetheless observed many individual instances of participants experiencing belief decay. However, this trend could also be explained by an alternative mechanism: individuals may have simply mimicked the data shown or shallowly incorporated it into the working memory such that it doesn’t impact their truly-held beliefs. In such cases, the posterior beliefs would be more accurately termed as “*elicited beliefs*” or “*reported beliefs*” as they may not represent what participants actually believe. However, based on the qualitative analysis, we find it unlikely that participants simply mimicked the trends they observed in the data since participants rationalized the data shown, updated their beliefs accordingly, and in most cases maintained that updated belief after the 5-minute delay. Nonetheless, future studies are required to understand the true underlying mechanism.

Ambiguity in Characterizing Belief Movement In this paper, we provide multiple ways of characterizing belief movement in exploratory analysis. In the quantitative analysis, we analyzed multiple noise thresholds ϵ as well as characterizations of belief movement that rely on the uncertainty intervals participants specified (Section 4.4), each of which led to varied interpretations of belief movement with 31%-70% of participants exhibiting belief maintenance. The qualitative analysis relied on coding characteristics of free-form text expressed by participants in response to each elicitation (T_p , T_0 , and T_5). The accounting of participants experiencing belief movement based on qualitative analysis (Figure 9) may also

be noisy. For instance, there may be noise in the qualitative codes as individuals’ responses vary with some expressing a higher level of detail. Further, there may also be a discrepancy between individuals’ *actual* belief movement and their *perceived* or *elicited* belief movement, depending on individuals’ self-awareness, the strength of their prior beliefs, their interpretation of the task, etc. We emphasize that none of our formulations of belief movement gave statistically significant evidence to support belief decay. Nonetheless, we find that we lack a concrete definition of what constitutes belief maintenance. Future studies could explore to what extent different quantitative and qualitative formulations reflect an individual’s truly-held beliefs.

Limitations and Future Work While our findings yield no statistically significant results to support the presence of belief decay over time, a plausible explanation could be the shorter time interval (5 minutes) we utilized. Although this was based on the findings of our pilot studies, we believe larger durations (days/weeks/months) should also be tested to identify the persistence of belief updates for variable time frames. Similarly, various visualizations for both belief elicitation and providing information should be explored in future studies. Since the qualitative analysis is limited by the crowdsourcing apparatus, we believe in-lab studies may be able to achieve richer participant responses and robust insights. Furthermore, real data rather than simulated (belief-incongruent) data and richer contextual information (e.g., such as narrowing down geographic locality of trends in education x income) can aid in validating the movements observed in our study by reflecting more realistic scenarios. Furthermore, based on the belief movement trends observed in this study, we posit that exploration of time-dependent models that account for decay (e.g., due to forgetting) may be worthwhile. Hence, this study opens the door to an abundance of future work attempting to understand the underlying mechanisms leading to belief movement and in turn belief decay.

6. Conclusion

In this paper, we sought to analyze the belief movement for individuals in response to new information, when posterior beliefs are elicited immediately after seeing the new data and after a brief distractor task. We hypothesized that individuals would experience a *decay*, where posterior beliefs elicited after a distractor task would systematically move back towards the prior compared to a posterior belief elicited immediately after seeing the data. We reported the results of an empirical study with 101 participants. We found no statistically significant evidence to support our hypothesis. However, our qualitative analysis revealed 9 prototypical reasons for instances of belief movement, including 4 instances where beliefs moved towards the prior, 2 instances where beliefs moved towards the data, and 3 instances where beliefs were maintained. This work fills a gap in prior work by beginning to scrutinize assumptions made (or left unsaid) about the nature of belief movement in belief update studies. In summary, while we find no evidence to support a systematic trend of belief movement towards the prior after a distraction task, we are able to shed light on the reasons for belief movement when it occurs.

References

- [And13] ANDERSON J. R.: *The adaptive character of thought*. Psychology Press, 2013. 2
- [AR21] AUGENBLICK N., RABIN M.: Belief movement, uncertainty reduction, and rational updating. *The Quarterly Journal of Economics* 136, 2 (2021), 933–985. 2, 3
- [BC16] BOLSTAD W. M., CURRAN J. M.: *Introduction to Bayesian statistics*. John Wiley & Sons, 2016. 2
- [BKIS*19] BEIERLE C., KERN-ISBERNER G., SAUERWALD K., BOCK T., RAGNI M.: Towards a general framework for kinds of forgetting in common-sense belief management. *KI-Künstliche Intelligenz* 33, 1 (2019), 57–68. 2
- [BLE19] BÖCHERER-LINDER K., EICHLER A.: How to improve performance in bayesian inference tasks: a comparison of five visualizations. *Frontiers in Psychology* 10 (2019), 267. 3
- [ESZ20] ENKE B., SCHWERTER F., ZIMMERMANN F.: *Associative memory and belief formation*. Tech. rep., National Bureau of Economic Research, 2020. 3
- [FCN04] FERRARI S., CRIBARI-NETO F.: Beta regression for modelling rates and proportions. *Journal of applied statistics* 31, 7 (2004), 799–815. 5
- [Fox83] FOX C. J.: Information and misinformation: An investigation of the notions of information, misinformation, informing, and misinforming. 2
- [GCK*10] GRIFFITHS T. L., CHATER N., KEMP C., PERFOR S. A., TENENBAUM J. B.: Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in cognitive sciences* 14, 8 (2010), 357–364. 2
- [GRF09] GREEN T. M., RIBARSKY W., FISHER B.: Building and applying a human cognition model for visual analytics. *Information visualization* 8, 1 (2009), 1–13. 3
- [GRH13] GARCIA-RETAMERO R., HOFFRAGE U.: Visual representation of statistical information improves diagnostic inferences in doctors and their patients. *Social Science & Medicine* 83 (2013), 27–33. 3
- [Gri06] GRIFFITHS T. L.: Statistics and the bayesian mind. *Significance* 3, 3 (2006), 130–133. 2
- [GT06] GRIFFITHS T. L., TENENBAUM J. B.: Optimal predictions in everyday cognition. *Psychological science* 17, 9 (2006), 767–773. 2
- [GTFG08] GOODMAN N. D., TENENBAUM J. B., FELDMAN J., GRIFFITHS T. L.: A rational analysis of rule-based concept learning. *Cognitive science* 32, 1 (2008), 108–154. 2
- [HE92] HOGARTH R. M., EINHORN H. J.: Order effects in belief updating: The belief-adjustment model. *Cognitive psychology* 24, 1 (1992), 1–55. 2
- [HRA15] HULLMAN J., RESNICK P., ADAR E.: Hypothetical outcome plots outperform error bars and violin plots for inferences about reliability of variable ordering. *PLoS one* 10, 11 (2015), e0142444. 3
- [JHVDS21] JARRETT D., HÜYÜK A., VAN DER SCHAAR M.: Inverse decision modeling: Learning interpretable representations of behavior. In *International Conference on Machine Learning* (2021), PMLR, pp. 4755–4771. 3
- [JKCP15] JOSHI A., KALE S., CHANDEL S., PAL D. K.: Likert scale: Explored and explained. *British journal of applied science & technology* 7, 4 (2015), 396. 4
- [KBGH15] KHAN A., BRESLAV S., GLUECK M., HORNBAEK K.: Benefits of visualization in the mammography problem. *International Journal of Human-Computer Studies* 83 (2015), 94–113. 3
- [KBH18] KHAN A., BRESLAV S., HORNBAEK K.: Interactive instruction in bayesian inference. *Human-Computer Interaction* 33, 3 (2018), 207–233. 3
- [KCF07] KELLEN V., CHAN S., FANG X.: Facilitating conditional probability problems with visuals. In *International Conference on Human-Computer Interaction* (2007), Springer, pp. 63–71. 3
- [KKGMH20] KIM Y.-S., KAYONGO P., GRUNDE-MCLAUGHLIN M., HULLMAN J.: Bayesian-assisted inference from visualized data. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 989–999. 2, 3
- [KMWD20] KARDUNI A., MARKANT D., WESSLEN R., DOU W.: A bayesian cognition approach for belief updating of correlation judgement through uncertainty visualizations. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 978–988. 2, 3, 4
- [KOSL12] KNOX W. B., OTTO A. R., STONE P., LOVE B. C.: The nature of belief-directed exploratory choice in human decision-making. *Frontiers in psychology* 2 (2012), 398. 2
- [KWMD21] KARDUNI A., WESSLEN R., MARKANT D., DOU W.: Images, emotions, and credibility: Effect of emotional facial images on perceptions of news content bias and source credibility in social media. *arXiv preprint arXiv:2102.13167* (2021). 2
- [LGS*22] LO L. Y.-H., GUPTA A., SHIGYO K., WU A., BERTINI E., QU H.: Misinformed by visualization: What do we learn from misinformative visualizations? In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 515–525. 2
- [LS10] LIU Z., STASKO J.: Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 999–1008. 3
- [MCK*22] MAHAJAN S., CHEN B., KARDUNI A., KIM Y.-S., WALL E.: Vibe: A design space for visual belief elicitation in data journalism. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 477–488. 3
- [MDF12] MICALLEF L., DRAGICEVIC P., FEKETE J.-D.: Assessing the effect of visualizations on bayesian reasoning through crowdsourcing. *IEEE transactions on visualization and computer graphics* 18, 12 (2012), 2536–2545. 2, 3
- [MGO20] MONADJEMI S., GARNETT R., OTTLEY A.: Competing models: Inferring exploration patterns and information relevance via bayesian model selection. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 412–421. 3
- [MOC21] MOSCA A., OTTLEY A., CHANG R.: Does interaction improve bayesian reasoning with visualization? In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (2021), pp. 1–14. 2, 3
- [MRK*22] MARKANT D. B., ROGHA M., KARDUNI A., WESSLEN R., DOU W.: Can data visualizations change minds? identifying mechanisms of elaborative thinking and persuasion. In *2022 IEEE Workshop on Visualization for Social Good (VIS4Good)* (2022), IEEE, pp. 1–5. 4
- [Nov79] NOV S.: Prospect theory: An analysis of decision under risk daniel kahneman; amos tversky. *Econometrica* 47, 2 (1979), 263–292. 2
- [OMHC12] OTTLEY A., METEVIER B., HAN P., CHANG R.: Visually communicating bayesian statistics to laypersons. In *Technical Report*. Citeseer, 2012. 2
- [OPH*15] OTTLEY A., PECK E. M., HARRISON L. T., AFERGAN D., ZIEMKIEWICZ C., TAYLOR H. A., HAN P. K., CHANG R.: Improving bayesian reasoning: The effects of phrasing, visualization, and spatial ability. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 529–538. 2, 3
- [PBG*14] PATTERSON R. E., BLAHA L. M., GRINSTEIN G. G., LIGGETT K. K., KAVENEY D. E., SHELDON K. C., HAVIG P. R., MOORE J. A.: A human cognition framework for information visualization. *Computers & Graphics* 42 (2014), 42–58. 3
- [PEM*21] PENNYCOOK G., EPSTEIN Z., MOSLEH M., ARECHAR A. A., ECKLES D., RAND D. G.: Shifting attention to accuracy can reduce misinformation online. *Nature* 592, 7855 (2021), 590–595. 4

- [PMN*14] PANDEY A. V., MANIVANNAN A., NOV O., SATTERTHWAITE M., BERTINI E.: The persuasive power of data visualization. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2211–2220. 2
- [RM*04] ROGERS T. T., MCCLELLAND J. L., ET AL.: *Semantic cognition: A parallel distributed processing approach*. MIT press, 2004. 2
- [Shn03] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *The craft of information visualization*. Elsevier, 2003, pp. 364–371. 3
- [SPG16] SUCHOW J. W., PACER M. D., GRIFFITHS T. L.: Design from zeroth principles. In *CogSci* (2016). 2
- [ST14] SCHOTTER A., TREVINO I.: Belief elicitation in the laboratory. *Annu. Rev. Econ.* 6, 1 (2014), 103–128. 3
- [STWB03] STEYVERS M., TENENBAUM J. B., WAGENMAKERS E.-J., BLUM B.: Inferring causal networks from observations and interventions. *Cognitive science* 27, 3 (2003), 453–489. 3
- [TMK11] TSAI J., MILLER S., KIRLIK A.: Interactive visualizations to improve bayesian reasoning. In *Proceedings of the human factors and ergonomics society annual meeting* (2011), vol. 55, SAGE Publications Sage CA: Los Angeles, CA, pp. 385–389. 3
- [Tve05] TVERSKY B.: Visuospatial reasoning. *The Cambridge handbook of thinking and reasoning* (2005), 209–240. 3
- [WSM08] WU Y., SHIH W. J., MOORE D. F.: Elicitation of a beta prior for bayesian inference in clinical trials. *Biometrical Journal: Journal of Mathematical Methods in Biosciences* 50, 2 (2008), 212–223. 3
- [WXCW17] WU Y., XU L., CHANG R., WU E.: Towards a bayesian model of data visualization cognition. In *IEEE Visualization Workshop on Dealing with Cognitive Biases in Visualisations (DECISIVe)* (2017). 2