Imagining the Future: Memory, Simulation and Beliefs

Pedro Bordalo, Giovanni Burro, Katherine Coffman, Nicola Gennaioli and Andrei Shleifer[1]

November 2023

Abstract

How do people form beliefs about novel risks, with which they have little or no experience? Motivated by survey data on beliefs about Covid we collected in 2020, we build a model based on the psychology of selective memory. When a person thinks about an event, different experiences compete for retrieval, and retrieved experiences are used to simulate the event based on how similar they are to it. The model predicts that different experiences interfere with each other in recall and that non domain-specific experiences can bias beliefs based on their similarity to the assessed event. We test these predictions using data from our Covid survey and from a primed-recall experiment about cyberattack risk. In line with our theory of similarity-based retrieval and simulation, experiences and their measured similarity to the cued event help account for experience effects, priming effects, and the interaction of the two in shaping beliefs.

Keywords: Similarity, selective recall, disagreement

[1] Bordalo: pedro.bordalo@sbs.ox.ac.uk, University of Oxford. Burro: giovanni.burro@unibocconi.it, Bocconi University. Coffman: kcoffman@hbs.edu, Harvard Business School. Gennaioli: nicola.gennaioli@unibocconi.it, Bocconi University. Shleifer: shleifer@fas.harvard.edu, Harvard University. This paper replaces the NBER Working paper by Bordalo et al. (2020) which presented the results of our first survey on beliefs about covid without explaining the puzzles, or presenting and testing a model. We are grateful to Sam Gershman for directing us to psychology research on simulation from memory, and to Ben Enke, John Conlon, Edgard DeWitte, Thomas Graeber, Spencer Kwon, Ulrike Malmendier, Dev Patel, Kunal Sangani, Josh Schwartzstein, Jesse Shapiro, and five anonymous referees for helpful comments. Gennaioli thanks the European Research Council (GA 101097578) for financial support.

Introduction

People regularly face novel shocks that change the world in significant and persistent ways, such as global warming, the advent of AI, the fall of the Berlin Wall, or the Covid pandemic. The response to such shocks, at the individual and collective levels, requires an estimation of the risks they entail. The standard approach to such estimation is Bayesian learning, which involves updating using statistical priors and likelihoods. But in entirely novel situations, where do likelihoods and priors come from? An alternative approach is to use personal experiences, as opposed to statistical data (Schacter, Addis, and Buckner 2007; Malmendier and Nagel 2011). But for novel risks, there may be few, if any, closely related personal experiences to draw on to form beliefs. How do people form beliefs in such cases? And does this process shed light on belief formation more generally?

In a 2020 survey of US respondents we found that beliefs about the lethality of Covid, a novel risk at the time, depended on a range of personal experiences in unrelated domains. Motivated by this fact, we build a model of beliefs based on the psychology of selective memory. When a person thinks about an event, different experiences compete for retrieval, including non-domain specific (NDS) ones, and some are neglected, including domain specific (DS) ones. Retrieved experiences are then used to simulate the event based on their similarity to it. The model predicts a similarity-based hierarchy of experience effects, and interference of NDS experiences with the use of DS information. We test these predictions by measuring a wide range of respondents' experiences and their perceived similarity to the target event, both in the Covid surveys and in a pre-registered primed-recall experiment on cyberattack risk. Consistent with our predictions, NDS experiences have strong explanatory power for beliefs, accounting for experience effects, priming effects, and their interaction. Selective memory puts testable structure on otherwise atheoretical influences on beliefs.

Section 2 presents motivating evidence from our Covid survey. It reveals mean overestimation of Covid's lethality (infection fatality rate) but also large disagreement. Three aspects are puzzling. First, people who overestimate an entirely unrelated rare event -- the share of Americans who have red hair -- were more pessimistic about Covid. This fact points to the role of cognitive factors, such as the ease of imagining rare events, as opposed to information, motivations, or preferences, in shaping beliefs. Second, pessimism for oneself and others was stronger if a person had recently experienced her own or a family member's non-Covid hospitalization. This fact points to memory: recalling sick family members may help imagining Covid deaths. Third, there was a striking age gradient: older people were much *less* pessimistic about Covid's lethality than younger people, when estimating lethality both for themselves (contrary to reality) and for others. This may also be due to memory: compared to the young, the elderly may recall many adversities dissimilar to Covid that they lived and survived, making it harder for them to imagine death from Covid. The last two facts point to a role of NDS experiences but go in different directions: NDS experiences boost estimates in the first case but dampen them in the second. Are these conflicting effects consistent with selective memory? If so, are there additional predictions that can be tested in the data?

Section 3 presents a model addressing these findings. When thinking about a risk, people recall either statistics about it heard in the media, or their own experiences. Recall of experiences is driven by three well-known forces: similarity, frequency, and interference (Kahana 2012, Bordalo et al. 2023). Recollections are then used to simulate the risk. Simulation is a form of reasoning by analogy well documented in psychology and neuroscience (Dougherty et al. 1997; Brown et al. 2000, Schacter, Addis, and Bruckner 2008, Biderman, Bakkour, and Shohamy 2020). It recruits experiences, both domain and non-domain-specific, to imagine the target event and hence the future (Schacter et al 2012). Imagination increases in the similarity of retrieved experiences to the target event (Kahneman and Tversky 1981). Similarity thus plays a dual role: it fosters recall *and* simulation. Ours is the first paper to bring similarity-based memory simulation into economics.

Simulation entails a novel, fundamental trade-off in the role of experiences: any experience creates material for simulation, which boosts probability estimates, but also interferes with retrieval of other experiences, which dampens estimates if the latter provide better simulation material. This trade-off between simulation and interference yields two testable predictions.

The first is a similarity-based hierarchy of "experience effects": experiences very similar to the event boost its believed probability more than less similar experiences, while experiences that are even less similar may dampen assessments. Exposure to Covid deaths should boost Covid pessimism more than do the less similar non-Covid health adversities. In fact, non-health adversities, being least similar, may even reduce the perception of Covid risks by interfering with recall of the health risks. This prediction can be tested by measuring, for each respondent, a range of DS and NDS experiences and their perceived similarity to the target event, and accounting for them in predicting beliefs.

Second, the effect of a given experience varies across people due to interference from other experiences in their databases, including NDS ones. For instance, the reaction of beliefs to the severity of the local Covid pandemic is dampened by exposure to, and hence recall of, a non-Covid health adversity, and vice versa. Interference implies that even DS experiences may not come to mind. This prediction can be tested by showing how the effect of a given experience lived by a respondent depends on the frequency and similarity of other experiences she has lived, including NDS ones.

We test these predictions using our Covid surveys in Section 4, and an experiment on cyberattack risk in Section 5. The cyberattack survey provides a sharp test of the model by measuring a broad range of experiences and their similarity to the target event, and by exogenously priming recall of DS and NDS experiences. The results reveal the central role of selective recall and simulation. In both surveys more similar experiences boost estimates more than less similar ones. Contrary to standard theories of belief formation, NDS experiences shape beliefs, both by increasing estimates through simulation, and by interfering with DS experiences. In the Covid survey, experiencing non-Covid hospitalization boosts pessimism about others' risk of dying from Covid. Yet people who (like the elderly) have survived many non-health adversities underestimate risks.[2]

[2] We unify an average tendency to overestimate unlikely risks with strong disagreement among people. Models of overestimation of unlikely events, such as Prospect Theory (Kahneman and Tversky 1979) or noise (Enke and Graeber 2023, Khaw et al. 2020) do not explain why a group of people, such as the elderly in our Covid survey, systematically underestimate an unlikely risk, while other people such as the young systematically overestimate the same risk.

The cyberattack experiment provides a more precise test of these predictions. In our model, priming can only affect beliefs if the primed experience is not otherwise retrieved spontaneously. Thus, a non-zero priming effect already shows that memory is selective. The priming experiment showcases the memory mechanism through two new predictions. First, priming an experience should boost estimates more if it is more similar to the event, even if the experience is NDS. Second, priming an experience should interfere with recall of other experiences, dampening their effect on beliefs. Our results confirm these predictions. We put a theoretical structure of priming effects, which have attracted considerable scepticism due to their instability. Similarity helps characterize the interaction of priming and experience effects, and sheds light on such instability (Cohen and Marechal 2016).

Our results mark a significant departure from Bayesian or noisy Bayesian models (e.g. Sims 2003, Woodford 2003) or Case-Based learning (Schank 1986, Gilboa and Schmeidler 1995), in which beliefs are shaped by DS information that, when present, is not interfered with by NDS information. Our evidence instead shows that beliefs are shaped endogenously by what is *recalled* and how it is *used*, and in particular that DS experiences may fail to be retrieved. From the theoretical standpoint, we build on our prior work on memory and probability judgments (Bordalo et al. 2023), but introduce the mechanism of similarity-based simulation, which is key for understanding how experiences are used, particularly NDS ones. Empirically, this innovation turns out to be crucial.

A vast body of social science research has documented experience effects on beliefs and decisions (e.g., Weinstein 1989). In economics, these have been linked to insurance demand (Kuhnreuther 1978), IPO investing (Kaustia and Knüpfer 2008), demand for redistribution (Alesina and Fuchs-Schündeln 2007), and stock market participation or inflation expectations (Malmendier and Nagel 2011, 2016). This evidence sometimes stresses domain-specificity, to the point that bond market experiences do not affect beliefs about stocks (Malmendier 2021). Other times this research invokes broad effects, such as people becoming more individualistic after randomly receiving land titles (Di Tella et al. 2007) or relying on the experiences of past generations, such as immigration or slavery, in forming attitudes towards redistribution (Chinoy et al 2023). Work on priming effects raises similar concerns due to the disparity of effects across studies and to a poor understanding of the underlying mechanism (Cohen and Marechal 2016). These findings challenge a mechanical – and therefore stable – role of past experiences: we need a theory of which memories are used and how. Our model offers a cognitive mechanism reconciling experience and priming effects, and novel empirical tests based on the measured frequency and similarity of DS and NDS experiences.

While our applications focus on beliefs about novel events, our approach is relevant more broadly. Versions of our model have been applied to understand beliefs about career and college major choices (Conlon and Patel 2022), gender and pro-sociality (Exley et al. 2022), and stock returns (Jiang et al. 2023), but also the effect of wholly irrelevant idiosyncratic experiences on a person's macroeconomic expectations (Cenzon 2023). Graeber et al (2022) use the same framework to compare learning from stories vs. statistics, and Colonnelli et al. (2023) to explore how messages shape public support for bailout of large firms. Across these different contexts, explanatory power comes from spontaneous or cued retrieval of experiences, often NDS ones, which creates systematic disagreement and instability in beliefs.[3]

2. Motivating Evidence: Puzzles in Beliefs about Covid

Our model is motivated by three puzzling facts on beliefs about Covid we documented in 2020, the year the pandemic started and before vaccines became available. We describe the structure of the surveys, the facts, and their broader relevance for studying beliefs.

2.1 The Covid surveys

We ran three surveys, in May, July and November/December 2020 for a total of 4525 US participants. Qualtrics collected data, stratifying the sample on year of birth, gender, race (White,

[3] Simulation of has been linked to how people discount the future (Becker and Mulligan 1997, Gabaix and Laibson 2022), or how entrepreneurs imagine future outcomes (Ashraf et al. 2022). These papers neither consider memory nor similarity.

Black, Asian, Latino/a), household income, and region (Northeast, South, Midwest, West). The full surveys, including details on measurements and quality controls, are in Online Appendix B.

Beliefs about Covid-19 Risks. Our key outcome is the believed Covid fatality rate, $\hat{\pi}$, for the general US population, which we refer to as beliefs about "others." We elicit the distribution of $\hat{\pi}$ along three demographics: age, race, and gender. We ask subjects to consider "1,000 people in each of the following [AGE/RACE/GENDER] categories who contract Covid-19 in the next 9 weeks" and then to assess, for buckets in each category, how many people would die from Covid. Respondents also assess the fatality rate among people similar to them (in terms of age, gender, race, socioeconomic status, zip code, health, etc.).

Experiences. The second block measures demographics and personal experiences. We asked whether respondents – and separately, a family member – have been hospitalized for non-Covid related reasons in the last year. In waves 2 and 3, to study the belief formation mechanism, we elicit a broader set of experiences. We describe them when testing the theory in Section 4.

Estimating a cued rare event. At the beginning of the survey, participants were asked to estimate how many Americans have red hair, both out of 1,000 and out of 10,000 (these two answer fields appeared in a random order). This question works as a quality control and to familiarize respondents with the question format, but it more generally proxies for one's tendency to overestimate a cued rare event, which captures a key aspect of our framework.

2.2 Basic Facts

Figure 1 reports the binned frequency distribution of $\hat{\pi}$ for others, restricting to subjects who reported an estimate below 1000 (i.e. below 100%). The distribution for self is similar.

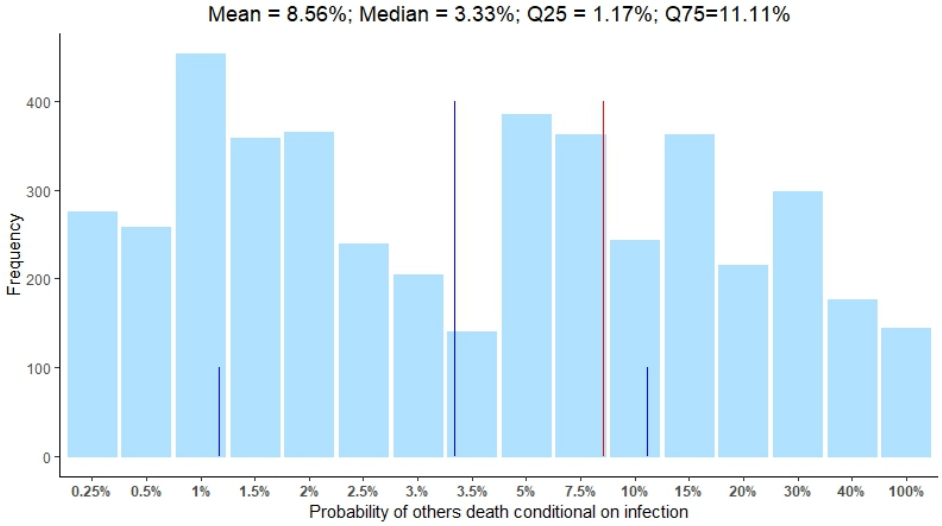

Figure 1

The Figure reports the distribution of $\hat{\pi}$ estimates for others, namely the estimated number of people, out of 1000 infected with Covid, who will die in the next 9 weeks. We elicit estimates for gender groups (male/female), age groups (0-39; 40-69; 70+) and race groups (White; African-American; Asian-American; Latinx-American) and average across them. To accommodate the skewed distribution, we use non-linear binning, and ticks on the x-axis refer to the upper limit of each bin. The vertical blue and red bars report the median and the mean, respectively. The small blue bars mark the interquartile range.

Two facts stand out. First, there is systematic overestimation of Covid's fatality rate. Median and mean estimates are at 3.3% and 8.6%, respectively. Scientific estimates at that time were about 0.68% (Meyerowitz-Katz and Merone 2020).[4] Second, there is large dispersion in estimates. The interquartile range of believed risk is 1.2% to 11%. Disagreement, in the form of a large mass of very pessimistic subjects, is responsible for the average overestimation of this risk.

Where do disagreement and overestimation come from? When first looking at the data, we found three factors: a respondent's age, their non-Covid bad health experiences, and their tendency to estimate a large share of red haired Americans. Figure 2 reports these facts. Panel A shows the age gradient, here documented for risks about self: older people are sharply less pessimistic about Covid risks for themselves than the young are for themselves. The 18-30 age group reports a mean fatality rate for self of 8% (median 2%), compared to 3.6% for the 69+ group (median 1%). This is very counterfactual: Covid death risk is much higher for the elderly. The young hugely overestimate their risk, which is at 0.01%, while the elderly underestimate theirs, which is at 4.7% (Levin et al. 2020).

[4] Similar estimates appear in contemporaneous papers, see Covid CDC (2020), Russell et al (2020), Modi et al (2021).

The young are also more pessimistic about others. This is puzzling: one may have expected pessimism by the elderly, due to possibly greater exposure to Covid deaths of other elderly people.

Panel B reports the effect of health adversities: a recent non-Covid hospitalization increases estimated fatality rates for *others* by nearly 50%, from 7.9% to 11.8% (similar effects arise for other non-Covid health adversities, see Section 4). It is puzzling that such an idiosyncratic health shock so strongly influences perceptions of Covid risks for the general population. One can perhaps argue that some hospitalizations are for respiratory disease, and as such rationally influence Covid death estimates, but the 11.8% is a much larger overestimate of general population risks than the 7.9% for people who do not share this experience. In the experiment in section 5, we examine the causal effect on beliefs of priming an objectively unrelated, but perceived to be similar, NDS experience.

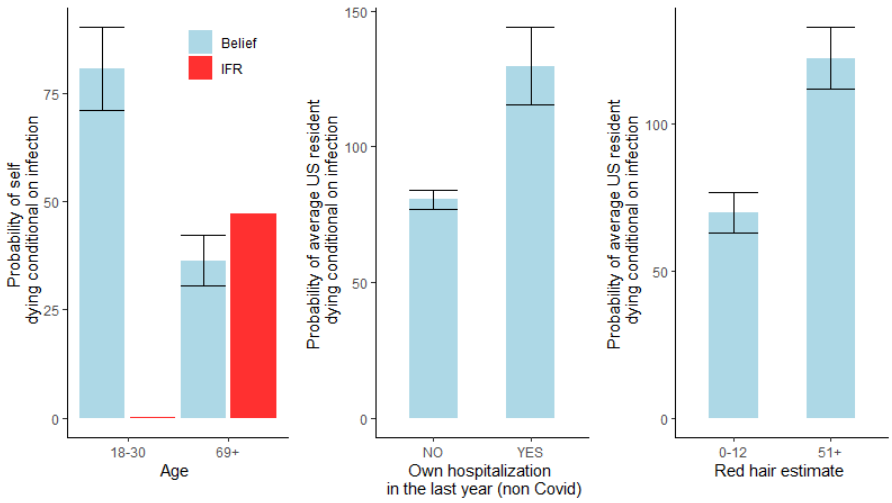

Figure 2

Panel A reports median and mean estimates of Covid fatality rate for self in the lowest and in the highest quintiles of age. The benchmark infection fatality ratio (IFR) is calculated for the sample of respondent, by using the formula IFR $= 10^{-3.27+0.0524*Age}$, derived in the meta-analysis of Levin et al. (2020). Panels B and C report estimated fatality rate for others with 95% confidence intervals. In Panel B, data are split based on the respondent having been hospitalized in the last year (not for Covid). In Panel C, data are split based on the respondent estimating the share of red-haired Americans in the top versus bottom tercile.

Lastly, panel C shows that respondents estimating a greater share of red-haired Americans are more pessimistic about Covid: going from the lowest to highest tercile of red hair estimates increases estimated fatality rates by 74% (from 7.0% to 12.2%).

These findings raise three key challenges for existing theories of belief formation. First, there is no mechanical tendency to over- or underestimate low probabilities. The young systematically over-estimate lethality π, while the elderly underestimate it. Second, the tendency to overestimate Covid risks appears related, among other things, to superficially similar health problems. This is challenging for DS experience effects, in which events in one setting, such as stocks, affect beliefs in the same setting but not in similar ones such as bonds (Malmendier 2021). Third, the tendency to overestimate probabilities is correlated across domains, including those without personal risks or motivated content, such as estimating the share of red haired Americans.

The evidence on red-haired Americans points to the importance of cognitive factors for beliefs, as opposed to risk preferences or motivations.[5] The roles of health experiences and age point to memory. On the one hand, when thinking of Covid, some people associatively retrieve their own or a loved one's recent illness, prompting pessimism about the new disease. On the other hand, the many lived experiences of the elderly – including surviving other adversities – make it harder for them to focus on Covid as a specific source of risk, compared to the young for whom Covid faces relatively less interference from pre-existing experiences. As intuitive these effects may seem, they also pose a puzzle because they go in different directions, the former boosting and the latter dampening estimates. How can we determine which experiences have which effect? Existing work offers no guidance, either because it does not consider NDS experiences, or because it documents their effect without providing a theoretical framework.

To solve this impasse, and place testable structure on experience effects, we explicitly model the influence of experiences on beliefs based on the psychology of selective memory. The model accounts for the puzzles presented in this section in a unified way and offers new predictions. It explains the role of experiences we measured in survey waves 2 and 3 for beliefs about Covid, reconciling the conflicting expects of different experiences based on measured similarity. Its

[5] Politics has little explanatory power for Covid risk perceptions in our data (see Tables C4 and C5 in Appendix C).

mechanisms are supported by the evidence from a follow-up experiment, in which exogenously priming some respondents to recall NDS experiences shapes beliefs about the likelihood of a severe cyberattack. In both domains there is systematic disagreement based on NDS experiences. Selective memory and simulation unify, and shed new light on, priming and experience effects.

3. The model

A decision maker (DM) estimates the probability of an event H (vs. alternative $\overline{H}$), such as dying from Covid conditional on infection, whose true probability is π. She selectively recalls two types of information. The first type is statistical, captured by an accurate numerical estimate π acquired through news or experts. The second type is experiences, pertaining to oneself, one's social circle, or learned from the media. These are stored in a database E. When cued to assess H, with probability $1 - \theta$ the DM samples the statistic π, and reports its value. With probability θ, she samples experiences and uses them to simulate the target event. The easier it is to simulate, the higher the estimated likelihood; θ thus captures the DM's reliance on experiences.

3.1 Recall of Experiences and Simulation

Following a large body of memory research, recall of an experience depends on similarity, frequency and interference (Kahana 2012). Specifically, the event being forecast, H, and the current context act as cues for recall. Experiences more similar to this cue or more frequent in the database E are more likely to be retrieved, and they inhibit recall of less similar or less frequent ones.

As in Bordalo et al (2023), a symmetric function $S: E \times E \to [0,1]$ measures the similarity between experiences $u, v \in E$. Similarity increases in the number of features shared by u and v, and is maximal, equal to 1, when $u = v$. For instance, a death from Covid is more similar to a death from pneumonia than to a death from homicide, and even less similar to non-adverse, non-health-related experiences such as finding a job. Recency is also a form of similarity: recent experiences are more similar to the present because they occurred in a similar context. We define the similarity between sets of experiences $A, B \subset E$ as the average pairwise similarity of their elements,

$$S(A,B) = \sum_{u \in A} \sum_{v \in B} S(u,v) \frac{1}{|A|} \frac{1}{|B|}. \tag{1}$$

$S(A,B)$ is symmetric and rises in feature overlap between the members of A and B. The similarity between two disjoint subsets of E can be positive when their elements share some features.

The DM evaluates event H, such as death conditional on Covid infection, denoted by $H = D|C$. Denote by $S(e) \equiv S(e,H)$ the similarity between experience e and H defined as in (1).

Assumption 1. Cued Recall: *When thinking about H, the probability that the DM recalls experience $e \in E$, denoted $r(e)$, is proportional to its similarity to the event, $S(e)$*:

$$r(e) = \frac{S(e)}{\sum_{u \in E} S(u)}. \tag{2}$$

From the numerator of (2), experience $e \in E$ is sampled more frequently when it is more similar to H. When thinking about death from Covid, due to similarity we are likely to recall Covid deaths in the news or among acquaintances. The denominator of (2) captures interference: all experiences $u \in E$ compete for retrieval, and may inhibit recall of e, especially experiences u that that are either similar to the cue or frequent. A person exposed to many deaths from hunger may retrieve those rather than deaths from Covid, even when thinking about $D|C$. Interference is well-established in memory research (e.g., Jenkins and Dallenbach 1924; McGeoch 1932; Underwood 1957).[6] It reflects the fact that we cannot fully control what we recall, which causes forgetting. Interference is central to understanding why even DS experiences may be underweighted.

In Bordalo et al. (2023), the likelihood of H is assessed based on the number of its instances that are recalled, as in Equation (2). We offer a more general theory in which even an experience not belonging to H, $e' \notin H$, can be used to simulate the event, boosting its estimated probability.

[6] For example, recall from a target list of words suffers intrusions from other lists studied at the same time, particularly for words that are similar to the target list, resulting in lower likelihood of retrieval (Shiffrin 1970; Lohnas et al. 2015).

Simulation is known to be central for thinking about the future (Dougherty et al. 1997; Brown et al. 2000, Hassabis et al. 2007a,b, Schacter et al. 2012, Biderman, Bakkour, and Shohamy 2020). In our context, recalling a death from pneumonia may boost the DM's ability to imagine death from Covid, making her more pessimistic, even though the experience is not in $D|C$. In cognitive science, the ease of simulation increases with the similarity between the retrieved memory and the target event (Kahneman and Tversky 1981, Schacter et al. 2012, Woltz and Gardner 2015). Simulation entails using the intrinsic features of e to imagine H. The higher the number of features the two events have in common, i.e. the higher their similarity, the stronger is simulation.[7] We rely on this insight to formalize the simulation function.

Assumption 2. Simulation: *Based on the retrieved experience $e \in E$, the DM simulates H with a probability $\sigma(e) \in [0,1]$ that increases in similarity: $\sigma(e) \geq \sigma(u)$ if and only if $S(e) \geq S(u)$.*

Similarity has two roles: as a driver of recall (Assumption 1) and of simulation (Assumption 2). By (1) and (2), when sampling from E, the DM recalls $e \in E$ with probability $r(e)$, and successfully simulates H with probability $\sigma(e)$. The average simulation of H is then given by:

$$\hat{\pi}_E = \sum_{e \in E} r(e)\sigma(e) = \frac{\sum_{e \in E} \sigma(e) \cdot S(e)}{\sum_{e \in E} S(e)}. \tag{3}$$

The dual role of similarity entails a fundamental trade-off between simulation and interference in how past experience shape beliefs. We now explore these implications.

3.2 The Properties of Memory Based Beliefs

By combining the use of statistics and the use of experiences, a population of DMs with identical recall and simulation parameters produces an average assessment of H given by:

$$\hat{\pi} = (1-\theta)\pi + \theta\hat{\pi}_E, \tag{4}$$

[7] Simulation may weigh features differently than recall. For instance, heart attacks may be more conducive at simulating a Covid death than the flu because they are deadly, even though they are less similar overall.

which combines the statistical "truth" π with the experience-based estimate $\hat{\pi}_E$.[8] The target event is overestimated on average when $\hat{\pi}_E > \pi$ and underestimated otherwise.[9]

A key driver of belief distortions in our model is non-domain specific, NDS, experiences. To understand what we mean by this, suppose that only domain-specific, DS, experiences are retrieved and used. These are straightforwardly defined as the events constituting H, which in our running example are the number of Covid infections and fatalities used for estimating the infection fatality rate. Specifically, assume that only DS events are recalled, $e \in H \cup \overline{H} = D|C \cup \overline{D}|C$, and only recollections in H are used to simulate H itself. This occurs if similarity and simulation are "narrow", namely $S\big(e, H \cup \overline{H}\big) = \sigma(e, H) = 1$ and zero otherwise. In this case, (4) yields the frequentist assessment $\hat{\pi} = |H| / \big|H \cup \overline{H}\big|$. As long as the DS database is unbiased (contains the true frequency of H and $\overline{H}$), beliefs are unbiased: the DM estimates the true conditional probability, $\hat{\pi} = \pi$. There is also no disagreement if everyone's DS experiences are the same.

NDS experiences matter, and create biases, because in reality neither similarity nor simulation are narrow. Thinking about a Covid death may cause the DM to think about other, NDS, adversities she experienced earlier in life. Even if her DS experiences may in principle enable her to produce a correct judgment, biases arise because DS ones are neglected, NDS ones are retrieved, and simulation is based on the latter. NDS experiences may explain why many survey respondents vastly overestimate Covid risks compared to the count of Covid fatalities. In the overall population, though, will over-or under-estimation prevail? And which kind of disagreement will emerge?

---

8 We assume that, when forming beliefs about a target event, the DM does not think about the alternative hypothesis $\overline{H}$ of it not happening. Formally, this is not material because, as discussed in the text, by ruling out NDS the DM is still capable of reaching the correct conditional assessment. Results are similar if, as in Bordalo et al. (2023) we relax this assumption, and the intuition is that recall errors matter more for distorting the unlikely $D|C$ event than its alternative. Note that this assumption is also in line with our survey and our experiment, in which respondents assess "Covid death" and "cyberattack" but the alternative hypothesis is not mentioned.

9 Here θ is exogenous. Graeber, Roth, and Zimmermann (2023) show that recall of experiences vs. statistics can also be understood based on memory: experiences are associated with more and diverse features, facilitating recall.

Proposition 1 *Suppose that the domain relevant database is unbiased,* $|H|/|H \cup \overline{H}| = \pi$*. If irrelevant experiences are recalled and used to simulate* H*,* $S(e) > 0$ *for* $e \neq H \cup \overline{H}$ *and* $\sigma(e) > 0$ *for* $e \neq H$*, there exists* π^* *such that* H *is overestimated if and only if its true frequency is low enough, namely* $\pi < \pi^*$*. If* $\pi < \pi^*$*, the overestimation increases in the DM's reliance on experience,* $\partial\hat{\pi}/\partial\theta > 0$.

NDS experiences exert two conflicting forces. On the one hand, they foster simulation of H, which boosts the memory-based estimate $\hat{\pi}$. On the other hand, they interfere with recall of other experiences, including DS experiences in H, which reduces $\hat{\pi}$. If H is rare, the effect of forgetting its few instances is weak, so over-estimation obtains. Even if there are very few Covid deaths, seeing many people in ICUs fosters simulation of death from Covid. People put positive probability on events they had never seen, provided they are similar enough to things they had seen.

Proposition 1 explains the observed tendency to overestimate rare events across many domains, including those without risk. Most people have few, if any, experiences of a rare event but simulation based on NDS experiences encourages overestimation. The reliance on NDS experiences is due to the key role of similarity and is not part of the approaches in which beliefs depend only on domain specific information (both mechanical experience effects and Bayesian models). Proposition 1 also shows a source of heterogeneity in this bias: a person's reliance on experience θ. In the Covid survey, we interpret the "red hair" variable as proxy for θ. In our experiment, we develop another proxy for θ, based on overestimation of an unrelated event, and find similar results.

Consider next the effect of specific DS and NDS experiences on beliefs:

Proposition 2 *Memory-based experience effects are shaped by similarity and interference:*

1. Similarity: experience e *boosts estimate* $\hat{\pi}$ *when added to database* E *if and only if it is sufficiently similar to* H *compared to an average member of* E*,* $\sigma(e) > \hat{\pi}_E$.

2. Interference: suppose $\sigma(e), \sigma(e') > \hat{\pi}_E$*. Adding* e *to* E *boosts* $\hat{\pi}$ *less if* e' *is also added to* E.

To see the effect of past experiences we must understand which of them are *recalled*, and how these are *used* for simulation. Crucially, even if DS experiences are available, NDS ones continue to matter. Point 1 says that one key driver of experience effects is perceived similarity. More similar experiences are more likely to produce successful simulation (they are also more likely to be recalled), so having similar experiences boosts estimates. This principle creates a similarity-based hierarchy of experience effects. At one extreme, DS experiences have maximal similarity, so they boost estimates, but NDS experiences can also boost estimates provided they are similar enough. At the opposite end, highly dissimilar experiences reduce estimates. Increasing their frequency hinders recall of better simulation material. Direct experiences with Covid should thus boost Covid pessimism more than experiences with other diseases, which are less similar, and even more compared to even less similar non-health adversities. This is a key prediction. It is a priori difficult to determine whether a given experience should be relevant and, if so, whether it should boost or dampen beliefs. Our model offers a solution: measure the perceived similarity of the experience with the target event. We later show the power of this method.

The second driver is interference across experiences due to competition for retrieval. Having past experiences conducive to simulation and hence to higher $\hat{\pi}$ reduces the sensitivity of beliefs to other experiences also conducive to simulation. Critically, this implies that a NDS experience can even interfere with a DS one (e.g. Proposition 2 holds even if $\sigma(e) > \sigma(e')$). People may exhibit muted reaction to relevant information due to irrelevant information in their database. Experiences with other adversities may reduce a person's sensitivity to direct experiences with Covid. Different pieces of information are selected, seldom integrated. This is another key way in which selective memory places further testable structure on experience effects. We later test this prediction as well.

3.3 Broader Implications and Roadmap

The role of selective recall and simulation for belief formation represents a major departure from existing frameworks. In conventional Bayesian models, beliefs are shaped by priors and react to data using likelihoods. For novel events such as Covid, the prior is typically agnostic and beliefs react strongly to the first instances of domain relevant data. This is also the case with Case-Based learning (Schank 1986, Gilboa and Schmeidler 1995), where people may be biased if the DS part of the database is biased, but biases cannot be produced by neglect of available DS data. Our model builds on a different perspective. People associatively sample their database, including experiences from different domains, and then use these experiences as material for imagination, sometimes interfering with the integration of domain-specific data. This approach is especially valuable for thinking about new risks, but it can shed light on heterogeneity and instability in other domains as well. When thinking about a fast growing firm, we may simulate its future success by thinking about Google, and forget less favourable data about the firm itself.

There are three drivers of disagreement in our model. The first is the database E, which varies across people. The second is similarity $S(u, v)$, which can also vary due to differential attention to features. For instance, a person focusing on the "death" and "respiratory" features of "Covid" will view it as less similar to "cancer" than a person focusing on "death" only. Third, reliance on experiences θ also varies across people. To assess these forces empirically, the database, similarity, and responsiveness to experiences must be measured. The range of personal experiences elicited should be broad, and include NDS ones that can matter for simulation and interference. Similarity must also be measured. The rest of the paper offers two illustrations of how to do this, including measuring similarity, which enables tests of Predictions 1 and 2.

Section 4 returns to beliefs about Covid and shows the usefulness of measuring the database E (survey waves 2 and 3), shedding light on the puzzles of Section 2. The Covid surveys were conducted before developing the model, and do not measure some important parameters. Section 5 presents an experiment on beliefs about cyberattacks, designed as a test of the model. Here we measure the database E and also similarity $S(e, H)$. The experiment includes a controlled change of the recall function $r(e)$ by priming respondents to recall different experiences before estimating the risk. In both of these domains, our results support Predictions 1 and 2.

4. Tests of Model Predictions in the Covid Surveys

We use the Covid surveys to test four predictions. Section 4.1 tests how similarity predicts the role of NDS experiences. Section 4.2 tests the role of interference across experiences. Section 4.3 tests two additional predictions that speak to the drivers of the age and red hair gradients.

To begin, we describe the measurement of experiences in survey waves 2 and 3, which offer a proxy for the databases E_i of each respondent i. After giving their estimate $\hat{\pi}_i$ respondents were asked whether they lived each of the following adversities: a serious life-threatening illness, a serious life-threatening accident or injury, having experienced poverty, a dangerous job, military service, or the untimely death or serious illness/injury of a loved one. We also ask them whether they have had Covid. To measure the severity of local conditions we use publicly available state-level data to build an index of pandemic severity, the cumulative level of deaths in the respondent's state at the time of maximal weekly case growth, and an index of recency, the days that have passed since that peak.[10] Table B.1 in Online Appendix C describes these covariates.

4.1 Similarity and Experience Effects in Beliefs about Covid

Point 1 in Proposition 2 yields the following prediction, proved (together with the others) in Appendix A. Remember that we measure beliefs about the event $D|C$ of Covid death.

Prediction 4.1 *If experiences e and e' are sources of Covid pessimism (higher $\hat{\pi}$), e predicts more pessimism than e' if and only if $S(e) > S(e')$. If e'' is a source of Covid optimism, then e'' is least similar, $S(e) > S(e') > S(e'')$.*

To test this prediction, we generate three sets of experiences with varying degrees of similarity to $D|C$. The first set consists of proxies for Covid experiences, including a *Had Covid* dummy and the index of pandemic severity (and recency). The second set consists of non-Covid health adversities, proxied by a *Health Adversities* index capturing having had a *serious illness* or a *serious injury*. The third set is an index of *Non Health Adversities* capturing whether the respondent has: i) experienced

[10] Accessible from the New York Times counts, https://www.nytimes.com/interactive/2021/us/covid-cases.html.

poverty, ii) worked at a job that carried serious health or safety risks, iii) performed military service, or iv) faced a serious injury, illness or untimely death of a loved one.[11] Proposition 1 implies that personal Covid experiences should boost Covid pessimism because they are more similar to $D|C$ than other adversities experienced in life. Crucially, this should also hold for the experience of having had, and survived, Covid. Non-Covid health adversities should be less of a booster of Covid pessimism, due to their lower similarity to $D|C$.[12] Finally, non health-adversities such as poverty, war, dangerous jobs, etc, should be associated with even less Covid pessimism because they are intuitively least similar to $D|C$. If these experiences come to mind, they do not help simulating $D|C$ but they block retrieval of better simulation material such as Covid deaths or other health problems. This hierarchy is consistent with a measurement of similarity we performed in May 2022 (see Appendix B) in which the average respondent rated experiences in the *Health Adversities* index as more similar to Covid death than those in the *Non Health Adversities* index.[13] In the cyberattack experiment we perform a more precise measurement, which can be systematically used for data analysis (see Section 5).

We regress beliefs for others $\hat{\pi}$ on the experience proxies above, as well as other experiences measured in all three Covid surveys: indicators for a recent own hospitalization, or hospitalization of a family member, for non-Covid reasons, the number of an individual's own health conditions, and subjective adversity.[14] We also include the red hair proxy, age, and a set of controls. We select controls using standard methods that pick the most reliable predictors of $\hat{\pi}_i$ from our full dataset, including all variables and all waves (Guyon and Elisseeff, 2003; James et al., 2013). See Online Appendix D for details. Table 1 reports the estimated coefficients for the theory-based predictors.

---

[11] We include *untimely death of a loved one* for it reflects enduring personal hardship, creating a non-health adversity. If we omit this experience the non-health adversities index retains the negative sign with a p value of 0.06 (see Table C4).

[12] Absent information on the database E, our model does not predict absolute effects of an experience. It is in particular consistent with a negative association between *Had Covid* and beliefs. The robust prediction of the model is that experiences of other illnesses should be associated with less pessimism than experiences with Covid.

[13] Full details of this survey are in Appendix B. The average rank (low rank means high similarity to Covid fatality) for the two components of *Health Adversities* is 3.4; that for the four components of *Non-Health Adversities* is 5.11.

[14] These health experiences are contemporaneous with Covid, so we do not include them in *Health Adversities*. Being focused on remote experiences, this index offers a stronger test of memory. *Subjective Adversity* captures (on a 1-7 scale) whether the person agrees with "Over the course of my life, I have experienced significant adversity".

Except for dummies, covariates are standardized to render coefficients comparable (Table C4 in Appendix C shows the full output).

Table 1. The Impact of Experiences on Covid Fatality Estimates

	OLS Predicting Beliefs Covid Fatality for Others
Had Covid	0.441***
	(0.167)
Health adversities	0.047**
	(0.019)
Non health adversities	-0.039***
	(0.015)
Hospitalization (self)	0.157**
	(0.073)
Hospitalization (family)	0.058
	(0.044)
No. Health Conditions	0.012
	(0.017)
Subjective Adversity	0.043**
	(0.019)
State Covid Level	0.059***
	(0.023)
Days	-0.097***
	(0.023)
Red hair	0.165***
	(0.033)
Age	-0.212***
	(0.021)
Controls	Y
Observations	2,953
Adjusted R^2	0.133

Notes: The dependent variable is the Covid fatality rate estimate for others, as defined in footnote 3. All variables, except for dummies, are standardized. Health adversities is an index given by the sum of two dummies indicating if the respondent ever suffered 1) a serious, life-threatening accident or injury; 2) a serious, life-threatening illness. Non health adversities is an index given by the sum of four dummies indicating if the respondent 1) worked a job that carried serious health or safety risks; 2) experienced military service; 3) experienced poverty; 4) experienced serious injury, illness, or untimely death of a loved one. The controls are the remaining selected variables (family hospitalization, number of health conditions, sex, race/ethnicity, education, income, region), which we omit together with the constant for readability. Clustered standard errors at state level. * denotes p<0.1, ** p<0.05, *** p<0.01.

Consistent with Prediction 4.1, *Had Covid* strongly predicts Covid pessimism (its coefficient cannot be directly compared to that of non-dummy regressors, which are standardized). This is a distinctive consequence of simulation: having had Covid, especially if severe, can make it easier to imagine less lucky or more vulnerable people dying from it.[15] In a Bayesian world, by contrast, surviving Covid should arguably promote optimism. The number of peak Covid deaths in a state, *level*, lead to pessimism but the effect fades over time as implied by the negative coefficient of *days* (which captures days since the growth peak).

Also consistent with Prediction 4.1, *Non-Covid Health Adversities* predict pessimism, but less than *Level*, and also than *Had Covid* when coefficients are comparable, i.e. in the case of the self-hospital dummy. Finally, consistent with Prediction 4.1, *Non Health Adversities* act as a source of Covid optimism. The model attributes this to interference: having gone through a bumpy life makes it easier to retrieve non-Covid risks, reducing simulation of Covid deaths, and lowering estimates.

Consistent with the model, both DS and NDS experiences matter, with similarity predicting their effects. Comparing their quantitative effects one can see that moving from zero to four Non-Health Adversities reduces pessimism to an extent equivalent to reducing the number of cumulative deaths in the state from 17000 to 0. This is a large number, given that the maximum number of cumulative Covid deaths at peak in the data is 15700. The effect of Health Adversities is also substantial, estimated at about half this effect. NDS experiences can explain why some people may be optimistic/scared even in regions with very virulent/mild pandemic conditions. Memory can thus account for large disagreement about the same event, even for people currently experiencing similar conditions and exposed to similar information. In the same vein, Table 1 also confirms the role of older age as a driver of optimism and red hair as a driver of pessimism, to which we return below.

One objection to the results in Table 1 is that experienced adversities may be endogenous and driven by a factor, such as risk tolerance, that also affects beliefs about Covid. Although we cannot rule out endogeneity of experiences, this explanation is unpersuasive for three reasons. First, it cannot explain why personal adversities affect beliefs about others. Second, endogeneity may also affect

---

[15] We also measure indirect Covid experiences by asking whether the respondent knows someone who had Covid, someone who was hospitalized for Covid, or someone who died from Covid. All these controls have positive coefficients (consistent with simulation) but only the last one is statistically significant. When we ran our surveys Covid was relatively rare, so local Covid conditions ("State Covid Level") may better capture indirect Covid experiences.

health adversities, such as illness, injury, and of course having had Covid. It is unlikely that risk tolerance generates pessimism for these experiences but optimism for others. Third, and crucially, risk tolerance cannot explain the role of age and red hair estimates which we study later, nor the results of our cyberattack experiment.

The results in Table 1 do not consider political affiliations, which may be important for decisions. We measure, and in robustness tests control for, political affiliation, and show that it has little explanatory power for estimates of Covid risks (Table C4 in Appendix C). We also collected data on self-reported individual behaviour, and find that Covid pessimism, instrumented by red hair estimates, explains more cautious behaviour. Thus, cognitive factors and experiences influence behaviour through beliefs. Political affiliation instead matters for attitudes toward lockdown policies (Table C5), in line with existing evidence (Allcott et al 2020, Bursztyn et al 2020).

4.2 Interference Across Experiences

Proposition 2.2 makes the following prediction for interference in beliefs about Covid.

Prediction 4.2 *The beliefs of a respondent who has experienced a non-Covid health adversity should, ceteris paribus, be less responsive to the severity of local pandemic conditions.*

A person who experienced a source of pessimism, such as a non-Covid health adversity, should react less to a given pandemic severity than a person who did not have such an experience. This is because recall of a DS experience can be interfered with by recall of an NDS one (and vice versa). As a test, we study interference of *Health Adversities* with local pandemic severity, as measured by *level*. In Figure 3, we assess the impact of *level* on beliefs, comparing respondents in the bottom tercile of local severity to those in the top tercile, for people with no versus some positive number of reported health adversities. Each dot reports the average Covid pessimism in the corresponding sample, measured by the average residual obtained from regressing $\hat{\pi}_i$ on all regressors of Table 1 Column 2 except for these two variables.

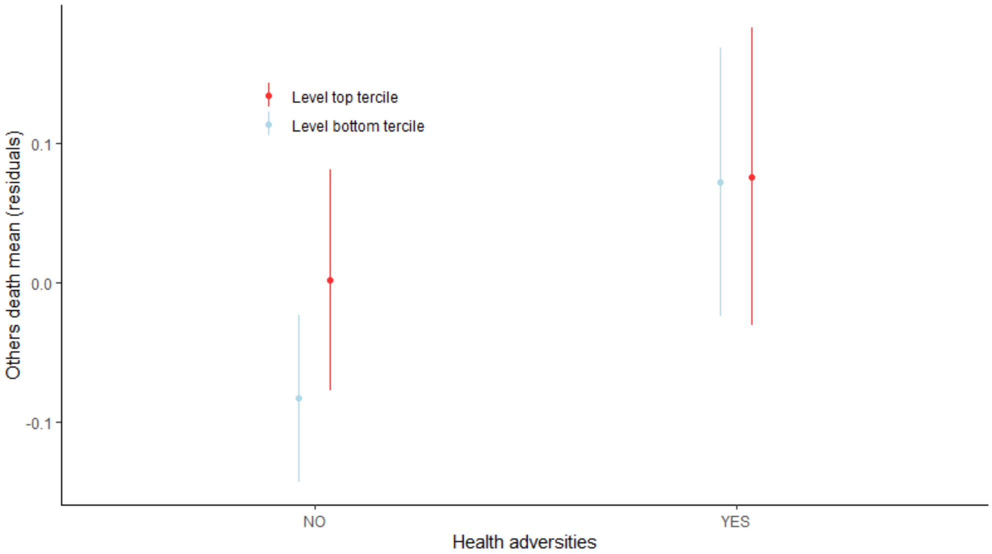

Figure 3.

The Figure reports the residuals of the standardized beliefs of Covid fatality rate (for others), estimated by removing from the model in column 2 of Table 2 the variables level, Hosp self, and Health adversities. The dummy on Health adversities measures whether the respondent reported 1 or more adversities (which in this specification include hospitalization, serious injury, and serious illness). Bottom and top tercile of level refer to the terciles of the distribution of level defined on waves 2 and 3 (when all health adversities are measured). Reported values are average residuals in each cell.

For respondents who have had no *Health Adversities*, moving from the bottom to the top tercile of *level* is associated with an increase in pessimism of $0.08 = 0.00 - (-0.08)$ ($p = 0.09$) of a standard deviation in beliefs. For respondents who have had *Health Adversities*, the same change in *level* has no impact on beliefs about Covid. Consistent with interference, having had non-Covid health adversities increases pessimism but also dampens the sensitivity of beliefs to DS experiences. Consistent with Proposition 2, interference is mutual: the effect of *Health Adversities* on beliefs is also dampened when the pandemic is severe. This points to a key property of selective memory: people do not integrate different pieces of information. They think about one or the other. Appendix C extends this analysis to the sources of pessimism in Table 1, finding consistent results.

4.3 Experiences, Age and Red Hair

Having documented the role of experience-based simulation and interference, we now show that they shed light on the other two puzzles of Section 2, the age and red hair gradients.

Prediction 4.3 *Age and Red Hair shape the impact of experiences on beliefs as follows.*

i) Older people, with a larger database of Non-Covid experiences, should ceteris paribus be on average more optimistic (lower $\hat{\pi}$). Furthermore, their beliefs should be less sensitive to any given experience e, whether it is a source of Covid pessimism or optimism.

ii) Higher reliance on experiences θ implies that people who estimate more red haired Americans should also be more pessimistic about Covid. In turn, these people should be more sensitive to any given experience e, whether it is a source of Covid pessimism or optimism.

Interference can explain the striking age effect (Prediction 4.3i). The database of the elderly is populated by many non-Covid experiences, since Covid is a new shock. These experiences create interference in retrieving Covid deaths, causing optimism. Critically, the same mechanism implies that the elderly should be less sensitive than the young to any specific experience they lived, as the latter is interfered with by many other experiences over a long lifetime. This account is consistent with memory research, which finds that the failure to remember specific events is to a large extent caused by a failure of cued retrieval (Shiffrin 1970).[16] An older person forgetting whether they locked the door earlier that day is failing to retrieve the exact event among many similar ones in the past (Wingfield and Kahana 2016). Our model captures this phenomenon.

Consider next Proposition 4.3 ii). Our model explains the role of the red hair proxy as capturing greater reliance on experience, θ. Critically, this implies that respondents with high red hair estimates should be disproportionally pessimistic if they experience sources of pessimism, and disproportionally optimistic if they experience sources of optimism. This prediction links the red hair proxy to recall and use of lived experiences, ruling out its interpretation as a mechanical tendency toward insensitivity or to report "high numbers", due to noise or other mechanisms (Kahneman and Tversky 1979, Enke and Graeber 2023, Khaw et al 2020, Abdellaoui et al 2011).

[16] There is evidence that memories "physically" degrade, which also causes forgetting and reduces the size of the database of the elderly compared to what it could have been otherwise. Our analysis requires that such degrading be sufficiently low that the elderly have a larger database of non-Covid experiences than the young. Consistent with this, in our data the elderly report having lived, on average, a larger number of the experiences we ask about than the young. Also consistent with this account is the finding that the elderly had lower stress levels and depression than the young during the pandemic (Fields et al., 2022).

We test prediction 4.3 by estimating separately the specifications of Table 1 for the top age tercile (people 62 or older) and the rest, and for the top red hair tercile and the rest, using all the available waves for the relevant experience. For each measured experience, we compute the difference in its estimated coefficient between the top tercile and the rest. This difference captures the differential "reactivity of beliefs to an experience" of respondents in the top group. Prediction 4.3 implies that the difference between older and younger people should be negative for sources of pessimism (e.g. the elderly should react less pessimistically than the young to health adversities), and positive for sources of optimism (e.g. the elderly should react less optimistically to non-health adversities). The opposite pattern should occur in the red hair split, with those in the top tercile of reactivity of beliefs adjusting more in response to both sources of optimism and pessimism.

Figure 4 reports the difference in the estimated coefficients for various measured experiences, where sources of pessimism and optimism are coloured in red and blue, respectively. On the left, we report the old-young difference, on the right the high-low red hair difference.

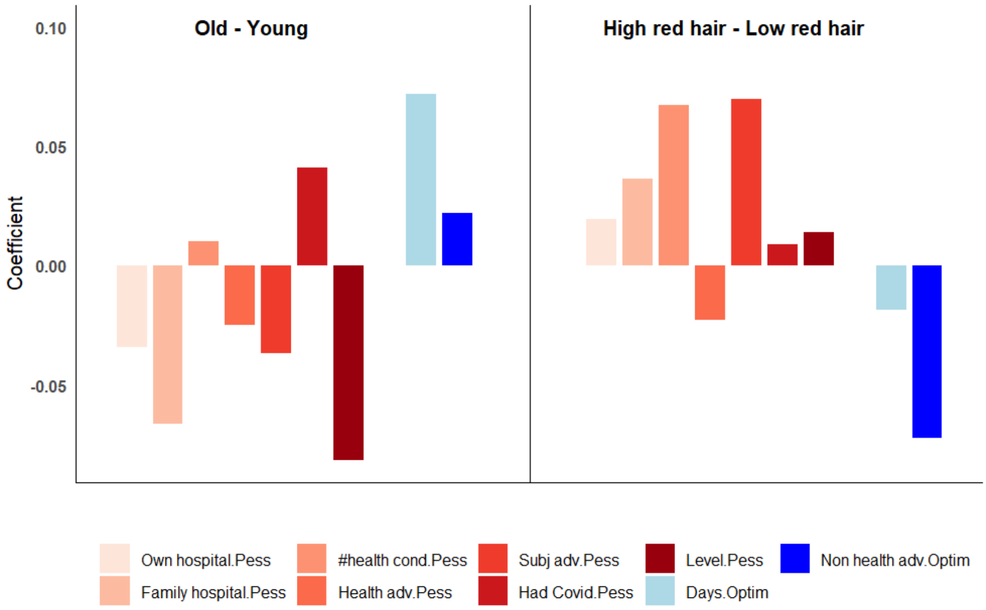

Figure 4

The left panel reports the difference between the coefficients of the specification for Covid fatality rate (others) of Tables 1 and 2 estimated in the top tercile of age (62+) and those estimated in the first two terciles of age (18-61). The right panel reports the difference between the coefficients of the specification for Covid fatality rate (others) of Tables 1 and 2 estimated in the top tercile of red hair estimates (more than 50) and those estimated in the first two terciles of red hair estimates (up to 50 out of 1000). Coefficients for variables available in all waves (hospital self, hospital family, no. health conditions, age, level, days) were obtained by estimating the model from column 2 in Table 1. Coefficients for variables available in waves 2 & 3 only (health adversities, non-health adversities, had Covid) were obtained by estimating the model from column 2 in Table 2. For comparability, all variables (including dummies) were standardized. Variables inducing pessimism (optimism) in the estimates for beliefs of others death from Tables 1 and 2 are in shaded of red (blue).

The results are broadly consistent with Predictions 4.3i) and 4.3ii). The elderly tend to react less pessimistically to sources of pessimism such as own and family non covid hospitalization, and less optimistically to sources of optimism such as non-health adversities.[17] Two exceptions to the pattern are the number of health conditions and having had Covid, experiences to which the elderly react more than the young.[18] Also consistent with our predictions, high red hair respondents tend to be more sensitive to determinants of pessimism and of optimism than low red hair respondents. The only exception is the (statistically insignificant) health adversity dummy. Using an F-test, we can reject that the coefficients are identical across the age groups, or across the red hair groups.[19]

Mean overestimation and strong disagreement in beliefs about Covid, as well as the puzzles in Figure 1, can be explained by a fundamental cognitive mechanism: memory-based beliefs. This mechanism yields similarity-based experience effects and interference across experiences, two key forces that allow NDS experiences to shape beliefs, yielding biased and heterogeneous reactions.

5. Recall, Similarity and Beliefs about Cyberattacks: An Experiment

Our Covid survey was designed prior to the development of the model, and it does not fully explore the memory mechanisms described in Section 3. In May 2023 we ran an experiment designed to more directly test the model (pre-registration AEARCTR-0011344). The experiment elicits beliefs about a different novel risk: a severe cyberattack in the US. It shows that the structure of memory-based beliefs documented in the Covid data extends to other domains.

[17] Older people might react less to news because they have more information. However, they are not more accurate. The median person over 72 underestimates own lethality by 3.1%, the median person between 65 and 71 does so by 1.3%.
[18] Both effects are statistically insignificant. The elderly's stronger reaction to having had Covid may arise because Covid is much more severe for them than for the young, so having had Covid is more similar to a Covid death for an older respondent. This underlines the importance of measuring individual-level similarity, which we do in Section 5.
[19] A test on the interaction of age with all variables included in all waves gives $p = 0.01$. For red hair we obtain $p = 0.06$. For variables only included in waves 2 and 3, p values are 0.06 and 0.03, respectively.

The goal of this experiment is to identify the key role of simulation and similarity in shaping the effect of different personal experiences on beliefs. The experiment builds on Proposition 2, in which a past experience, e, shapes beliefs through endogenous recall $r_i(e)$, its perceived similarity $S_i(e)$ to event H, and its interference with other lived experiences, e'. To identify these effects, we include two key design features. First, we exogenously prompt some subjects to recall a past experience, e. In our model, this treatment not only affects the recall $r_i(e)$ of this experience, but it also interferes with recall of other lived experiences e' that are different from e. Second, at the end of the survey we measure a subject's perceived similarity $S_i(e)$ between the full set of measured experiences and the event H. This experiment allows us to connect priming and experience effects through similarity, which has not been done before. Such analysis may in turn be relevant for understanding "information interventions" more broadly. Section 5.1 describes the survey. Section 5.2 maps primed recall to Proposition 2. Section 5.3 tests the empirical predictions.

5.1 The Survey

We ran the study on Prolific in May 2023 with 3,000 participants. The survey instrument is in Appendix E. We start by collecting a proxy for reliance on experiences. To ensure that our results are not driven by specificities of assessments about red hair, subjects estimate the number of U.S. cities out of 100 that receive more than 1ft of snow in a typical year (we denote the assessment by *snow*). We interpret higher assessments as greater reliance on experiences, higher θ_i.

Respondents are randomly allocated to four groups. In three of these, respondents are primed to recall a specific adverse experience e_p, with p being, respectively: i) *ID theft*: personal experience with identity theft, a data breach, stolen credit card information, or a compromised password, ii) *financial struggle*: personal experience of struggling with finances, and iii) *loved loss*: having lost a loved one to illness. These respondents are first asked whether they have had experience e_p. If so, they answer 4 brief, open-ended questions about it. If they have not had the experience, they move on to the next stage. A fourth *control* group is not primed.

The forecasted event in this survey is $H =$"significant cyberattack". We chose adverse experiences that vary in their perceived similarity to a cyberattack: we expect that *ID theft* is on average judged more similar to a cyberattack than *financial struggle* and *loved loss*, but we also expect individual-level variation in perceived similarity. *ID theft* is thus the most domain-specific experience, as there have been no severe nation-wide cyberattacks to date. In contrast, *financial struggle* and *loved loss* are NDS ones, yet if perceived as similar may affect beliefs about the target.

We next elicit beliefs. We provide a definition of a significant cyberattack: one that significantly disrupts critical civilian infrastructure, such as power lines, hospitals, banking system, communication satellites, or manufacturing. To reduce measurement error, we ask for two estimates: i) the likelihood on a 0 – 100 scale that they will be personally and significantly impacted by a cyberattack over the next 5 years, and ii) how many out of 1,000 people like them in the United States would be significantly impacted over those 5 years. Following these estimates, we ask participants how vividly they imagined a cyberattack when producing their estimates (1 – 7 scale).

At the end of our experiment, we ask participants whether they have experienced a wider set of adverse experiences, including the three possible primed experiences from the first stage but also memories of the Sept 11, 2001 terrorist attacks, a recent extreme weather event, a recent hospitalization, an addiction, and a serious accident/injury. This offers a proxy for the database E_i of each respondent. It also allows for the second key addition relative to the Covid survey: we ask participants to rate, on a 1 – 7 scale, how similar they believe each of these experiences is to a cyberattack. This is a proxy for a respondent's perceived similarity $S_i(e)$, which shapes spontaneous recall $r_i(e)$ and simulation $\sigma_i(e)$, two key model parameters. Capturing individual level variation in similarity $S_i(e)$ allows us to run sharper tests of the model.

5.2 Model Predictions: Similarity, Priming and Experience Effects

We next map the primed recall treatments to the model. If respondent i is in priming treatment p and has lived experience e_p, then she recalls it for sure, $r_{ip}(e_p) = 1$. Priming increases recall of e_p from the baseline no-priming probability, $r_i(e_p)$. Boosting recall of e_p is the key effect of priming. It is present only if e_p is forgotten with some likelihood absent priming, i.e. $r_i(e_p) < 1$. This point may seem obvious, but it is a key innovation of our model: experiences do not mechanically affect beliefs. Their role must be formalized in a model of retrieval and simulation.

The second effect of priming is to influence recall of non-primed experiences, $e \neq e_p$. When thinking about a cyberattack, experiences other than e_p may spontaneously come to mind. Critically, though, primed recall of e_p "pollutes" the retrieval context, creating interference. For instance, after being primed with $e_p =$ *financial struggle* the DM may spontaneously think of financial losses from the cyberattack, and find it harder to retrieve different experiences such as $e =$*September 11th*. Such interference from e_p to $e \neq e_p$ is well-known in memory research (Lohnas et al 2015). For instance, when recalling white things in a kitchen, cueing subjects with "milk" makes it more likely that they recall "yogurt" compared to less similar category members, such as "paper towel". This occurs because the cued item "milk" is both contextually close and similar to the question "white things in kitchen." Thus, it helps retrieval of items that are most similar to itself, inhibiting retrieval of other items. This effect directly follows from models of similarity based-retrieval like that in Equation (2): priming e_P just before H, and in the same survey, causes these events to share contextual features, over and beyond their intrinsic similarity. That is, overall similarity between them becomes $S_{ip}(e_P) = (1 + \kappa) \cdot S_i(e_P)$, where $\kappa > 0$ is the boost due to common context while $S_i(e_P)$ is their similarity based on intrinsic features.[20] Plugging $S_{ip}(e_P)$ in Equation (2) yields the interference effect

[20] This follows immediately from standard multidimensional scaling models in which similarity takes the form:

$$S(e,u) = \exp\left[-\delta \cdot \sum_l w_l \cdot (f_{el} - f_{ul})^2\right],$$

where f_{xl} if the value of feature $l = 1, \dots, L$ for event $x = e, u$ where w_l is the salience-based weight of the same feature, and $\delta > 0$. Context is one of the features, so recalling e_p close to the assessment of H reduces the contextual distance between the two, increasing similarity between them in the way described in the text.

of primed recall: it reduces spontaneous recall of other experiences compared to the baseline probability $r_i(e)$, i.e. $\Delta r_i(e) = r_{ip}(e) - r_i(e) \leq 0$ for $e \neq e_p$. Again, this effect is only present if there is some likelihood that e_p is not recalled absent priming; otherwise, $r_i(e_p) = 1$ implies $r_{ip}(e) = r_i(e) = 0$ for $e \neq e_p$.

In sum, our primed recall treatments can be mapped to the model as exogenous changes in recall for primed and non-primed experiences. This memory foundation places structure on priming effects, yielding new testable predictions. To see these predictions, note that the rest of the model is as before: recall of e leads to simulation $\sigma_i(e)$ of the cyberattack, where the simulation function is the same in treatment and control.[21] Denoting by I_p a dummy equal to one if the respondent is primed and recalls e_p, and assuming for simplicity that the statistic is zero ($\pi = 0$), beliefs are then given by:

$$\hat{\pi}_{ip} = \theta_i \sum_{e \in E_i} \sigma_i(e) r_i(e) + \theta_i I_p \left[\sigma_i(e_p) \Delta r_i(e_p) + \sum_{e \in E_i \setminus e_p} \sigma_i(e) \Delta r_i(e) \right]. \tag{5}$$

Beliefs reflect two "priming effect" terms, in square brackets, and one "experience effects" term, also present for subjects in the control treatment c. These effects are as follows.

Experience Effects: $\sum_{e \in E_i} \sigma_i(e) r_i(e)$. This term captures the fact that each respondent i has a database of experiences E_i that can be retrieved and used, even when nothing is primed. If an experience e is more similar to a cyberattack, $S_i(e)$ is higher, then adding it to the database raises belief $\hat{\pi}_{ip}$. This occurs because similarity boosts simulation (Assumption 2). This is a second key innovation of our model: to examine experience effects we must check how their recollections are used. The role of similarity in simulation places testable restrictions on this mechanism.

Direct Effect of Priming: $\sigma_i(e_p) \Delta r_i(e_p)$. Primed recall of e_p makes it available for simulation, which boost estimates. If the primed experience is more similar to a cyberattack, $S_i(e_p)$

[21] Priming does not affect the simulation function $\sigma_i(e)$ because the latter depends on intrinsic similarity, which our experiment does not manipulate. It may be possible to devise experiments that change $\sigma_i(e_p)$ by increasing the salience of intrinsic features shared by e_p and H.

is higher, the effect is stronger due to better simulation, $\sigma_i(e_p)$ is higher. The effect of similarity is not necessarily monotonic: if e_p is very similar to a cyberattack, it has a high probability of recall absent priming. Thus, $\Delta r_i(e_p)$ is lower, which attenuates (or reverses) the magnitude of the direct priming effect. This is our model's third key implication, and one that is important for the debate on the efficacy of priming: one should not expect priming to always work or work to the same extent for all people, because the effect of activating a mental association (in this case by recalling an experience) is mediated by its perceived similarity to the target.

Interference Effect of Priming: $\sum_{e\in E_i\setminus e_p}\sigma_i(e)\Delta r_i(e)$. This term captures the fact that, in our model, priming and experience effects interact. Primed recall of e_p interferes with recall of experiences different from it. This effect tends to reduce simulation, lowering $\hat{\pi}_{ip}$, dampening or possibly reversing the overall effect of priming on beliefs, compared to the direct effect.[22] This is the fourth key implication of our model: priming may backfire because it can interfere with retrieval and use of better simulation material. This can also help explain, in a structured way, the variability of priming effects across people as well as across messages that may have the same intention but different content (and hence different feature-based similarity to the target).

In sum, our model places a testable and psychologically-grounded structure on experience and priming effects. We now empirically evaluate this structure.

5.3 Similarity, Priming and Experience Effects: Empirical Tests

We report the distribution of estimates of the probability of a cyberattack. We see large heterogeneity in beliefs, just as in the Covid survey.

[22] This connects priming to Proposition 2, which describes how adding experience e to E shapes estimates via simulation and interference. One difference between living an experience (Proposition 2) and priming its recall is that priming creates "free retrieval", while still allowing for spontaneous recall of non-primed experiences when thinking about H. In fact, primed recall increases the overall amount of simulation material by $\sum_{e\in E_i}\Delta r_i(e)=\sum_{e\in E_i\setminus e_p} r_{ip}(e)>0$. Thus, regardless of the primed experience, priming tends to increase estimates relative to control.

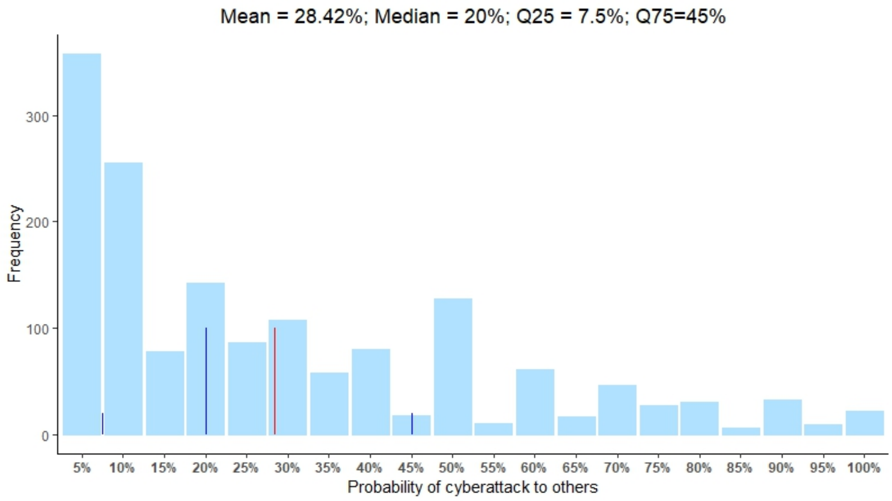

Figure 5

The Figure reports the distribution of $\hat{\pi}$ estimates for the likelihood of the respondent being personally and significantly impacted by a cyberattack over the next 5 years. The vertical blue and red bars report the median and the mean, respectively. The small blue bars mark the interquartile range.

To evaluate how selective recall of experiences (primed and non-primed) and similarity account for the observed belief heterogeneity, we first focus on the effect of priming experiences $e_p = ID\ theft, financial\ struggle$, *loved loss*, and then compare it to the role of these experiences when lived but not primed. We then develop a test to identify how priming interferes with other experiences, and include in this analysis all other measured experiences.

Neglecting interference and focusing only on direct priming and experience effects, yields:[23]

Prediction 5.1 Similarity in Priming and Experience Effects. *Living an experience e with higher measured similarity $S_i(e)$ with a cyberattack ceteris paribus increases the estimate $\hat{\pi}_{ip}$:*

i) due to the direct priming effect, if $e_p = e$, which eventually diminishes when $S_i(e_p)$ gets large,

ii) due to the experience effect, if $e_p \neq e$.

Similarity shapes priming and experience effects: priming a more similar lived experience tends to boost estimates more compared to control, but the priming effect eventually diminishes because a highly similar experience is recalled anyway. By the same token, having lived a more similar experience boosts simulation based on it, also boosting estimates more. This implies that,

[23] To aid intuition, here we focus on direct priming effects and neglect interference. The proof of Prediction 5.1 shows that the similarity hierarchy of priming effects also holds when interference is accounted for. Prediction 5.2 develops a test that separately identifies direct priming and interference effects.

looking across experiences, the strength of the effect of priming an experience should correlate with the strength of the effect of living that experience, even if not primed (due to spontaneous recall).

Table 2 assesses these predictions. In Column 1, we regress beliefs on a dummy equal to 1 if the respondent is randomly allocated to priming treatment p and zero otherwise. In this intent-to-treat (ITT) specification, some of the treated respondents have not lived e_p, so they are not prompted to recall it. The presence of these subjects dilutes the priming effect but ensures that the average database E_i is held constant across treatment and control, allowing us to focus on priming effects alone.

In Column 2, we restrict to respondents who in each treatment p have lived and hence recalled e_p, in a treatment-on-treated (TOT) approach. We regress beliefs on the recall dummy (I_p in Equation (5)) but also control for dummies for having had each of the experiences. This separates priming from experience effects, and allows us to ask whether experiences with higher average similarity to a cyberattack are associated with stronger priming *and* experience effects, in line with Prediction 5.1. In Column 3 we use individual-level measurement of similarity to construct, for each primed experience, a dummy for whether the subject attached to e_p a similarity that is below median compared to other subjects. This allows us to see whether subjects who perceive a given e_p as less similar to a cyberattack exhibit weaker priming effects compared to other participants.

We use data from high quality responses, defined as those who answer our numeracy questions correctly.[24] As a measure of beliefs, we aggregate the two cyberattack estimates by taking the z-score of each, averaging them, and then standardizing this average measure. We also control for the z-score of the respondent's *snow* estimate as well as demographics.[25]

Table 2. The Impact of Primed and Lived Experiences on Cyberattack Estimates

[24] Numeracy questions included converting a share (2%) into the absolute number of cases per 5000 observations, as well as giving a consistent answer to the snow question when framed in terms of "out of 100" and "out of 1000". We drop 3 people who did not say "yes" when asked to commit to providing thoughtful responses. Among high quality respondents, the TOT restriction drops 29% of the *ID theft*, 12% of *financial struggle*, and 32% of the *loved loss* groups, reflecting that a greater share of our respondents have experienced financial struggles relative to the other experiences.

[25] In Appendix F, we reproduce this and other results by loosening quality restrictions and by considering "treatment on treated" or "intent to treat" whenever not studied in the text. The results are qualitatively similar. We also present results on vividness.

	OLS Predicting Index of Cyberattack Estimates		
	High Quality Only		
	ITT	TOT	
	1	2	3
ID theft prime	0.15**	0.22***	0.19**
	(0.061)	(0.069)	(0.086)
Financial struggle prime	0.14**	0.12*	0.22***
	(0.062)	(0.066)	(0.083)
Loved loss prime	0.057	0.10	0.22**
	(0.061)	(0.071)	(0.092)
Below Median Similarity x *ID theft* prime			0.056
			(0.10)
Below Median Similarity x *Financial struggle* prime			-0.17*
			(0.094)
Below Median Similarity x *Loved loss* prime			-0.21**
			(0.11)
Had *ID theft*		0.40***	0.40***
		(0.064)	(0.064)
Had *Financial struggle*		0.30***	0.29***
		(0.084)	(0.084)
Had *Loved loss*		0.026	0.023
		(0.060)	(0.060)
Snow	0.16***	0.19***	0.18***
	(0.023)	(0.025)	(0.025)
Controls	Y	Y	Y
Observations	2090	1703	1703
Adjusted R-squared	0.049	0.089	0.092

Notes: * denotes p<0.1, ** p<0.05, *** p<0.01. Controls include age, sex, race/ethnicity, education, income, region. The index of the cyberattack estimates is constructed by taking the z-score of each cyberattack estimate for the individual, averaging them, and then computing the z-score of the averaged measure. *Snow* is the z-score of the individual's estimated share of U.S. cities receiving more than 1ft of snow in a typical year. The *prime* indicators take 1 if the individual was randomly-assigned to that treatment. The *had* indicators take 1 if the individual reported having had that personal experience. *Below median similarity* is an indicator that takes 1 if the individual reported a below-median similarity assessment of the primed experience compared to others in the same treatment.

Column 1 shows that priming a respondent to recall a personal experience such as *ID theft* or *financial struggle* boosts estimates. Consistent with selective recall and simulation, these experiences help the respondent to imagine a cyberattack, but with some probability they are not retrieved if not primed. There is instead no priming effect for the *loved loss* prime. This is consistent with the role of similarity in Prediction 5.1i): on average, *ID theft* and *financial struggle* are judged to be significantly more similar to a cyberattack (mean of 5.95 and 3.22 on the 1 – 7 scale, respectively) than to *loved loss* (mean of 1.95). The coefficients of *ID theft* and *financial struggle* are very similar, which in this ITT approach may be due to the lower frequency of people who have lived *ID theft* compared to *financial struggle* (70% vs 88%). This evidence confirms our key idea that NDS experiences, if similar enough to the target, can increase assessments.

Column 2 moves to TOT estimates, which allows us to better assess the role of similarity for priming (abstracting from the prevalence of the experiences in the population) as well as to compare priming and experience effects. Consistent with Proposition 5.1, priming *ID theft* has a stronger effect than priming the less similar *financial struggle*, and *loss* continues to be associated with no priming effect. The strong effect of priming ID theft shows that many subjects fail to recall having lived this experience if unprompted, again pointing to the importance of selective memory for experience effects. Critically, and consistent with Prediction 5.1, the coefficients of the experience dummies are in line with the similarity based hierarchy of experience effects we saw in the Covid survey: having lived a more similar experience such as ID theft is more conducive to pessimism than having lived the less similar *financial struggle* and even more so compared to the least similar *loved loss*. Both DS and NDS experience effects shape assessments.

Column 3 further highlights the role of similarity, showing that – for *financial struggle* and *loved loss* – subjects who regard the given primed experience, e_p, as less similar to a cyberattack exhibit weaker priming effects compared to other participants. Among those who perceive the primed experience as more similar, we see significant and substantial priming effects for all three primes, including *loved loss*. The model suggests that these similarity-driven effects are concentrated among the *financial struggle* and *loved loss* primes at least in part because these experience are unlikely to be recalled when thinking about a cyberattack if not prompted. On the other hand, there is no role of individual-level similarity for the *ID theft* prime. This is consistent with the diminishing effect of similarity in Proposition 5.1.

This result is important because it shows that, in line with the memory perspective, heterogeneity in beliefs emerges not just due to differences in the experience database E_i, but also due to differences in the perceived similarity of experiences to an event. This may help explain why priming effects are often elusive and unstructured: they can move people in different directions, depending on their perceived similarities and also, as shown below, on interference. When news or primes are viewed as memory cues rather than as abstract statistical signals, selective memory places a structure on these highly heterogeneous reactions.

Finally, Table 2 supports the role of reliance on experience θ_i: in all three columns, the higher the estimate of the number of snowy cities, the higher the estimated chance of a cyberattack. We later show that, as in the case of red hair in the Covid survey, this result cannot be explained by noise or by a mechanical tendency of some people to report large numbers.

We next consider the role of interference using all measured experiences, not only the primed ones, to proxy for the database E_i. As we show in the Appendix, Equation (5) then yields the following estimating equation.

Prediction 5.2 *Memory-based beliefs can be approximated by the regression:*

$$\hat{\pi}_{ip} = a_0 + a_1 \cdot \bar{S}_i(E_i) + a_2 \cdot I_p \cdot S_i(e_p) + a_3 \cdot I_p \cdot S_i(e_p) \cdot \bar{S}_i(E_i \backslash e_p), \tag{6}$$

where $\bar{S}_i(E_i)$ *is the average similarity of all lived experiences,* $S_i(e_p)$ *the similarity of the prime,* $\bar{S}_i(E_i \backslash e_p)$ *the average similarity of all non-primed but lived experiences, and* $a_1, a_2 > 0$*,* $a_3 < 0$*.*

The first regressor, $\bar{S}_i(E_i)$, linearly approximates the first experience effects term in Equation (5): people who ceteris paribus lived experiences they overall perceive as more similar to a cyberattack report higher estimates, $a_1 > 0$, in both treatment and control. The second regressor, $S_i(e_p)$, linearly approximates direct priming effects in Equation (5): people who are primed to recall an experience they perceive as more similar to a cyberattack report higher estimates, $a_2 > 0$. The third, non-linear term captures the interference in Equation (5): people primed with an experience they perceive as highly similar to a cyberattack exhibit a muted effect from lived, non-primed experiences, $a_3 < 0$. Once an experience highly similar to a cyberattack is top of mind, it is hard to think about anything else.

To test Equation (6), we aggregate experiences into individual level similarity indices. Note the usefulness of measuring similarity at the individual level: in the Covid survey we were restricted to distinguishing coarse categories of experiences (i.e. Health and Non Health Adversities) based on average similarity ratings. To construct $\bar{S}_i(E_i)$ and $\bar{S}_i(E_i \backslash e_p)$ we take the z-score of the average perceived similarity of the relevant experiences. To construct $S_i(e_p)$, we take the z-score of the perceived similarity of the primed experience. We predict the index of cyberattack estimates from these key terms, including in addition our standard controls and *snow.* We restrict to high quality observations and use a treatment-on-treated approach.

Column 1 of Table 3 estimates Equation (6) without the interference term and confirms the findings of Table 2: greater similarity of lived and primed experiences positively predict cyberattack estimates, even when these experiences are NDS. The priming effect confirms that living an experience, even if relevant, is not enough for it to be recalled and used to form beliefs. Lived but not primed similar experiences also boost estimates, again consistent with Table 2. Similarity is important, because it fosters both recall and simulation. Equally important, the data reveals significant interference: experiences cannot be treated in isolation, because they interact in recall, as shown by the negative and significant interaction term in Column 2. Consistent with Prediction 5.2, when one experience is primed, non-primed experiences are interfered with by the former and hence they are less impactful on beliefs.

Table 3. Similarity and Responsiveness to Experiences

	OLS Predicting Index of Cyberattack Estimates High Quality Only, Treatment-on-Treated		
	Pooled	Below Median *Snow*	Above Median *Snow*

	1	2	3	4	5	6
$\bar{S}_i(E_i)$, Total Similarity of Lived Experiences	0.19***	0.20***	0.14***	0.14***	0.24***	0.25***
	(0.028)	(0.028)	(0.039)	(0.039)	(0.040)	(0.040)
$S_i(e_p)$, Similarity of Primed Experience	0.13***	0.19***	0.17***	0.20***	0.10**	0.19***
	(0.028)	(0.034)	(0.039)	(0.047)	(0.040)	(0.049)
$\bar{S}_i(E_i)\ x\ S_i(e_p)$		-0.059***		-0.039		-0.080***
		(0.020)		(0.030)		(0.027)
Snow	0.18***	0.18***	0.36***	0.36***	0.088*	0.093*
	(0.025)	(0.025)	(0.090)	(0.090)	(0.050)	(0.050)
Controls	Y	Y	Y	Y	Y	Y
Observations	1706	1706	868	868	838	838
Adjusted R-squared	0.107	0.111	0.079	0.079	0.102	0.110

Notes: * denotes p<0.1, ** p<0.05, *** p<0.01. Controls include age, sex, race/ethnicity, education, income, region. The index of the cyberattack estimates is constructed by taking the z-score of each cyberattack estimate for the individual, averaging them, and then computing the z-score of the averaged measure. $\bar{S}_i(E_i)$ is the z-score of the average perceived similarity of all lived, non-primed experiences; an individual who reports 0 lived experiences has a pre-standardized $\bar{S}_i(E_i)$ of 0. $S_i(e_p)$ is the z-score of the perceived similarity of the primed experience; an unprimed individual has a pre-standardized $S_i(e_p)$ of 0. *Snow* is the z-score of the individual's estimated share of U.S. cities receiving more than 1ft of snow in a typical year.

To study the role of reliance on experience θ_i, in Columns 3 through 6 we split the sample according to *snow,* our responsiveness to experiences proxy. A straightforward implication of Equations (5) and (6) is that memory effects should be stronger for people with above median snow (higher θ_i). Both experience effects and interference are consistent with these predictions: they are stronger for people reporting above median *snow*. Priming effects are instead equal across the two groups, which is intuitive: priming e_p makes it highly available even to people who would typically not rely on it (low θ_i). This result is consistent with the notion that people reporting higher level of Snow are do not just exhibit a tendency to report larger numbers. Rather, they are more sensitive to experience and similarity-based simulation.

In sum, the results of our experiment are consistent with the key mechanism of our model, as well as with our findings from the Covid survey. Beliefs depend on retrieval, which is imperfect (as shown by the fact that priming matters), but also on simulation (as shown by the fact that NDS events matter depending on their similarity to cyberattack). Consistent with the model, priming and experience effects share the same similarity-based hierarchy, and there is interference between the two. Similarity-based retrieval and simulation help account for the structure of beliefs.

6 Conclusion

When we ran our first survey in 2020, we were surprised to find that older people were so much more optimistic than the young about Covid risks, for themselves as well as for others, and that own non-Covid health adversities had such a strong impact on Covid pessimism for others. We felt this had to do with experiences, so we measured a wider range of them in surveys 2 and 3, including non-health related ones. This confirmed that beliefs about Covid depended on a broad range of past experiences, including those from very different domains. To account for these facts, we developed a model based on the psychology of memory and simulation, where non-domain-specific experiences can shape beliefs via simulation as well as by interfering with retrieval of relevant experiences. The model helps explain our initial Covid puzzles and identifies novel mechanisms of belief formation. We test and confirm this mechanism in the Covid survey and in a primed-recall experiment on cyberattack risk, specifically devised as a test of the model. The evidence shows the importance of similarity and interference in modulating and unifying experience and priming effects.

One important message that emerges from our analysis is that selective memory makes it hard for people to integrate different pieces of information. Once retrieved, one piece interferes with retrieval of another. This perspective is very different from Bayesian models or rational inattention, in which all information is integrated but slanted toward the prior due to noise. Our evidence is consistent with that from statistical problems (Bordalo et al. 2023b) in which memory also matters. Selective use of remembered information explains both the average overestimation of small risks and the high disagreement often observed in survey data, as the result of simulation and interference coming from different experiences. It implies that people facing the same event may neglect publicly available relevant data, and instead focus on irrelevant experiences, creating bias and heterogeneity.

Our model also offers a structured way to study the role of memory in belief formation. This requires measuring a broad range of experiences, including non-domain specific ones, but also individual level similarity judgments between these experiences and the event whose probability is being assessed. This structured approach can help improve surveys and priming experiments, going beyond an intuition about directional effects and accounting for the probability with which different experiences are spontaneously retrieved – which limits the effect of priming them – and for how they are used in simulation. Memory brings both new data and new predictions to the table.

The mechanism of simulation, together with our account of priming, can shed light on how narratives or political advertising could change beliefs by activating otherwise neglected experiences. For decades, Avis Car Rental Company, which lagged Hertz in sales, advertised itself with "We are number two. We try harder." This simulation of quality from unrelated experiences with hard-driving underdogs apparently worked for some potential customers. Simulation and interference offer a mechanism for persuasion that fosters retrieval of experiences that are good for simulating what the persuader is interested in, and interferes with conflicting thoughts. Crucially, this mechanism clarifies the role of individuals' own experiences, as well as their subjective perceptions of similarity, in understanding heterogeneity in the response to these messages.

More broadly, memory is a key input into our cognitive activities. Even the distinction between beliefs and preferences may be more tenuous than one thinks. When we assess a political candidate, a consumer product, or a financial asset, we imagine what the candidate would do once in office, the uses of the product, or the returns of the asset based on the thoughts that come to mind, usually from past experiences (Bordalo, Gennaioli, and Shleifer 2020). Growing neuroscientific evidence shows the key role of memory in this process (Shadlen and Shohamy 2016). This perspective creates exciting opportunities to explain economic choice with new models and new data.

References

Abdellaoui, Mohammed, Aurélien Baillon, Laetitia Placido, and Peter Wakker. 2011. "The Rich Domain of Uncertainty: Source Functions and their Experimental Implementation." *American Economic Review,* 101 (2): 695-723.

Alesina, Alberto, and Nicola Fuchs-Schündeln. 2007. "Good-Bye Lenin (or not?): The Effect of Communism on People's Preferences." *American Economic Review,* 97 (4): 1507-1528.

Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang. 2020. "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic." *Journal of Public Economics*, 191: 104254.

Andre, Peter, Ingar Haaland, Christopher Roth, and Johannes Wohlfart. 2021. "Inflation Narratives". *ECONtribute Discussion Paper* 127.

Ashraf, Nava, Gharad Bryan, Alexia Delfino, Emily Holmes, Leonardo Iacovone, and Ashley Pople. 2022. "Learning to See the World's Opportunities: The Impact of Imagery on Entrepreneurial Success." Working Paper.

Becker, Gary S., and Casey B. Mulligan. 1997. "The Endogenous Determination of Time Preference.", *Quarterly Journal of Economics*, 112 (3): 729–758.

Biderman, Natalie, Akram Bakkour, and Daphna Shohamy. 2020. "What are Memories for? The Hippocampus Bridges Past Experience with Future Decisions." *Trends in Cognitive Sciences,* 24 (7): 542-556.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2020. "Memory, Attention, and Choice." *Quarterly Journal of Economics,* 135(3):1399-1442.

Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer. 2023. "Memory and Probability." *Quarterly Journal of Economics*, 138(1): 265-311.

Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer. 2023b. "How People Use Statistics." National Bureau of Economic Research, w31631.

Brown, Norman, Lori Buchanan, and Roberto Cabeza. 2000. "Estimating the Frequency of Nonevents: The Role of Recollection Failure in False Recognition." *Psychonomic Bulletin & Review,* 7(4): 684-691.

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2020. "Misinformation During a Pandemic." *NBER* Working Paper 27417.

Chinoy, Sahil, Nathan Nunn, Sandra Sequeira, and Stefanie Stantcheva. 2023. "Zero-Sum Thinking and the Roots of US Political Divides." *NBER* Working Paper 31688.

Cohn, Alain, and Michel Maréchal. 2016. "Priming in Economics." *Current Opinion in Psychology*, 12: 17-21.

Colonnelli, Emanuele, Niels Gormsen and Timothy McQuade. 2023. "Selfish Corporations," *The Review of Economics Studies*, forthcoming.

Conlon, John, and Dev Patel. 2022. "What Jobs Come to Mind? Stereotypes about Fields of Study." Working Paper.

Covid, C. D. C., and Response Team. 2020. "Severe Outcomes among Patients with Coronavirus Disease 2019 (COVID-19)—United States, February 12–March 16, 2020." *Morbidity and Mortality Weekly Report,* 69.12 (2020): 343-346.

Dougherty, Michael, Charles Gettys, and Rickey Thomas. 1997. "The Role of Mental Simulation in Judgments of Likelihood." *Organizational Behavior and Human Decision Processes,* 70(2): 135-148.

Dryhurst, Sarah, Claudia Schneider, John Kerr, Alexandra Freeman, Gabriel Recchia, Anne Marthe Van Der Bles, David Spiegelhalter, and Sander van der Linden. 2020. "Risk Perceptions of COVID-19 around the World." *Journal of Risk Research*, 23(7-8): 994-1006.

Enke, Benjamin, and Thomas Graeber. 2023. "Cognitive Uncertainty." *Quarterly Journal of Economics,* forthcoming.

Exley, Christine, Oliver Hauser, Molly Moore, and John-Henry Pezzuto. 2022. "Beliefs about Gender Differences in Social Preferences." Working Paper.

Fields, Eric, Elizabeth Kensinger, Sandry Garcia, Jaclyn Ford, and Tony Cunningham. 2022. "With Age Comes Well-Being: Older Age Associated with Lower Stress, Negative Affect, and Depression Throughout the COVID-19 Pandemic." *Aging & mental health*, 26(10): 2071–2079.

Gabaix, Xavier, and David Laibson. 2022. "Myopia and Discounting." Working Paper.

Gilboa, Itzhak, and David Schmeidler. 1995. "Case-Based Decision Theory." *Quarterly Journal of Economics*, 110(3): 605-639.

Graeber, Thomas, Florian Zimmermann, and Christopher Roth. 2022. "Stories, Statistics, and Memory." Working paper.

Guyon, Isabelle, and Andre Elisseeff. 2003. "An Introduction to Variable and Feature Selection." *Journal of Machine Learning Research* 3: 1157–1182.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein. "Learning through Noticing: Theory and Evidence from a Field Experiment." *Quarterly Journal of Economics* 129 (2014): 1311-1353.

Hassabis, Demis, and Eleanor Maguire. 2007. "Deconstructing Episodic Memory with Construction." *Trends in Cognitive Sciences,* 11(7): 299-306.

Hassabis, Demis, Dharshan Kumaran, Seralynne Vann, and Eleanor Maguire. 2007. "Patients with Hippocampal Amnesia Cannot Imagine New Experiences." *Proceedings of the National Academy of Sciences,* 104 (5): 1726-1731.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. New York: Springer

Jenkins, John, and Karl Dallenbach. 1924. “Obliviscence during Sleep and Waking.” *American Journal of Psychology*, 35(4): 605-612.

Jiang, Zhengyang, Hongqi Liu, Cameron Peng, and Hongjun Yan. 2023. “Investor Memory and Biased Beliefs: Evidence from the Field.” *Mimeo.*

Kahana, Michael. 2012. *Foundations of Human Memory*. Oxford, UK: Oxford University Press.

Kahneman, Daniel, and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision Under Risk.” *Econometrica*, 47(2): 263-291.

Kahneman, Daniel, and Amos Tversky. 1981. “The Simulation Heuristic.” Working Paper.

Kaustia, Markku, and Samuli Knüpfer. 2008. “Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions.” *Journal of Finance*, 63 (6): 2679-2702.

Khaw, Mel Win, Ziang Li, and Michael Woodford. 2021. "Cognitive Imprecision and Small-stakes Risk Aversion." *Review of Economic Studies* 88: 1979-2013.

Kunreuther, Howard. 1978. *Disaster Insurance Protection: Public Policy Lessons*. New York: Wiley.

Levin, Andrew, William Hanage, Nana Owusu-Boaitey, Kensington Cochran, Seamus Walsh, and Gideon Meyerowitz-Katz. 2020. “Assessing the Age Specificity of Infection Fatality Rates for COVID-19: Systematic Review, Meta-analysis, and Public Policy Implications.” *European Journal of Epidemiology* 35 (12),1123-1138.

Lohnas, Lynn, Sean Polyn, and Michael Kahana. 2015. “Expanding the Scope of Memory Search: Modeling Intralist and Interlist Effects in Free Recall.” *Psychological Review* 122(2): 337-363.

Malmendier, Ulrike. 2021. “Exposure, Experience, and Expertise: Why Personal Histories Matter in Economics.” *Journal of the European Economic Association*, 19(6): 2857–2894.

Malmendier Ulrike, and Stefan Nagel. 2011. “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *Quarterly Journal of Economics,* 126(1): 373-416.

Malmendier, Ulrike, and Stefan Nagel. 2016. “Learning from Inflation Experiences.” *Quarterly Journal of Economics*, 131(1): 53-87.

McGeoch, John. 1932. “Forgetting and the Law of Disuse.” *Psychological Review,* 39(4): 352–370.

Meyerowitz-Katz, Gideon, and Lea Merone. 2020. “A Systematic Review and Meta-Analysis of Published Research Data on COVID-19 Infection-fatality Rates.” *International Journal of Infectious Diseases*, 101: 138-148.

Modi, Chirag, Vanessa Bohm, Simone Ferraro, George Stein, and Uros Seljak. 2021. “Estimating COVID-19 Mortality in Italy early in the COVID-19 Pandemic.” *Nature Communications,* 12 (1): 2729.

Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer. 2008. “Coarse Thinking and Persuasion.” *Quarterly Journal of Economics,* 123(2): 577-619.

Russell, Timothy, Joel Hellewell, Christopher Jarvis, Kevin Van Zandvoort, Sam Abbott, Ruwan Ratnayake, Stefan Flasche et al. 2020. "Estimating the Infection and Case Fatality Ratio for Coronavirus Disease (COVID-19) using Age-adjusted Data from the Outbreak on the Diamond Princess Cruise Ship, February 2020." *Eurosurveillance,* 25 (12): 2000256.

Schacter, Daniel, Donna Addis, and Randy Buckner. 2007. "Remembering the Past to Imagine the Future: the Prospective Brain." *Nature Reviews Neuroscience,* 8 (9): 657-661.

Schacter, Daniel, Donna Addis, Randy Bruckner. 2008. "Episodic Simulation of Future Events." *Annals N.Y. Academy of Science,* 1124 (1): 39-60.

Schacter, Daniel, Donna Addis, Demis Hassabis, Victoria Martin, Nathan Spreng, and Karl Szpunar. 2012. "The Future of Memory: Remembering, Imagining, and the Brain." *Neuron,* 76 (4): 677-94.

Schank, Roger. 1986. *Explanation Patterns: Understanding Mechanically and Creatively.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Schwartzstein, Joshua. "Selective Attention and Learning," *Journal of the European Economic Association*, 12 (2014): 1423-1452.

Shadlen, Michael, and Daphna Shohamy. 2016. "Decision Making and Sequential Sampling from Memory." *Neuron*, 90(5): 927–939.

Shiffrin, Richard. 1970. "Forgetting: Trace Erosion or Retrieval Failure?" *Science,* 168(3939): 1601-1603.

Sims, Christopher. 2003. "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

Underwood, Benton. 1957. "Interference and Forgetting." *Psychological Review,* 64(1), 49–60.

Weinstein, Neil. 1989. "Effects of Personal Experience on Self-Protective Behavior." *Psychological Bulletin*, 105 (1): 31 – 50.

Wingfield, Arthur, and Kahana, Michael. 2002. "The Dynamics of Memory Retrieval in Older Adulthood." *Canadian Journal of Experimental Psychology*, 56: 187-199.

Woltz, Dan, and Michael Gardner. 2015. "Semantic Priming Increases Word Frequency Judgments: Evidence for the Role of Memory Strength in Frequency Estimation." *Acta Psychologica*, *160*: 152-160.

Woodford, Michael. 2003. "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information and Expectations in Modern Macroeconomics*, Princeton: Princeton University Press.

Appendix A. Proofs

Proof of Proposition 1 In the normative benchmark in which only H experiences can be used to simulate H and in which only $H \cup \overline{H}$ are recalled, $\hat{\pi}_E = |H|/|H \cup \overline{H}| = \pi$. If experiences other than H can be used to simulate H by a factor $\tilde{\sigma} > 0$ and if experiences outside $H \cup \overline{H}$ are according to similarity $\tilde{S} > 0$, then using Equation (3) we have:

$$\hat{\pi}_E = \frac{|H| + \tilde{\sigma}\left[|\overline{H}| + \tilde{S}|E \backslash H \cup \overline{H}|\right]}{|H \cup \overline{H}| + \tilde{S}|E \backslash H \cup \overline{H}|},$$

which is larger than the frequentist estimate if and only if the true probability of H is sufficiently low:

$$\frac{|H|}{|H \cup \overline{H}|} = \pi < \pi^* \equiv \frac{\tilde{\sigma}|\overline{H}|}{\tilde{S}|E \backslash H \cup \overline{H}|} + \tilde{\sigma}.$$

From Equation (4) we have $\frac{\partial \hat{\pi}}{\partial \theta} = \hat{\pi}_E - \pi$, which is positive if $\pi < \pi^*$.

Proof of Proposition 2 Part 1: adding a measure $|e|$ of experience e and a measure $|e'|$ of experiences e' to a baseline database E is equal to, using Equation (3):

$$\hat{\pi}_{E \cup |e| \cup |e'|} = \frac{\sigma(e)S(e)|e| + \sigma(e')S(e')|e'| + \mathbb{E}_E(\sigma S)|E|}{S(e)|e| + S(e')|e'| + \mathbb{E}_E(S)|E|}, \qquad (A.1)$$

where $\mathbb{E}_x(.)$ denotes the average in set x. The effect of increasing $|e|$ is:

$$\frac{\partial \hat{\pi}_{E \cup |e| \cup |e'|}}{\partial |e|} = S(e)\frac{[\sigma(e) - \sigma(e')]S(e')|e'| + [\sigma(e)\mathbb{E}_E(S) - \mathbb{E}_E(\sigma S)]|E|}{[S(e)|e| + S(e')|e'| + \mathbb{E}_E(S)|E|]^2} \qquad (A.2)$$

When a single experience e is added to a large database E, in which $|e'| = 0$, we have:

$$\left.\frac{\partial \hat{\pi}_{E \cup |e| \cup |e'|}}{\partial |e|}\right|_{|e|,|e'|=0} = \frac{|E|\mathbb{E}_E(S)S(e)}{[\mathbb{E}_E(S)|E|]^2}[\sigma(e) - \hat{\pi}_E], \qquad (A.3)$$

Which is positive if and only if $\sigma(e) - \hat{\pi}_E > 0$.

Part 2: to study the effect adding an additional experience e' on the marginal effect of adding e, we take the derivative of (A.2) with respect to $|e'|$, which yields:

$$\frac{\partial \hat{\pi}_{E \cup |e| \cup |e'|}}{\partial |e| \partial |e'|} \propto$$
$$[\sigma(e) - \sigma(e')]S(e')[S(e)|e| + \sigma(e')S(e')|e'| + \mathbb{E}_E(S)|E|] -$$
$$2S(e')\big[[\sigma(e) - \sigma(e')]S(e')|e'| + [\sigma(e)\mathbb{E}_E(S) - \mathbb{E}_E(\sigma S)]|E|\big], \qquad (A.2)$$

Which after some simplification yields:

$$\left.\frac{\partial \hat{\pi}_{E \cup |e|}}{\partial |e| \partial |e'|}\right|_{|e|,|e'|=0} \propto -\sigma(e') - \sigma(e) + 2\hat{\pi}_E. \qquad (A.3)$$

If $\sigma(e'), \sigma(e) > \hat{\pi}_E$, then (A.3) is negative, which proves the result.

Proof of Prediction 4.1. This prediction follows from (A.3), which after some algebra yields:

$$\left.\frac{\partial \hat{\pi}_{E \cup |e| \cup |e'|}}{\partial |e| \partial S(e)}\right|_{|e|,|e'|=0} \propto \sigma(e) - \hat{\pi}_E + S(e)\frac{\partial \sigma(e)}{\partial S(e)}$$

This expression is positive provided similarity $S(e)$ is sufficiently high that $\sigma(e) - \hat{\pi}_E \geq 0$, namely e is a source of pessimism, because $\partial \sigma(e)/\partial S(e) \geq 0$. It can be negative only if e is sufficiently dissimilar, namely $S(e)$ is sufficiently low, that $\sigma(e) - \hat{\pi}_E < 0$, so e is a source of optimism.

Proof of Prediction 4.2. It follows directly from inspection of Equation (A.3), given than both Non Covid health adversities and *Level* are both sources of Pessimism.

Proof of Prediction 4.3. Consider first part i). We first show why older people should ceteris paribus be more optimistic about Covid than young people. Denote by NC and C the set of Non-Covid and Covid experiences of a person. We have:

$$\hat{\pi}_{NC\cup C} = \frac{\mathbb{E}_C(\sigma S)|C| + \mathbb{E}_{NC}(\sigma S)|NC|}{\mathbb{E}_C(S)|C| + \mathbb{E}_{NC}(S)|NC|}$$

elderly people have the same set C of young people (Covid is a common shock), but a larger set of Non Covid experiences $|NC|$. But then:

$$\frac{\partial \hat{\pi}_{NC\cup C}}{\partial |NC|} = \frac{\mathbb{E}_{NC}(S)\mathbb{E}_C(S)[\hat{\pi}_{NC} - \hat{\pi}_C]|C|}{[\mathbb{E}_C(S)|C| + \mathbb{E}_{NC}(S)|NC|]^2} < 0,$$

because $\hat{\pi}_{NC} - \hat{\pi}_C < 0$, namely Covid experiences allow better simulation of Covid than Non Covid ones. Consider next the elderly's reactivity to experiences. In (A.2), setting $|e'| = 0$ we can compute:

$$\frac{\partial \hat{\pi}_{E\cup|e|}}{\partial |e| \partial |E|} = -S(e)\mathbb{E}_E(S)\frac{2[\sigma(e) - \hat{\pi}_E]}{[\mathbb{E}_E(S)|E|]^2},$$

which is negative for sources of pessimism, $\sigma(e) - \hat{\pi}_E > 0$, and positive for sources of optimism, $\sigma(e) - \hat{\pi}_E < 0$. Thus age dampens the sensitivity to any experience.

Part ii) follows directly from Equation (4), which implies:

$$\frac{\partial^2 \hat{\pi}}{\partial |e| \partial \theta} = \frac{\partial \hat{\pi}}{\partial |e|}$$

so for respondents with higher θ the impact of experience e (whether a source of optimism or pessimism) on their expected belief gets amplified. The same applies to the average belief of a group of respondents with the same θ.

Proof of Prediction 5.1. Consider first the direct priming effect. In Equation (5), this is equal to:

$$\sigma_i(e_p)\Delta r_i(e_p) = \sigma_i(e_p)[1 - r_i(e_p)].$$

Because both $r_i(e_p)$ and $\sigma_i(e_p)$ increase in similarity $S_i(e_p)$, higher similarity has a non-monotonic effect on direct priming. In general, because there are many experiences in the database, we expect $r_i(e_p)$ to be small. As a result, we expect direct priming to be increasing, but at a diminishing rate. To map the regressions of Table 2, consider the full priming effect, including also interference. From Equation (5) this is equal to $\sigma_i(e_p)\Delta r_i(e_p) + \sum_{e\in E_i\backslash e_p}\sigma_i(e)\Delta r_i(e)$, which can be rewritten as:

$$\sigma_i(e_p)[1 - r_{ip}(e_p)] + \sum_{e\in E_i}\sigma_i(e)[r_{ip}(e_p) - r_i(e_p)], \qquad (A.4)$$

Which exploits now the fact that the summation operator runs over all experiences. Using the fact that $\sum_{e\in E_i}[r_{ip}(e_p) - r_i(e_p)] = 0$, the overall priming effect can now be rewritten as:

$$\sigma_i(e_p)[1 - r_{ip}(e_p)] + \sum_{e\in E_i}\sigma_i(e)[r_{ip}(e_p) - r_i(e_p)], \qquad (A.5)$$

Using the same approximation we use for prediction 5.2 this is in turn equal to:

$$\sigma_i(e_p)[1 - r_{ip}(e_p)] + \kappa S_i(e_p)\left[\sigma_i(e_p) - \frac{\sum_{e\in E_i\backslash e_p}\sigma_i(e)}{|E| - 1}\right], \qquad (A.5)$$

Also the new second "interference" term increases in $\sigma_i(e_p)$ and hence in $S_i(e_p)$, confirming that priming e_p boosts beliefs more the more similar this experience is, namely the higher is $\sigma_i(e_p)$.

Consider next point ii), namely experience effects. In the group of subjects for which recall is not primed, either because they have not lived the primed experience or because they are in the control group, beliefs by equation (5) are equal to:

$$\hat{\pi}_{ip} = \theta_i \sum\nolimits_{e\in E_i} \sigma_i(e) r_i(e) \qquad (A.6)$$

This is the standard model without priming in which Prediction 4.1 holds. Thus, having lived more similar experiences boosts estimates provided these are sources of pessimism and sources of optimism are even less similar than sources of pessimism.

Consider now subjects who went through the priming treatment. For these subjects, nothing changes insofar as the dummy for primed recall is included in the regression, which is always the case in Table 2. In this case, the dummy I_p captures the term in square brackets, and the lived experiences dummies capture the effect of experiences in (A.6), which also follows Prediction 4.1.

Proof of Prediction 5.2. We obtain a linear approximation to Equation (5) around the zero similarity point $S_i^*(e) = 0$ for all i. Consider first the "experience effect" term:

$$\frac{\partial}{\partial S_i(u)} \frac{\sum_{e\in E_i} \sigma_i(e) S_i(e)}{\sum_{e\in E_i} S_i(e)}\bigg|_{S=0} = \frac{\partial \sigma_i(u)}{\partial S_i(u)} r_i(u) + \sum\nolimits_{e\in E_i} \sigma_i(e) \frac{\partial r_i(e)}{\partial S_i(u)}\bigg|_{S=0} = \frac{\sigma_i'}{|E_i|},$$

where $\sigma_i' > 0$ is the first derivative of the simulation function at $S = 0$ and experiences are equally retrievable under constant (equal to zero) similarity, $r_i(u) = 1/|E_i|$. The second, interference, term vanishes in aggregate (retrieval probabilities add to one). The linearized experience effect is then:

$$\sum\nolimits_{e\in E_i} \frac{\sigma_i'}{|E_i|} S_i(e) = \sigma_i' \cdot \overline{S}_i(E_i), \qquad (A.7)$$

where $\overline{S}_i(E_i)$ is the average similarity of all lived experiences. Consider next the direct priming term. Deriving with respect to $S_i(e_p)$ we obtain:

$$\frac{\partial}{\partial S_i(e_p)} \frac{\sigma_i(e_p) \sum_{e\in E_i\setminus e_p} S_i(e)}{\sum_{e\in E_i} S_i(e)}\bigg|_{S=0} = \frac{\partial \sigma_i(e_p)}{\partial S_i(e_p)} \frac{\sum_{e\in E_i\setminus e_p} S_i(e)}{\sum_{e\in E_i} S_i(e)} - \frac{\sigma_i(e_p) \sum_{e\in E_i\setminus e_p} S_i(e)}{\left[\sum_{e\in E_i} S_i(e)\right]^2}\bigg|_{S=0} =$$

$$\sigma_i' \frac{|E_i| - 1}{|E_i|} - \sigma_i \frac{|E_i| - 1}{|E_i|^2}, \qquad (A.8)$$

where σ_i is simulation at zero similarity, which is equal to zero, $\sigma_i = 0$. From Assumption 2 and in line with Table 2, we have that $\sigma_i' > 0$. Next, deriving with respect to and $S_i(e)$, $e \neq e_p$ we obtain:

$$\frac{\partial}{\partial S_i(e)} \frac{\sigma_i(e_p) \sum_{e\in E_i\setminus e_p} S_i(e)}{\sum_{e\in E_i} S_i(e)}\bigg|_{S=0} = \sigma_i \frac{S_i(e_p)}{\left[\sum_{e\in E_i} S_i(e)\right]^2}\bigg|_{S=0} = \sigma_i \frac{1}{|E_i|^2}, \qquad (A.9)$$

where again $\sigma_i = 0$. Thus, the direct priming effect is:

$$\grave{\sigma}_i' \cdot S_i(e_p) \equiv \sigma_i' \frac{|E_i| - 1}{|E_i|} \cdot S_i(e_p) > 0. \qquad (A.10)$$

We have thus obtained coefficients $a_1 = \sigma_i' > 0$ and $a_2 = \sigma_i' \frac{|E_i|-1}{|E_i|} > 0$ of Equation (6). To obtain the interference term, it is enough to replace $\Delta r_i(e) = -\kappa S_i(e_p)/[|E_i| - 1]$ in Equation (5) to obtain:

$$-\kappa \cdot S_i(e_p) \cdot \frac{\sum_{e\in E_i\setminus e_p} \sigma_i(e)}{|E_i| - 1},$$

whose linearization with respect to $S_i(e)$, $e \neq e_p$, yields:

$$-\kappa \cdot \sigma_i' \cdot S_i(e_p) \cdot \frac{\sum_{e\in E_i\setminus e_p} S_i(e)}{|E_i| - 1} = -\kappa \cdot \sigma_i' \cdot S_i(e_p) \cdot \overline{S}_i(E_i\setminus e_p),$$

which identifies $a_3 = -\kappa \cdot \sigma_i' < 0$.