



Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Imagining the Future: Memory, Simulation and Beliefs

*Pedro Bordalo, Giovanni Burro, Katherine Coffman, Nicola Gennaioli and
Andrei Shleifer*

Working Paper n. 701

This Version: June, 2023

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
<http://www.igier.unibocconi.it>

The opinions expressed in the working papers are those of the authors alone, and not those of the Institute, which takes non institutional policy position, nor those of CEPR, NBER or Università Bocconi.

Imagining the Future: Memory, Simulation and Beliefs

Pedro Bordalo, Giovanni Burro, Katherine Coffman, Nicola Gennaioli and Andrei Shleifer¹

July 2023

Abstract

How do people form beliefs about novel risks, with which they have little or no experience? Motivated by survey data we collected in 2020, which showed that beliefs about Covid's lethality depended on a range of personal experiences in unrelated domains, 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 yields predictions on how experiences interfere with each other in recall and how non domain-specific experiences 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. Experiences and their measured similarity to the cued event successfully help explain beliefs, with patterns consistent with our theory. Our approach offers a new, structured way to study and jointly account for systematic biases and substantial belief heterogeneity.

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 and IGIER. 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, and Jesse Shapiro 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). 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?

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 yields many predictions on how experiences interfere with each other in recall, and how NDS experiences create biases in beliefs based on their similarity to the event. We test these predictions using data from the Covid survey and in a pre-registered, primed-recall experiment where the cued event is the risk of a cyberattack. Personal experiences and their measured similarity to the cued event have strong explanatory power for beliefs, consistent with selective memory as a key mechanism.

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 of disagreement appear puzzling. First, people who estimate that a higher share of Americans have red hair were more pessimistic about Covid. This fact points to the importance of cognitive factors, such as the ease of imagining rare events, as opposed to information, motivations, or preferences, in shaping beliefs. Second, there was a striking age gradient: older people were

substantially *less* pessimistic about Covid’s lethality than younger people, both when estimating lethality for themselves and others, contrary to reality. Third, people were more pessimistic if they had experienced more non-Covid health adversities, such as their own or a family member’s recent hospitalization for non-Covid reasons; again, this holds true for lethality estimates for themselves or others. These two facts point to the importance of memory. When thinking about Covid, people may associatively recall severe diseases beyond Covid and use those to simulate Covid lethality (NDS boosting estimates). It may be harder for the elderly to simulate Covid deaths due to strong memory competition from many other experiences they lived and survived (DS experiences being neglected), compared to the young for whom Covid is more formative.

In Section 3 we present our model. People recall either statistical information about the risk, which they may have seen in the media, or their own past experiences. Recall of experiences is driven by three well-known regularities: similarity, frequency, and interference (Kahana 2012, Bordalo et al. 2023). Recalled experiences, which may or may not concern the risk being assessed, are then used to simulate that risk. Simulation is a form of reasoning by analogy well documented in psychology (Dougherty et al. 1997; Brown et al. 2000, Schacter, Addis, and Bruckner 2008, Biderman, Bakkour, and Shohamy 2020)². It recruits non-domain-specific (NDS) experiences to imagine the target event and to shape beliefs, especially if the experience is similar to that event (Kahneman and Tversky 1981). Similarity plays a dual role: it fosters recall and simulation.

In the model, having lived a specific experience entails a trade-off between simulation and interference: the experience potentially creates material for simulation, which boosts probability estimates, but interferes with retrieval of other experiences, which dampens estimates if the latter provide better simulation material. This trade-off yields two comparative statics.

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

² Economists have previously used the concept of simulation in modeling how people discount the future (Becker and Mulligan 1997, Gabaix and Laibson 2022). These papers do not connect memory and simulation of future events. Ashraf et al. (2022) present evidence that pictorial imagery helps potential entrepreneurs imagine future outcomes.

are even less similar may dampen assessments. For instance, experiences of Covid deaths should boost Covid pessimism more than non-Covid health adversities, which are less similar. In turn, non-Health Adversities, being least similar, may reduce the perception of Covid risks by interfering with recall of the latter. Both domain specific and NDS experiences shape beliefs, based on similarity.

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 be forgotten.

These roles of similarity and interference in shaping the effect of NDS experiences represent a major 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.

This framework also suggests a new strategy for running surveys that elicit beliefs: measure many experiences, including NDS ones, and measure perceptions of similarity between these experiences and the target event. We apply this strategy in Section 4 to our Covid surveys, and in Section 5 to a survey on the risk of a significant cyberattack. The latter survey exogenously primes respondents to recall DS and NDS experiences and measures individual perceptions of similarity, providing a targeted test of the model.

The results reveal the central role of selective recall and simulation. In both surveys we find the hierarchy of experience effects: more similar experiences boost estimates compared to less similar ones, and NDS experiences matter. The same hierarchy arises with priming: inducing a respondent to recall a more similar experience boosts estimates more, even if the experience is NDS. The fact that priming a DS experience boosts estimates shows that DS experiences may fail to come to mind. Crucially, the tests also show that such failure reflects, to some extent, interference. In the Covid survey, different people react differently to the same local pandemic conditions based

on their NDS experiences. In the cyberattack experiment, priming recall of an experience reduces the explanatory power of non-primed experiences.

A vast body of social science research has documented the effect of past experiences on beliefs and decisions (e.g., Weinstein 1989). Economic research links individual experiences to insurance demand (Kuhnreuther 1978) and IPO investing (Kaustia and Knüpfer 2008), political experiences to the demand for redistribution (Alesina and Fuchs-Schündeln 2007), and macroeconomic experiences to stock market participation and inflation expectations (Malmendier and Nagel 2011, 2016). These experience effects are mechanical and domain specific (Malmendier 2021): an encoded experience is always recalled, possibly with some time decay, and affects beliefs about its domain directly. Our evidence instead shows that DS experiences may fail to be retrieved, and beliefs are shaped endogenously by their *recall* and *use*. Using NDS experiences yields disagreement despite common information; cueing own NDS experiences, e.g. through advertising or persuasion, leads to instability.

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. In our model, memory explains different biases based on measurable differences in databases and similarity. Our approach can also help explain why seemingly irrelevant cues may affect choice, as in framing effects. The cues can trigger retrieval of NDS experiences, affecting simulation of the target event.

Versions of our model have been already applied to several domains. Conlon and Patel (2022) use it to understand beliefs about future career and college major choices, Jiang et al. (2023) to study beliefs about future stock returns, Exley et al. (2022) to describe beliefs about gender and prosociality, Graeber et al (2022) to compare memories of stories and statistics, Cenzon (2023) to study how wholly idiosyncratic experiences affect macroeconomic expectations, and Colonnelli et

al. (2023) to explore how different messages shape public support for bailout of large firms. Across these different contexts, explanatory power comes from spontaneous or cued retrieval of NDS experiences, which creates systematic disagreement and instability in beliefs.

2. Motivating Evidence: Puzzles in Beliefs about Covid

Our model is motivated by three puzzling facts on beliefs about Covid we discovered 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, 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.

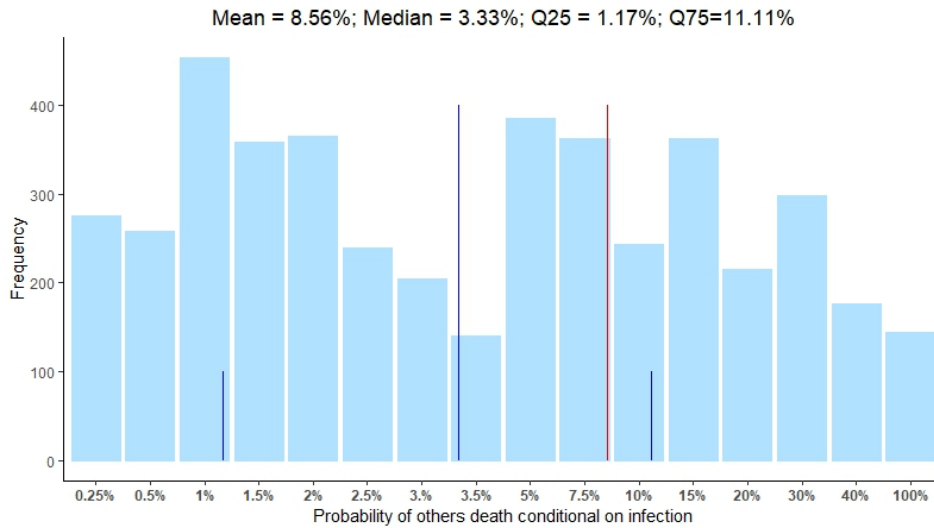


Figure 1

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).³ 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.

³ Similar estimates appear in contemporaneous papers, see Covid CDC (2020), Russell et al (2020), Modi et al (2021).

Where do disagreement and mean overestimation come from? When first looking at the data, we found three striking 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 highly 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, at 4.7% (Levin et al. 2020). 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.




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.

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%). It is puzzling that such an idiosyncratic health shock, entirely unrelated to Covid, so strongly influences perceptions of Covid

risks for the general population. Third, panel C shows that respondents estimating a greater share of red haired Americans are much 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 , 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 domain-specific 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 risks 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.⁴ The roles of health experiences and age point to memory. When thinking of Covid, some people associatively retrieve their own or a loved one's recent illness, prompting pessimism about the new disease. The many lived experiences of the elderly – including surviving other adversities – make it harder for them to imagine Covid risks, compared to the young for whom Covid faces relatively less interference from pre-existing experiences. Disagreement may stem from selective retrieval of non-Covid experiences.

We present a memory-based model of beliefs that accounts for these puzzles and offers new predictions. It explains the role of experiences we measured in survey waves 2 and 3 for beliefs about Covid. Its 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 unifies, and sheds new light on, priming and experience effects.

⁴ Politics has little explanatory power for Covid risk perceptions in our data (see Tables C4 and C5 in Appendix C).

3. The model

A decision maker (DM) estimates the probability of an event H (vs. alternative \bar{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.

A symmetric function $S: E \times E \rightarrow [0,1]$ 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 it is to a death from homicide; a death from Covid is 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,

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

$S(A, B)$ is symmetric and increases 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, which we denote by $H = CD$. Based on (1), denote by $S(e) \equiv S(e, H)$ the similarity between experience e and H .

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, $S(e)$:*

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

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 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 CD . Interference is well-established in memory research (e.g., Jenkins and Dallenbach 1924; McGeoch 1932; Underwood 1957).⁵ 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.

Bordalo et al. (2023) present a theory in which 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. 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). In our context, recalling a death from pneumonia may boost the DM’s ability to imagine death from Covid, making

⁵ 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).

her more pessimistic, even though the experience is not in CD . In cognitive science, the ease of simulation increases with the similarity between the retrieved memory and the target event (Kahneman and Tversky 1981, Woltz and Gardner 2015). 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 $S(e) \geq S(u)$.*

Simulation entails using the intrinsic features of e to imagine H . The higher the number of features the two events have in common, the stronger is simulation. It is easier to imagine a Covid death based on pneumonia than on hunger, because the former is more similar to the target.⁶

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:

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

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:

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

which combines the statistical “truth” π with the experience-based estimate $\hat{\pi}_E$.⁷ The target event is overestimated on average when $\hat{\pi}_E > \pi$ and underestimated otherwise.⁸

A key driver of belief distortions in our model is non-domain specific, NDS, experiences. Suppose instead that only domain-specific, DS, experiences are retrieved and used (e.g., Covid

⁶ 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.

⁷ We assume that, when forming beliefs about a target event, the DM does not think about the alternative hypothesis of it not happening. This is consistent with both our survey and our experiment, in which respondents assess “Covid death” and “cyberattack” but the alternative hypothesis is not mentioned. Results are similar if we relax this assumption.

⁸ 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.

infections and fatalities). In particular, only DS events are recalled, $e \in H \cup \overline{H}$, and only recollections in H are used to simulate H itself. This occurs if similarity and simulation are “narrow”, namely $S(e, H \cup \overline{H}) = \sigma(e, H) = 1$ and zero otherwise. In this case, (4) yields the frequentist assessment $\hat{\pi} = |H|/|H \cup \overline{H}|$. As long as the DS database is unbiased (contains the true frequency of H and \overline{H}), beliefs are unbiased, $\hat{\pi} = \pi$. 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 due to forgetting of them, retrieval of NDS ones, and simulation based on the latter. Will over-or under-estimation prevail? And which kind of disagreement will emerge?

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. Proposition 1 shows a source of heterogeneity in this bias: a person's reliance on experience θ . In our Covid survey, the "red hair" variable can be interpreted as proxy for θ . Consider now the effect of specific experiences.

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. Point 1 says that one key driver 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. We later test this prediction.

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. We later test this prediction as well.

3.3 Broader Implications and Roadmap

The pillar of our analysis is that selective memory, specifically recall and simulation, is key for understanding beliefs in many domains. This 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 an entirely different perspective. Novel events are judged using pre-existing, albeit unrelated, experiences and these experiences interfere with the integration of domain-specific data.

Our approach builds on both the limitations of human memory and the creativity of the human mind. People associatively sample their database, including experiences from different domains, and then use these experiences as material for imagination. 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.

Memory relates to disagreement in a structured way, based on three drivers. 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 memory mechanisms empirically, both the database and similarity 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, enabling 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 a survey experiment on beliefs about severe cyberattacks, designed as an explicit 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.

Before discussing the tests, we describe the richer 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 respondents whether they have had Covid. To measure the severity of local pandemic conditions we use publicly available state-level data to construct 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 pandemic recency, the days that have passed since that peak.⁹ Table B.1 in Online Appendix C describes these covariates.

4.1 Similarity and Experience Effects in Beliefs about Covid

⁹ Accessible from the New York Times counts, <https://www.nytimes.com/interactive/2021/us/covid-cases.html>.

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 *CD* 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 *CD*. 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 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.¹⁰ Proposition 1 implies that personal Covid experiences should boost Covid pessimism because they are more similar to *CD* 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 *CD*.¹¹ 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 *CD*. If these experiences come to mind, they do not help simulating *CD* 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*

¹⁰ We include *untimely death of a loved one* for it reflects enduring personal hardship, creating a non-health adversity. Omitting this experience, the non-health adversities index retains the negative sign with a p value of 0.06 (see Table C4).

¹¹ 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.

Health Adversities index.¹² 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.¹³ 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. 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	
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)

¹² 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.

¹³ 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”.

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.¹⁴ In a Bayesian world, by contrast, surviving Covid should, if anything, 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. This is due to interference: having gone through a bumpy life makes it easier to retrieve risks different from Covid. This reduces simulation of Covid deaths, lowering estimates.

Consistent with the model, both DS and NDS experiences matter based on similarity. 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

¹⁴ 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.

Covid deaths at peak in the data is 15700. The effect of Health Adversities is about half this effect, so it is also large. 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. 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, for whom personal preferences are obviously irrelevant. Second, endogeneity may also affect 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, 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.*

This is a direct consequence of Proposition 2.2. A person who experienced a source of pessimism, such as a non-Covid health adversity, should react less to a given severity of the pandemic 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 the severity of local pandemic conditions, as measured by *level*. In Figure 3, we assess the impact of *level* on beliefs, comparing respondents in the bottom tercile of local pandemic 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.

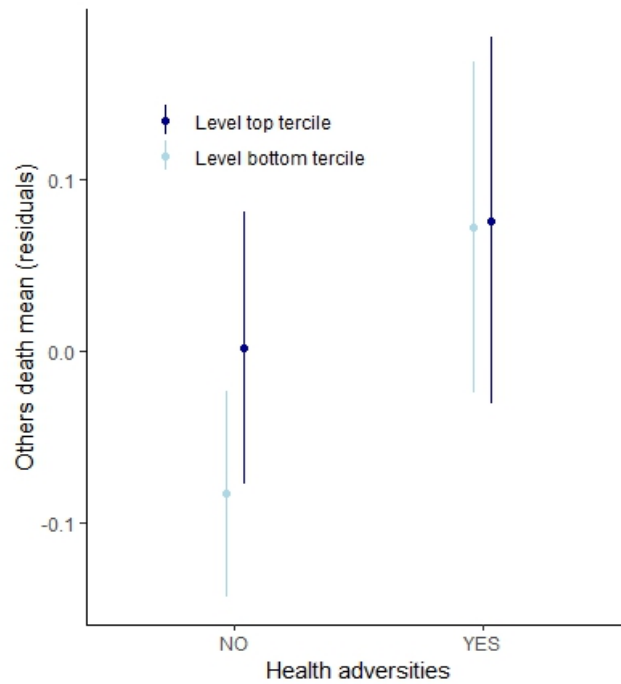


Figure 3.

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 explains the striking age effect (Prediction 4.3i). The database of the elderly is populated overwhelmingly with 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).¹⁵ 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). We argued that 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 disproportionately pessimistic if they experience sources of pessimism, and disproportionately optimistic if they experience sources of optimism. This prediction links the red hair proxy to recall and use of lived experiences, allowing to rule out its interpretation as a mechanical tendency toward insensitivity, due to noise or other mechanisms (Kahneman and Tversky 1979, Enke and Graeber 2023, Khaw et al 2020, Abdellaoui et al 2011).

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.

¹⁵ 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.

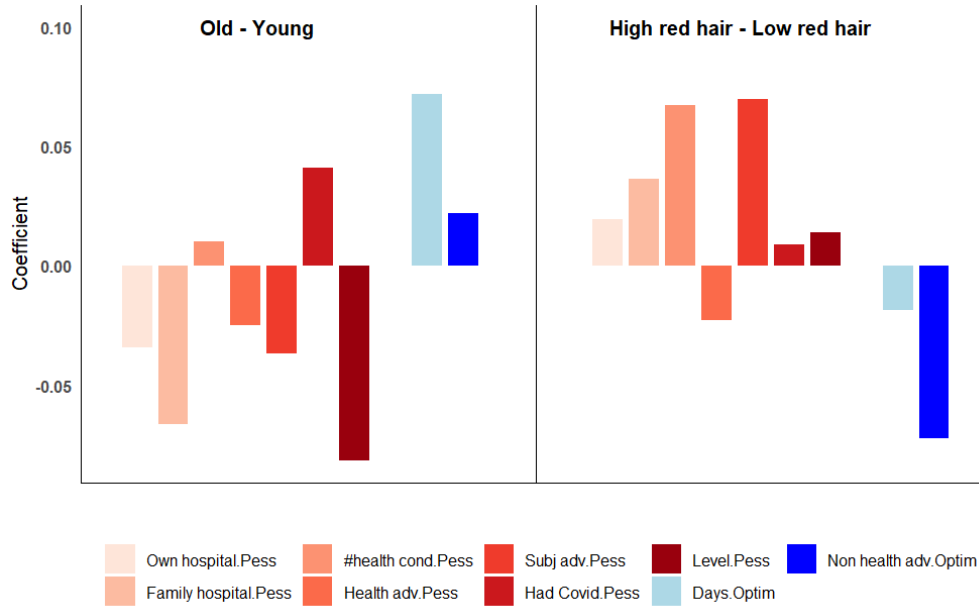


Figure 4

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.¹⁶ 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.¹⁷ 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.¹⁸

¹⁶ 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%.

¹⁷ 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.

¹⁸ 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.

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 before developing the model, so we do not measure or control for its key mechanisms. 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. Studying beliefs about cyberattacks also allows us to assess whether the role of experiences documented in the Covid data extends to other domains.

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 all reported experiences and the event H . This measure allows us to test for the role of similarity in simulation and in spontaneous recall. 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*). Higher assessments capture more reliance on real or fictional snow 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) *finance*: personal experience of struggling with finances, and iii) *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” which we denote by SC . We chose adverse experiences that vary in their perceived similarity to SC : we expect that *ID theft* is on average judged more similar to SC than *finance* and *loss*, but we also expect individual-level variation in perceived similarity. *ID theft* is also the most domain-specific experience, as there have been no severe nation-wide cyberattacks to date.

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 SC . 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. $r_i(e_p) < 1$.

The second effect of priming is to influence recall of non-primed experiences, $e \neq e_p$. When thinking about SC , 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 recalling $e_p = \textit{finance}$ the DM may spontaneously think of financial losses from the cyberattack, while instead finding it harder to retrieve different experiences such as $e = \textit{September 11}^{\text{th}}$. 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 less likely that they spontaneously retrieve dissimilar items, 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 similar to itself (e.g. “yogurt”), while inhibiting retrieval of items that are different. 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.¹⁹ Plugging $S_{ip}(e_p)$ in Equation (2) yields the interference effect 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 $e \neq e_p$. 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. The rest of the model is then 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.²⁰ 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:

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

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 SC , $S_i(e)$ is higher, then adding it to the database raises belief $\hat{\pi}_{ip}$. This occurs because similarity boosts simulation (Assumption 2).

¹⁹ 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} is 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.

²⁰ 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 .

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 SC , $S_i(e_p)$ 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 SC , 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.

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.

In sum, in our model priming is tightly connected to memory and exerts two conflicting forces. It creates material for simulation from e_p , but also interferes with recall of other experiences, dampening simulation from $e \neq e_p$. This structure 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 $\sum_{e \in E_i} \Delta r_i(e) = \sum_{e \in E_i \setminus e_p} r_{ip}(e) > 0$. This implies that, regardless of the primed experience, priming tends to (but does not necessarily) increase estimates relative to control.

5.3 Similarity, Priming and Experience Effects: Empirical Tests

We first report the distribution of estimates of the probability of a cyberattack. Unlike in our Covid survey, there is no objective benchmark, but the data also shows large heterogeneity in beliefs.

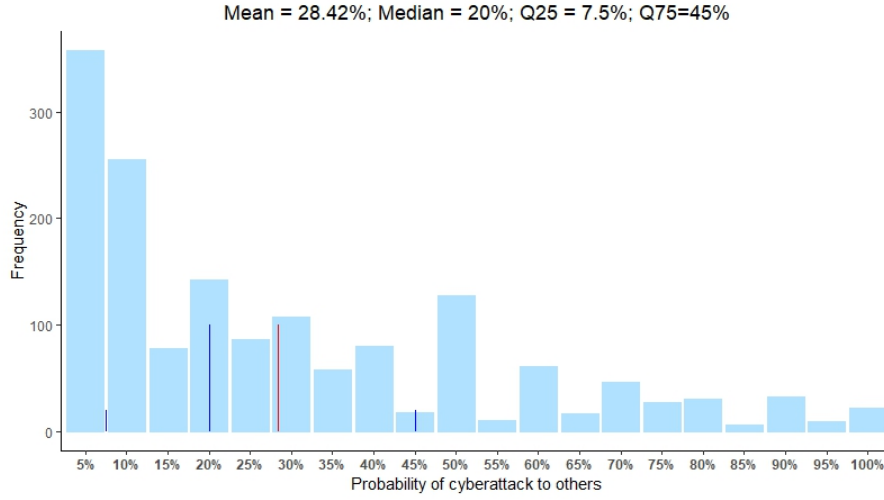


Figure 5

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 $e_p = ID\ theft, finance, loss$, 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:²¹

Prediction 5.1 Similarity in Priming and Experience Effects. *Living an experience e with higher measured similarity $S_i(e)$ with SC 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,

²¹ 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 also allows us to ask whether experiences with higher average similarity to SC are associated with stronger priming *and* experience effects, in line with Prediction 5.1. Finally, 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 dummy allows us to see whether subjects who perceive a given e_p as less similar to SC exhibit weaker priming effects compared to other participants.

We use data from high quality responses, defined as those who answer our numeracy questions correctly.²² 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.²³

²² 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 *finance*, and 32% of the *loss* groups, reflecting that a greater share of our respondents have experienced financial struggles relative to the other experiences.

²³ 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.

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

	OLS Predicting Index of Cyberattack Estimates		
	High Quality Only		TOT
	ITT 1	2	
<i>ID theft</i> prime	0.15** (0.061)	0.22*** (0.069)	0.19** (0.086)
<i>Finance</i> prime	0.14** (0.062)	0.12* (0.066)	0.22*** (0.083)
<i>Loss</i> prime	0.057 (0.061)	0.10 (0.071)	0.22** (0.092)
Below Median Similarity x <i>ID theft</i> prime			0.056 (0.10)
Below Median Similarity x <i>Finance</i> prime			-0.17* (0.094)
Below Median Similarity x <i>Loss</i> prime			-0.21** (0.11)
Had <i>ID theft</i>		0.40*** (0.064)	0.40*** (0.064)
Had <i>finance</i>		0.30*** (0.084)	0.29*** (0.084)
Had <i>loss</i>		0.026 (0.060)	0.023 (0.060)
<i>Snow</i>	0.16*** (0.023)	0.19*** (0.025)	0.18*** (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 *finance* boosts estimates. Consistent with selective recall and simulation, these experiences help

the respondent to imagine *SC*, but with some probability they are not retrieved if not primed. There is instead no priming effect for the *loss* prime. This is consistent with the role of similarity in Prediction 5.1i): on average, *ID theft* and *finance* 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 *loss* (mean of 1.95). The coefficients of *ID theft* and *finance* are very similar, which in this ITT approach may be due to the lower frequency of people who have lived *ID theft* compared to *finance* (70% vs 88%).

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 *finance*, 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 *finance* and even more so compared to the least similar *loss*.

Column 3 further highlights the role of similarity, showing that – for *finance* and *loss* – subjects who regard the given primed experience, e_p , as less similar to *SC* 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 *loss*. The model suggests that these similarity-driven effects are concentrated among the *finance* and *loss* primes at least in part because these experience are unlikely to be recalled when thinking about *SC* 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 next consider the role of interference using all measured experiences, not only the primed ones, to proxy for the database E_i . We develop a proxy for the interference term in Equation (5). As shown above, priming e_p causes – for the purpose of recall – an increase in similarity $S_{ip}(e_p) - S_i(e_p) = \kappa \cdot S_i(e_p)$, capturing the extent to which recall of non-primed experiences is dampened in Equation (2). Assuming that interference is equal across all non-primed experiences $e \neq e_p$, this implies that the recall probability of each experience falls by $\Delta r_i(e) = -\kappa S_i(e_p) / [|E_i| - 1]$. Equation (5) then yields the following estimating equation.

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

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

where $\bar{S}_i(E_i)$ is the average similarity of all lived experiences, $S_i(e_p)$ the similarity of the prime, $\bar{S}_i(E_i \setminus e_p)$ 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 SC 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 SC 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 SC exhibit a muted effect from lived, non-primed experiences, $a_3 < 0$. Once an experience highly similar to SC 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 $\bar{S}_i(E_i \setminus e_p)$ 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. 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 <i>Snow</i>		Above Median <i>Snow</i>	
	1	2	3	4	5	6
$\bar{S}_i(E_i)$, Total Similarity of Lived Experiences	0.19*** (0.028)	0.20*** (0.028)	0.14*** (0.039)	0.14*** (0.039)	0.24*** (0.040)	0.25*** (0.040)

$S_i(e_p)$, Similarity of Primed Experience	0.13*** (0.028)	0.19*** (0.034)	0.17*** (0.039)	0.20*** (0.047)	0.10** (0.040)	0.19*** (0.049)
$\bar{S}_i(E_i) \times S_i(e_p)$		-0.059*** (0.020)		-0.039 (0.030)		-0.080*** (0.027)
<i>Snow</i>	0.18*** (0.025)	0.18*** (0.025)	0.36*** (0.090)	0.36*** (0.090)	0.088* (0.050)	0.093* (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 1 ft 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).

In sum, the results of our experiment are consistent with both our findings from the Covid survey and the key mechanism of our model. Beliefs depend on both DS and NDS experiences because retrieval is imperfect (priming matters) and NDS events are used in simulation (similarity matters). Consistent with the model, the similarity hierarchy of priming and experience effects is congruent, and there is interference between the two. Interference matters for belief formation.

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, 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 yields many new predictions. We test and confirm these predictions 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, but consistent with evidence 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 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 generally, memory is a key input into all of our cognitive activities, so its effect can be far-reaching. 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 that memory is crucial for such assessments (Shadlen and Shohamy 2016). This perspective creates exciting opportunities to explain economic behaviour 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.", *The 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." Mimeo, Harvard University.
- 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.
- Burszty, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2020. "Misinformation During a Pandemic." *NBER Working Paper* 27417.
- Colonnelli, Emanuele, Niels Gormsen and Timothy McQuade. 2023. "Selfish Corporations," NBER w30576.
- 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." *MMWR Morb Mortal Wkly Rep*, 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.
- 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 J., Sean Polyn, and Michael J. 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(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, 657-661.

- Schacter, Daniel, Donna Addis, Randy Bruckner. 2008. "Episodic Simulation of Future Events." *Annals N.Y. Academy of Science*, 1124, 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: 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 N., 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 $H \cup \bar{H}$ are recalled, $\hat{\pi}_E = |H|/|H \cup \bar{H}| = \pi$. If experiences other than H can be used to simulate H by a factor $\tilde{\sigma} > 0$ and if experiences outside $H \cup \bar{H}$ are according to similarity $\tilde{S} > 0$, then using Equation (3) we have:

$$\hat{\pi}_E = \frac{|H| + \tilde{\sigma}[|\bar{H}| + \tilde{S}|E \setminus H \cup \bar{H}|]}{|H \cup \bar{H}| + \tilde{S}|E \setminus H \cup \bar{H}|},$$

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

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

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):

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

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

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

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

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

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:

$$\begin{aligned} \frac{\partial \hat{\pi}_{E \cup |e| \cup |e'|}}{\partial |e| \partial |e'|} &\propto \\ &[\sigma(e) - \sigma(e')]S(e')[S(e)|e| + \sigma(e')S(e')|e'| + \mathbb{E}_E(S)|E|] - \\ &2S(e')[[\sigma(e) - \sigma(e')]S(e')|e'| + [\sigma(e)\mathbb{E}_E(S) - \mathbb{E}_E(\sigma S)]|E|], \end{aligned} \quad (\text{A.2})$$

Which after some simplification yields:

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

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 that 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 $\sigma_i(e_p)\Delta r_i(e_p) + \sum_{e \in E_i \setminus e_p} \sigma_i(e)\Delta r_i(e)$, which can be rewritten as:

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

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:

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

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

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

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:

$$\hat{\pi}_{ip} = \theta_i \sum_{e \in E_i} \sigma_i(e) r_i(e) \quad (\text{A.6})$$

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:

$$\left. \frac{\partial}{\partial S_i(u)} \frac{\sum_{e \in E_i} \sigma_i(e) S_i(e)}{\sum_{e \in E_i} S_i(e)} \right|_{S=0} = \frac{\partial \sigma_i(u)}{\partial S_i(u)} r_i(u) + \sum_{e \in E_i} \sigma_i(e) \left. \frac{\partial r_i(e)}{\partial S_i(u)} \right|_{S=0} = \frac{\sigma'_i}{|E_i|},$$

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:

$$\sum_{e \in E_i} \frac{\sigma'_i}{|E_i|} S_i(e) = \sigma'_i \cdot \bar{S}_i(E_i), \quad (\text{A.7})$$

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

$$\left. \frac{\partial}{\partial S_i(e_p)} \frac{\sigma_i(e_p) \sum_{e \in E_i \setminus e_p} S_i(e)}{\sum_{e \in E_i} S_i(e)} \right|_{S=0} = \frac{\partial \sigma_i(e_p)}{\partial S_i(e_p)} \frac{\sum_{e \in E_i \setminus e_p} S_i(e)}{\sum_{e \in E_i} S_i(e)} - \frac{\sigma_i(e_p) \sum_{e \in E_i \setminus e_p} S_i(e)}{[\sum_{e \in E_i} S_i(e)]^2} \Big|_{S=0} = \sigma'_i \frac{|E_i| - 1}{|E_i|} - \sigma_i \frac{|E_i| - 1}{|E_i|^2}, \quad (\text{A.8})$$

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:

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

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

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

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:

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

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

ONLINE APPENDIX

Appendix B. The Covid Surveys

To assess risk perceptions during the Covid-19 pandemic, we conducted a survey of a diverse sample of over 1,500 Americans. The survey asked an array of questions related to beliefs, preferences and behavioral responses, as well as sociodemographic characteristics. We do not incentivize participants for accuracy given the large uncertainty surrounding the data on many of these issues. We first describe the structure and implementation of the first survey we ran, in May 2020, and then discuss the changes made in Waves 2 and 3. The survey instruments can be found at the conclusion of this section.

WAVE 1 SURVEY

To reach a diverse sample of Americans, we partnered with Qualtrics, who handled the recruitment and compensation of our participants. We specified a desired 1,500 respondents, who met the following quotas:

- Gender: Female (~50%); Male (~50%)
- Age: 18-34 (~25%); 35-49 (~25%); 50 - 69 (~30%); 70 and older (~20%)
- Household Income: <\$50K (~35%); \$50K-100K (~35%); >100K (~30%)
- Region: Midwest (~20%); Northeast (~20%); South (~40%); West (~20%)
- Race: White (~66%); Black (~12%); Latinx (~12%); Asian (~10%)

To guarantee representation in line with these quotas, the 5 demographic questions requesting this information were presented immediately following the consent form, allowing for screening out of participants as quotas were met. In addition, any participant who indicated they were younger than 18 years old or resided outside of the United States was screened out.

We also wanted to guarantee a minimum level of quality and thoughtfulness of participant responses. Immediately following the demographic screener questions, participants were told: “We care about the quality of our survey data and hope to receive the most accurate measures of your opinions. It is important to us that you provide thoughtful, careful answers to each question in the survey. Do you commit to providing your thoughtful and careful answers to the questions in this

survey?” Participants had to select “I commit to providing thoughtful and careful answers” from 3 possible options in order to continue in the survey.

Finally, we wanted to familiarize participants with the question format they would see on much of the survey, while providing a further screen of their thoughtfulness and quality. Because objective likelihoods of suffering particular health consequences related to Covid-19 are in some cases quite small, it could be difficult for a typical participant to express their beliefs in a probability or percentage format. More generally, individuals often have difficulty interpreting probabilities, particularly in more abstract contexts. Gigerenzer and Hoffrage (1995) suggest that presenting or eliciting frequencies, rather than probabilities, improves participant understanding.

To address these concerns, we asked questions in terms of frequencies, but also began by familiarizing participants with the question format. We told respondents: “Many of the questions on this survey will ask you to make your best estimate as to how many out of 1,000 Americans will experience different events or have different features. To give you some practice and get you used to thinking in these terms, we have a few example questions for you to work through.”

For the first example, participants were told that, according to the United States Census, approximately 20 out of 1,000 Americans live in Massachusetts, and that this is equivalent to approximately 2% or 2 out of every 100. We then asked them, using this estimate, to tell us how many out of 5,000 Americans live in Massachusetts. Participants had to provide an answer of 100 (i.e. 2% of 5,000) in order to continue in the survey.

For the second example, participants were told that they would estimate the size of a group of Americans with a certain attribute. In particular, they were asked to provide their guess of 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). Only participants who estimated that fewer than 1,000 out of 1,000 Americans had red hair could continue in the survey. Participants also had to provide consistent answers: their answer to the “out of 10,000” question had to be 10 times their answer to the “out of 1,000” question in order to continue in the survey.

Following their successful completion of this question, we informed participants of what their red hair estimate implied both as a percentage and in terms of how many Americans out of 100, out of 1,000, and out of 100,000 would have red hair. We also provided an accurate estimate as a useful reference point: roughly 15 out of 1,000 Americans are estimated to have red hair, which we described to them as 1.5%, 1.5 out of 100, 15 out of 1,000, or 1,500 out of 100,000.

After completing these questions in line with our specified quality conditions, participants continued to our questions of interest. Qualtrics did not provide us with data on the participants who were screened out, nor did they inform us of the rate at which participants were screened out.

Participants completed several blocks of questions: Covid-19 Related Health Risks for People Like Self, Other Health Risks for People Like Self, Economic and Other Risks, Covid-19 Related Health Risks for Others, Demographics, and Preferences and Behavior. We asked about many sources of risk to assess whether the salience of Covid-19 health risks influences how other health and economic risks are judged.

A. Covid-19 Related Health Risks for People Like Self

In this block, we first ask participants to think about 1,000 people “very similar to you (i.e., in terms of age, gender, race socioeconomic status, zip code, health status, etc.)”. We then ask “of these 1,000 people, how many do you believe will contract Covid-19 in the next 9 weeks?” We provide a time-frame to make the question more concrete, and we choose 9 weeks because we anticipate running multiple waves of this survey over time, approximately 9 weeks apart. We do not bound participants’ answers.

Because this is the first risk elicitation question of this form, we contextualize this answer for all participants. In particular, after they provide their response, they are taken to a new survey page that informs them about the answer they just gave. Suppose they answered that they believe 300 of 1,000 people similar to them will contract Covid-19 in the next 9 weeks. The survey then repeats to them: “Just to clarify, by entering 300 for the question on the previous page, you are indicating that you believe 300 out of 1,000 people very similar to you will contract Covid-19 in the next 9 weeks. This is equivalent to 30%.” Each participant is then asked if they would like to revise their answer, and if they indicate that they would, they have the opportunity to provide a new answer. In our analysis, we replace initial estimates with revised estimates for all participants who indicated they wished to revise their answer.

This block on Covid-19 related health risks for self includes two other risk assessment questions. Each asks people to consider 1,000 people very similar to them *who contract Covid-19 in the next 9 weeks*. They are then asked to estimate how many of these 1,000 people very similar to them who contract Covid-19 will require hospitalization. They are also asked to estimate how many of 1,000 people very similar to them who contract Covid-19 will die. The questions about hospitalization and death due to Covid-19 are both conditional on contracting Covid-19. These questions attempt to

isolate beliefs about potential health consequences due to Covid-19 from beliefs about its prevalence or contagiousness.

B. Other Health Risks for People Like Self

We are interested in understanding how perceptions of Covid-19 related health risks compare to and interact with beliefs about other serious health risks faced by this same population. In this next block of questions, we adapt a similar question format to assessing other health risks. For each of the questions, participants are again prompted to consider 1,000 people “very similar to you (i.e., in terms of age, gender, race socioeconomic status, zip code, health status, etc.)”. They are asked to estimate, out of those 1,000, how many will: (i) require hospitalization for a reason other than Covid-19 in the next 5 years, (ii) die for a reason other than Covid-19 in the next 5 years, (iii) have a heart attack in the next 5 years, and (iv) develop cancer in the next 5 years.

C. Economic Risks and Other Threats

We would also like to understand how participants perceive the economic risks surrounding the Covid-19 pandemic. Because these questions do not easily lend themselves to the “out of 1,000” format used for the health questions, we use the Likert-scale. For four different economic outcomes, we ask participants to assess the likelihood of this outcome on a 1 – 7 scale, where 1 indicates extremely unlikely and 7 indicates extremely likely.

We present two pairs of questions, the first related to the stock market and the second related to the unemployment rate. Within each pair, we present both a favourable and unfavourable outcome. For the stock market the two outcomes are: (i) the U.S. stock market drops by 10% or more in the next 9 weeks, (ii) the U.S. stock market grows by 10% or more in the next 9 weeks. For the unemployment rate the two outcomes are: (i) the U.S. unemployment rate reaches 20% or more in the next 9 weeks, and (ii) the U.S. unemployment rate falls below 5% in the next 9 weeks. By eliciting beliefs about good and bad outcomes we can assess not only general optimism or pessimism, but also perceived tail uncertainty.

D. Covid-19 Related Health Risks for Others

Participants’ assessments of their own personal risk of dying from Covid-19 likely depend on their beliefs about the relative importance of different risk factors. We assess how participants believe the chances of *dying* from Covid-19 vary for different demographic groups. For the sake of

simplicity, respondent time, and statistical power, we focus on three easy-to-describe demographic characteristics: age, race, and gender.

We craft the questions to parallel those from the first block of the survey, assessing Covid-19 death risks for people like the respondents themselves. This time, we ask participants to consider “1,000 people in each of the following [AGE/RACE/GENDER] categories who contract Covid-19 in the next 9 weeks.” We ask them, within each category, to assess how many of the 1,000 Americans who contract Covid-19 in the next 9 weeks will pass away due to Covid-19. For the age category, participants make a forecast for 1,000 Americans under 40 years old, for 1,000 Americans between the ages of 40 – 69 years old, and for 1,000 Americans ages 70 and older. For the race category, participants make a forecast for 1,000 white Americans, for 1,000 Black Americans, for 1,000 Asian Americans, and for 1,000 Latinx Americans. For the gender category, participants make a forecast for 1,000 American men and for 1,000 American women.

E. Sociodemographic Characteristics

Recall that at the beginning of the survey, all participants are asked to report: year of birth, gender, race (White, Black, Asian, Latinx, check all that apply), approximate annual household income (choose from buckets of \$25,000 increments), and region of the country (Northeast, South, Midwest, West). These questions appear as the very first five survey questions, so that Qualtrics can use them as screener questions in order to guarantee a stratified sample.

We also ask non-required sociodemographic questions at the end of the survey: state of residence, whether their current place of residence is best described as urban, suburban, or rural, their educational attainment, whether they have been diagnosed with diabetes, heart disease, lung disease, hypertension, obesity, cancer, or another serious immunocompromising condition, whether they have been hospitalized for non-Covid-19 related reasons within the last year, whether a member of their family has been hospitalized for non-Covid-19 related reasons within the last year, and whether they have been unemployed anytime over the last 9 weeks.

F. Preferences and Behavior

Finally, we ask participants about their behavioral responses to the Covid-19 pandemic, and about their preferences regarding policy responses. We ask them how soon they believe “stay at home” measures should be lifted, and whether they would resume their normal activities if stay at home measures were lifted today. We ask about avoidance of medical care, specifically, how reluctant they would be to go to the emergency room today if they or someone in their family had an urgent

medical issue, and whether they have avoided filling prescriptions, doctor's appointments, or other forms of medical care in the last few weeks. We then ask them approximately how many times per week over the last few weeks they have left their home to shop, do errands, socialize, etc. (specifically excluding work or exercise). Finally, we ask them, in their opinion, how likely is a significant resurgence of Covid-19 in the fall/winter of 2020.

G. Treatment Assignment and Order

We were also interested in assessing whether the salience of a certain demographic categorization (age, race, or gender) influenced individual perceptions of Covid-19 risks about oneself. For this reason we randomly assigned each participant to one of four treatments that tweaks the order of questions so that the subject is asked to assess Covid-19 risks for certain demographic groups before answering the Covid-19 Related Health Risks for People Like Self.

Specifically, in the control condition the order is exactly as described above, and we randomly assign, at the participant level, the age, race, and gender questions within the Covid-19 Related Health Risks for Others. In the other three treatments, we extract one of the three questions about others – either the age question, the race question, or the gender question – and move it to the front of the survey, immediately preceding the Covid-19 Related Health Risks for People Like Self block. The idea is to prime participants to think about risks in terms of age, race, or gender, before thinking about risks for people like themselves. For participants assigned to one of these three treatments, the remaining 2 questions about others are kept in their original place, in a random order, within the Covid-19 Related Health Risks for Others block later in the survey.

H. Implementation

Qualtrics obtained 1,526 responses to our survey between May 6 and May 13, 2020. Of those 1,526, we drop 4 observations: (i) two of these observations did not provide an answer to our first Covid-19 question asking for beliefs of contracting Covid-19 in the next 9 weeks, and (ii) two of these observations consistently provided answers greater than 1,000 to our questions asking for Covid-19 risk assessments out of 1,000 people.²⁴ The median time taken to complete our survey is approximately 10.5 minutes.

²⁴ As part of our IRB approval, respondents were permitted to skip questions. As a result, our number of observations for any particular question is often fewer than our total number of respondents, but typically close to the full sample.

WAVES 2 AND 3 SURVEYS

After analysing the data from our first wave, we conducted two additional waves of our survey. The most significant changes are the inclusion of additional questions, aimed at unpacking the surprising age result, an additional treatment related to question block order, and the addition of an information experiment (only in the Wave 3 survey). We describe these changes below.

Additional Questions

Waves 2 and 3 feature additional questions focused on personal experiences and activities. These questions are placed after the questions that appeared on the original survey, allowing for cleaner comparisons of answers to the original questions across survey waves.²⁵

The first additional questions ask about interactions with individuals who might be perceived to be more vulnerable to Covid-19. In particular, we ask whether the individual has at least one young child at home (under 2), has at least one child under 18 at home, has elderly family members at home, or sees parents or other older family members on a regular basis.

We then turn our attention to three factors that we hypothesized might help to explain our age effect. We ask participants their extent of agreement (1 – 7 scale) with three statements: “at this stage in my life, it is possible/realistic to minimize risks,” over the course of my life, I’ve experienced significant adversity,” and “I was extremely surprised by the emergence of the Covid-19 pandemic.” Following this, we ask specifically about experience with six particular forms of adversity: a serious, life-threatening illness, a serious life-threatening accident or injury, working a job that carries serious health or safety risks, serious illness, injury or untimely death of a loved one, military service, and poverty.

We also ask about personal experiences with Covid-19, asking participants whether they have been infected with Covid-19 (diagnosed by a medical professional), whether they personally know someone who has been infected by Covid-19, and separately, who has been hospitalized due to Covid-19, and separately, who has died due to Covid-19.

We close by asking about political orientation and news sources. Participants are asked to describe their political orientation, choosing from a list ranging from strongly democratic to strongly

²⁵ The one exception to this is that directly following the question asking how many times per week have you left your home, we add a follow-up questions that asks them specifically about different outside of the home activities (i.e. left home for work, went to a bar, ate indoors at a restaurant, etc.). The only “original” question that appears after this follow-up question is their beliefs about the likelihood of a resurgence.

republican. They are then asked about their frequency of consumption of Covid-19 related information from a variety of sources, as well as their degree of trust in those sources.

New Treatment Variation

In the first wave, we randomized the order in which certain survey blocks appeared. In particular, participants either answered questions about their own Covid-19 related health risks first, or saw one of the three blocks asking them to assess others (by age, race, or gender). In Waves 2 and 3, we introduce a new order variation. In particular, we randomize one-fourth of participants into seeing the block that asks about general health risks before they answer questions about their own Covid-19 related health risks. This allows us to ask how thinking about Covid-19 influences estimates of other health risks. We eliminate the treatment that asks participants to assess Covid-19 risks by gender as the first block, replacing it with this new treatment variation.

Information Experiment

In the third wave of the survey, we introduced an information experiment. This information experiment is placed right before the extended block of demographic and personal experience questions that previously closed the survey. In order to implement the experiment, we moved the question asking participants about their state of residence to the front of the survey (alongside our screening questions). Note that all respondents receive this information experiment.

In this experiment, we ask individuals for their best guess of how many people in their state died from Covid-19 between August 1, 2020 – October 1, 2020. Then, we provide them with truthful information about the number of Covid-19 deaths in their state during that time period (according to the Worldometer Covid-19 data tracker; this source is listed as the source for participants).

We then give participants an opportunity to provide a revised estimate of the Covid-19 hospitalization rate and death rate for Americans like themselves (as asked in the own Covid-19 health risks section of the survey). This allows us to consider reaction to information.

Implementation

Waves 2 and 3 were both implemented in partnership with Qualtrics under the same parameters as Wave 1. Qualtrics was instructed to exclude from participation any individual who had participated in a previous wave of our survey.

Wave 2 was conducted between July 15 – July 22, 2020. We were provided with a total of 1,557 responses. One response was dropped from analysis based upon providing multiple answers that exceeded 1,000 to questions that asked about rates out of 1,000; three responses were dropped from analysis because they skipped several consecutive questions.

Wave 3 was launched on October 30, 2020. Unfortunately, Qualtrics had difficulty fielding our targeted sample size of 1,500 respondents. Recruiting slowed significantly and we decided to close the survey with 1,453 responses on December 13, 2020. We dropped one response from analysis because they skipped several consecutive questions.

SUPPLEMENTARY SIMILARITY SURVEY

In May 2022, we ran a simple additional survey, aimed solely at assessing the subjective similarity of different experiences from our original surveys to a severe Covid outcome. We wanted to understand whether our intuitions about perceived similarity aligned with the views of a large, diverse sample, matched in terms of demographics to our original survey population.

Respondents were provided with a list of eight experiences, each of which was asked about in our original 2020 survey waves. The eight experiences were the two components of our “Health Adversities” index (if the respondent ever suffered a serious, life-threatening accident or injury; if the respondent ever suffered a serious, life-threatening illness), the four components of our “Non health adversities” index (if the respondent worked a job that carried serious health or safety risks; if the respondent experienced military service; if the respondent experienced poverty; if the respondent experienced serious injury, illness, or untimely death of a loved one), and two additional adverse experiences: having experienced a non-Covid hospitalization and having experienced a family member hospitalization. The listed order of these experiences was randomized at the individual level.

We asked respondents to force rank the eight experiences according to how similar they perceived each to be to a serious Covid outcome in 2020, where 1 indicated most similar and 8 indicated least similar. We randomized respondents into one of three survey options. The first asked the respondent to rank the experiences according to how similar they were to a severe Covid case in 2020. The second asked the respondent to rank the experiences according to how similar they were to a Covid

hospitalization in 2020. The third asked the respondent to rank the experiences according to how similar they were to a Covid death in 2020.

In order to enable Qualtrics to field a panel matched on demographics to our previous survey waves, respondents were asked to provide their sex, race/ethnicity, income, region, and age in the first block of the survey. In addition, participants had to indicate that they were willing to provide thoughtful answers in order to proceed.

Implementation

The similarity survey was implemented in partnership with Qualtrics under the same parameters as Waves 1 – 3 of our original survey. Data was collected from 1,046 respondents from May 24 – May 26, 2022. Median completion time for the survey was just over two minutes. We pre-registered the survey using AsPredicted; the pre-registration is available here: <https://aspredicted.org/nu8xv.pdf>. We pre-registered the plan to report the mean similarity ranks for each of the eight experiences, without updating our specifications for Table 2.

Results

In Table B1, we report the average rank assigned to each experience, alongside the 95% confidence interval, using each of the individual-level observations. The table is sorted according to perceived similarity. Recall that lower numbers indicate greater perceived similarity.

Table B1. Average Subjective Similarity Rank

	Average Rank at Individual Level	95% CI	
Serious Illness	3.26	3.13	3.39
Loss of Loved One	3.42	3.28	3.55
Accident or Injury	3.83	3.71	3.95
Family Hospitalization	4.29	4.17	4.41
Non-Covid Hospitalization	4.43	4.31	4.56
Dangerous Job	4.89	4.76	5.01
Poverty	5.54	5.41	5.67
Military Service	6.35	6.22	6.47

These results are quite similar when broken out separately according to similarity to a severe Covid case, similarity to a Covid hospitalization, or similarity to a Covid death. See Table B2 below.

Table B2. Average Subjective Similarity Rank, split by Type of Covid Experience

	Average Subjective Similarity Rank		
	Serious Covid Case	Covid Hospitalization	Covid Death
Serious Illness	3.14	3.46	3.20
Loss of Loved One	3.47	3.50	3.27
Accident or Injury	3.96	3.86	3.67
Family Hospitalization	4.36	4.14	4.36
Non-Covid Hospitalization	4.42	4.37	4.50
Dangerous Job	5.01	4.70	4.95
Poverty	5.36	5.68	5.58
Military Service	6.27	6.29	6.47

Appendix C. Summary Statistics and Robustness on Covid Surveys

In this appendix we present:

1. Summary statistics, correlations, and description of the variables included in our analysis in Section 4;
2. The full version of tables 1 and 2. These include all the controls which were not shown in the main text, and regressions for beliefs on Covid infection and hospitalization.
3. A robustness exercise on interference.

Table C1

Summary statistics. The table describes if the variable was collected in all waves or just in waves 2 and 3 of the survey.

Variable	Waves	Min	Max	Mean	sd
Beliefs others death	All	0	1000	85.64	121.87
Beliefs own death	All	0	1000	53.12	114.78
Age	All	18	116	48.89	18.22
Red hair	All	0	1000	55.64	93.56
State Level	All	7	15669	4750.79	5086.03
Days since Peak	All	1	217	42.1	58
No. health conditions	All	0	7	0.88	0.83
Hospital self	All	0	1	0.1	0.3
Hospital family	All	0	1	0.18	0.38
Had Covid	2 & 3	0	1	0.04	0.2
Health adversities	2 & 3	0	2	0.37	0.56
Non health adversities	2 & 3	0	4	0.9	0.78
Subjective adversity	2 & 3	1	7	4.41	1.64

Table C1 presents summary statistics of our variables. Table C2 presents Pearson's correlation coefficients among them. We now give a fine-grained description of them:

- Beliefs others death is the belief on the number of deaths, out of 1000, conditional on contracting Covid in the next 9 weeks, averaging over estimates for gender groups (males/females), age groups (0-39; 40-69; 70+) and race groups (White; African-American; Asian-American; Latinx-American).
- Beliefs own death is the belief on the number of deaths, out of 1000, for “people like self” conditional on contracting Covid in the next 9 weeks.
- Age is the age of the respondent.
- Red hair is the belief of the respondent on the number of Americans, out of 1000, with red hair.
- Level is the cumulative number of deaths for Covid in the respondent's state, at the time of maximum weekly growth of deaths in the state. Maximum weekly growth is defined as the day with

the highest increase in 7 days rolling average of daily deaths increases, (death number on day t minus death number of day $t-7$).

- Peak is the number of days since the time of maximum weekly growth of cases in the State, where maximum weekly growth is defined in the same fashion as for deaths.
- Number of health conditions takes values from 0 to 7 and considers: diabetes; heart disease; lung disease; hypertension; obesity, cancer; other serious immunocompromising condition.
- Hospital self is a dummy equal to 1 if the respondent was hospitalized, not for Covid, in the last year.
- Hospital family is a dummy equal to 1 if a family member of the respondent was hospitalized, not for Covid, in the last year.
- Had Covid is a dummy equal to 1 if the respondent has been infected with Covid-19 (diagnosed by a medical professional).
- Health adversities takes values from 0 to 2 and considers if the respondent has personally experienced i) a serious, life-threatening accident or injury; ii) a serious, life-threatening illness.
- Non health adversities takes values from 0 to 4 and considers if the respondent has personally experienced any of the following: i) worked a job that carried serious health or safety risks; ii) serious illness, injury, or untimely death of a loved one; iii) military service; iv) poverty.
- Subjective adversity is the rate of agreement with the statement “Over the course of my life, I've experienced significant adversity”. It takes values from 1 (not at all) to 7 (completely agree).

Table C2

Correlations among variables. Green correlation coefficient are significant at 5% level.

	Others death	Age	Red hair	Level	Days	Health cond	Hosp self	Hosp fam	Had Covid	Health adv	Non h adv	Subj adv
Beliefs others death	0.56	-0.28	0.18	0.09	0.02	-0.02	0.12	0.12	0.13	0.06	-0.05	0.11
Beliefs others death		-0.15	0.18	0.05	0.01	0.06	0.11	0.08	0.1	0.08	0	0.1
Age			-0.09	-0.2	-0.14	0.26	-0.14	-0.23	-0.11	0.06	0.09	-0.14
Red hair				0.05	0.05	0	0.05	0.03	0.03	0.02	-0.03	0.07
State Level					0.66	0	0.15	0.17	0.03	-0.02	-0.08	0.09
Days since Peak						0.03	0.14	0.15	0.03	0	-0.04	0.08
No. health conditions							0.11	0.06	0.06	0.28	0.19	0.13
Hosp self								0.39	0.13	0.17	0.01	0.13

Hosp fam									0.09	0.11	0.06	0.13
Had Covid										0.13	-0.02	0.09
Health adversities											0.07	0.21
Non health adversities												0.19

Table C3 presents the full list of selected regressors for beliefs about Covid fatality for others, for self, the risk of own infection and of own hospitalization. The risk of own infection is the belief on the number of Covid infections, out of 1000, for “people like self” in the next 9 weeks. Own hospitalization is the belief on the number of Covid hospitalizations, out of 1000, for “people like self” conditional on contracting Covid in the next 9 weeks. We can see that all the results regarding fatality also hold for infections and hospitalization.

Table C3

Own death is the belief on the number of deaths, out of 1000, for “people like self” conditional on contracting Covid in the next 9 weeks. Others death is the belief on the number of deaths, out of 1000, conditional on contracting Covid in the next 9 weeks, averaging over estimates for gender groups (males/females), age groups (0-39; 40-69; 70+) and race groups (White; African-American; Asian-American; Latinx-American). Own infection is the belief on the number of Covid infections, out of 1000, for “people like self” in the next 9 weeks. Own hosp is the belief on the number of Covid hospitalizations, out of 1000, for “people like self” conditional on contracting Covid in the next 9 weeks. All variables are standardized except for dummy variables (Hosp self; Hosp fam; Black; Asian; Rural). Red hair is the belief of the respondent on the percentage of red-haired Americans. Level is the cumulative number of deaths for Covid in the state, at the time of maximum weekly growth in the state. Days is the number of days since the peak of cases in the state. No. of health conditions takes values from 0 to 7 and considers: diabetes; heart disease; lung disease; hypertension; obesity, cancer; other serious immunocompromising condition. Hosp self (fam) is a dummy equal to 1 if the respondent (a family member) was hospitalized, not for Covid, in the last year. Income is the income of the respondent. Rural, Asian, and Black are dummies referring to the residential area or ethnicity of the respondent.

	<i>Dependent variable:</i>			
	Own death (1)	Others death (2)	Own infection (3)	Own hosp (4)
Age	-0.131*** (0.019)	-0.236*** (0.015)	-0.183*** (0.013)	-0.112*** (0.015)
Red hair	0.163*** (0.032)	0.155*** (0.019)	0.171*** (0.029)	0.130*** (0.026)
State Level	0.037** (0.015)	0.073*** (0.014)	0.071*** (0.014)	0.077*** (0.017)
Days since Peak	-0.057*** (0.013)	-0.084*** (0.015)	-0.088*** (0.012)	-0.083*** (0.018)
No. health conditions	0.090*** (0.015)	0.032*** (0.011)	0.027** (0.013)	0.039*** (0.013)
Hosp (self.)	0.245*** (0.078)	0.231*** (0.062)		0.319*** (0.065)
Hosp (fam.)		0.093*** (0.036)	0.156*** (0.048)	0.099*** (0.038)

Income	-0.036** (0.016)	-0.044*** (0.016)	-0.083*** (0.013)	-0.043** (0.019)
Black	0.111** (0.053)	0.164*** (0.048)		0.084** (0.042)
Asian		0.205*** (0.060)		
Rural	0.123*** (0.033)	0.068** (0.030)		0.064* (0.035)
Constant	-0.084*** (0.022)	-0.103*** (0.022)	-0.027* (0.014)	-0.086*** (0.018)
Observations	4,514	4,477	4,506	4,511
R ²	0.073	0.122	0.081	0.063
Adjusted R ²	0.071	0.120	0.080	0.060

Note:

*p<0.1; **p<0.05; ***p<0.01

Clustered standard errors at state level

Table C4 presents the full output of Table 1 in column 1. Column 2 shows that our results, that higher non health adversities lead to lower pessimism, hold if we omit “serious injury, illness or untimely death of a loved one” from non-health adversities. Column 3 shows that political affiliation does not predict beliefs, nor does it affect the role of the other regressors.

Table C4

Others death is the belief on the number of deaths, out of 1000, conditional on contracting Covid in the next 9 weeks, averaging over estimates for gender/age/race groups. More precisely, a first estimate is obtained averaging over beliefs for males and females; a second estimate is obtained averaging over beliefs for three age groups (0-39; 40-69; 70+); a third estimate is obtained averaging over beliefs for four race groups (White; African-American; Asian-American; Latinx-American). The final estimate is obtained averaging these three estimates. All variables, but dummies, are standardized. Health adversities is an index given by the sum of two dummies indicating 1) if the respondent ever suffered a serious, life-threatening accident or injury; 2) if the respondent ever suffered a serious, life-threatening illness. Non health adversities is an index given by the sum of four dummies: indicating 1) if the respondent worked a job that carried serious health or safety risks; 2) if the respondent experienced military service; 3) if the respondent experienced poverty; 4) if the respondent experienced serious injury, illness, or untimely death of a loved one. Non health adversities (small) does not consider the fourth one. Subjective adversity is the rate of agreement with the sentence “Over the course of my life, I’ve experienced significant adversity.” Republican degree is a variable which measures political orientation of the respondent and it takes values from 1 (Strongly Democratic) to 7 (Strongly Republican).

	<i>Dependent variable:</i>		
	Others death		
	(1)	(2)	(3)
Had Covid	0.441*** (0.167)	0.446*** (0.167)	0.442*** (0.167)
Health adversities	0.047** (0.019)	0.046** (0.019)	0.044** (0.019)

Non health adv.	-0.039*** (0.015)		-0.037** (0.015)
Non health adv. (small)		-0.031* (0.016)	
Subj. adversity	0.043** (0.019)	0.041** (0.019)	0.043** (0.020)
No. health cond.	0.012 (0.017)	0.010 (0.016)	0.014 (0.016)
Hosp (self.)	0.157** (0.073)	0.160** (0.073)	0.158** (0.073)
Hosp (fam.)	0.058 (0.044)	0.050 (0.044)	0.055 (0.045)
State Level	0.059*** (0.023)	0.061*** (0.023)	0.059** (0.023)
Days since Peak	-0.097*** (0.023)	-0.097*** (0.023)	-0.096*** (0.023)
Red hair	0.165*** (0.033)	0.166*** (0.033)	0.166*** (0.032)
Age	-0.212*** (0.021)	-0.216*** (0.020)	-0.211*** (0.023)
How Republican			-0.022 (0.026)
Income	-0.043* (0.023)	-0.042* (0.022)	-0.039* (0.024)
Black	0.133** (0.054)	0.136** (0.054)	0.119** (0.053)
Asian	0.249*** (0.092)	0.252*** (0.091)	0.249*** (0.089)
Rural	0.113** (0.044)	0.116*** (0.044)	0.101** (0.044)
Constant	-0.128*** (0.030)	-0.129*** (0.029)	-0.123*** (0.027)
Observations	2,953	2,953	2,944
Adjusted R ²	0.133	0.132	0.133
<i>Note:</i> *p<0.01; ** p<0.05; *** p<0.01 Clustered standard errors at state level			

Beliefs, political affiliation, and behavior. Existing work on the pandemic has stressed the importance of political beliefs in shaping behaviour (e.g. wearing a mask) and policy views. Do memory-based beliefs about the lethality of Covid, which are only modestly influenced by politics, affect behaviour?

Our survey measured behaviour and attitudes, including how often respondents leave home for reasons other than work or exercise, whether they have recently forfeited medical care to avoid leaving home, and whether they favour lifting the lockdown measures in place at the time of the survey. Of course, past experiences may affect behaviour through a variety of channels. For instance, respondents with past health adversities may refrain from going out because it is harder for them to do so, not necessarily because they are more pessimistic about Covid. To address this issue, we use the “red hair” proxy as an instrument for beliefs. The idea is that “red hair” captures respondent’s general tendency to overestimate unlikely events, regardless of whether they concern risk or not. As a result, if “red hair” helps explain behaviour, it arguably does so via beliefs.²⁶

Table C5 reports our regressions. In columns (1), (3) and (5) we show the role of beliefs in OLS specifications in which we control for the best predictors of behaviour selected by our method. In columns (2), (4) and (6) we instrument beliefs using the red hair proxy. Relative to Table 1, we add political affiliation (‘how republican’) which, while not selected as a predictor of beliefs, is a commonly cited predictor of attitudes towards the pandemic (Bursztyn et al 2020).

Respondents who estimate higher “red hair”, and hence have more pessimistic beliefs about Covid, behave more cautiously. Interference in retrieval affects beliefs and, through this channel, memory affects behaviour. This only occurs, however, for individual decisions, not for a policy preference such as whether to lift the lockdown. Political affiliation instead emerges as a key predictor for policy preferences, consistent with existing work.

Table C5

Going out is the answer to the question “Over the last few weeks, approximately how many times per week have you left your home to shop, do errands, socialize, etc.?”. It takes values 1 (never), 2 (once a week), 3 (twice a week), 4 (three or more times a week). Med avoid is the answer to the question “Have you avoided filling prescriptions at the pharmacy, doctor’s appointments, or other forms of medical care in the last few weeks?”. It takes values 1 (Yes, completely), 2 (Somewhat), 3 (Not at all). Lift lockdown is the answer to the question “Would you resume your normal activities if lockdown or “stay-at-home” measures were lifted today?”. It takes value from 1 (Definitely yes) to 5 (Definitely not). Death others is the belief on Covid death for others, as described in tables 1 and 2. It is obtained as the average of the estimated risk of death for separate age, ethnicity and gender classes. This is instrumented with the estimated number of red-haired Americans ($F >> 10$ in all cases). Republican degree is a variable which measures

²⁶ Red hair also has a low correlation with the other predictors of beliefs. It has a -0.09 correlation with “Age”. The next variable in the survey whose correlation with red hair is highest in magnitude is “Subjective Adversities” which has a 0.07 correlation with red hair.

political orientation of the respondent and it takes values from 1 (Strongly Democratic) to 7 (Strongly Republican). All variables are standardized and controls include variable which were selected by performing a dependent variable specific model selection algorithm. Max weekly growth death is the maximum weekly growth of Covid deaths in the state. Days since weekly death peak is the number of days since Covid deaths peak in the state. Current level death is the current cumulative level of Covid deaths in the state. Unemployment is a dummy equal to 1 if the respondent experienced unemployment in the last nine weeks.

	<i>Dependent variable:</i>					
	Going out	Going out	Med avoid	Med avoid	Lift Lockdown	Lift Lockdown
	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Death others	-0.071*** (0.023)	-0.228** (0.112)	-0.057** (0.023)	-0.278** (0.114)	-0.002 (0.019)	-0.119 (0.098)
Max weekly growth death	-0.057*** (0.014)	- 0.055*** (0.014)				
Days since wk death peak	0.044* (0.023)	0.036 (0.023)				
Current level death			-0.019 (0.023)	-0.028 (0.020)		
Age	0.065*** (0.018)	0.023 (0.031)	0.227*** (0.016)	0.169*** (0.031)		
Age squared			0.065*** (0.016)	0.076*** (0.015)		
Female	-0.051*** (0.019)	-0.049** (0.020)			0.113*** (0.020)	0.115*** (0.021)
Black					0.026 (0.018)	0.034* (0.019)
Asian	-0.071*** (0.019)	- 0.062*** (0.017)			0.056*** (0.014)	0.066*** (0.018)
Rural			-0.102*** (0.020)	-0.089*** (0.019)		
Education			-0.092*** (0.017)	-0.093*** (0.019)		
West					0.025	0.022

					(0.023)	(0.024)
Suburban					0.083***	0.072***
					(0.016)	(0.017)
Income					-0.092***	-0.091***
					(0.017)	(0.018)
No. health conditions	-0.083***	-0.076***	-0.084***	-0.076***	0.056**	0.056**
	(0.020)	(0.018)	(0.020)	(0.022)	(0.022)	(0.022)
Hosp (fam)	0.056***	0.064***				
	(0.016)	(0.017)				
Hosp (self)			-0.082***	-0.067***		
			(0.020)	(0.025)		
Unemployment			-0.032*	-0.028		
			(0.019)	(0.018)		
State population	-0.038**	-0.035**	-0.034**	-0.026*	-0.079**	-0.064
	(0.017)	(0.015)	(0.014)	(0.016)	(0.038)	(0.042)
Republican degree	0.090***	0.083***	0.012	0.003	-0.261***	-0.267***
	(0.022)	(0.023)	(0.018)	(0.025)	(0.043)	(0.047)
Constant	0.115***		-0.042*		0.082***	
	(0.018)		(0.024)		(0.021)	
Controls	YES	YES	YES	YES	YES	YES
Observations	2,962		2,960		2,963	
R ²	0.043		0.141		0.122	
Adjusted R ²	0.039		0.138		0.119	

Note:

*p<0.1; **p<0.05; ***p<0.01

Clustered standard errors at state level

Table C6 presents a more complete analysis of interference. It reports the coefficient of the interaction among all Covid and non-Covid adversities. We also report the coefficient of the interaction of a variable with itself, obtained by adding the square of that variable to the corresponding regression. For the sake of clarity and brevity, health adversities include serious injury, serious illness, and hospital self. Hence, it is defined from 0 to 3, differently from Table 2. Green indicates agreement with our theory, yellow disagreement. A darker color corresponds to a

lower p-value. We can see that, consistent with Figure 4, interference is present across the board, with the strongest ones being among i) Level and family hospital; ii) health conditions and family hospital. The square of the number of health conditions has a strong and negative coefficient, meaning that numerous health conditions interfere one with the other in shaping pessimism.

Table C6

Each cell reports the interaction estimated between the row and the column, together with their p values in parentheses. A green cell indicates that the sign of the coefficient directionally matches the prediction of the theory, a yellow cell indicates that it does not. Darker colors indicate lower p value. Interactions were estimated adding them to the model presented in table 1 column 2, if the two variables were available in all waves. They were estimated adding them to the model presented in table 2 column 2, if at least one of the two variables was available only in waves 2 and 3. The interaction of a variable with itself represents the coefficient of the square of the variable. Health adversities takes values from 0 to 3 and it includes serious injury, serious illness, and own hospital.

Others Death	Level	Health cond	Family hosp	Health adv	Had Covid
Level	-0.009 (0.399)	-0.007 (0.572)	-0.072 (0.000)	-0.032 (0.061)	-0.153 (0.052)
Health conditions		-0.011 (0.006)	-0.112 (0.000)	-0.015 (0.298)	-0.077 (0.459)
Family hospital				-0.013 (0.762)	-0.132 (0.714)
Health adversities				-0.007 (0.660)	0.022 (0.875)

Figure C1 extends the analysis of Figure 3 by examining interference between non-Covid health adversities and the experience of having had Covid. As in Figure 3, having had Covid reduces the marginal impact of non-Covid health experiences, and vice versa, except for the index of health adversities (serious illness or injury, or own hospitalization).

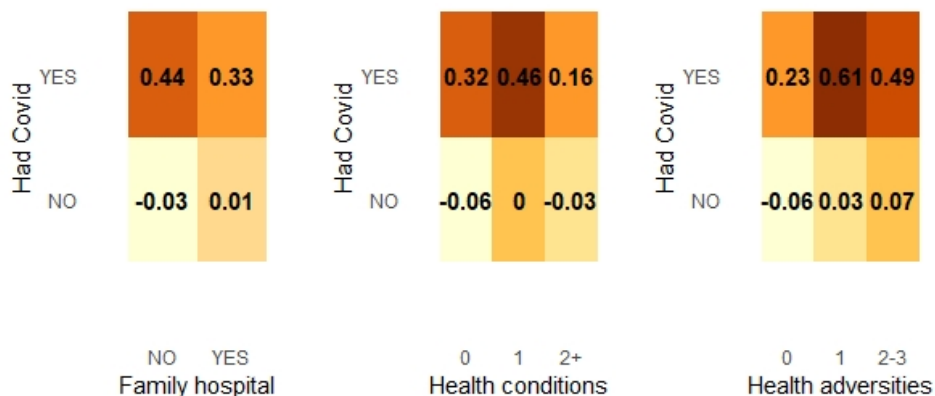


Figure C1.

The Figure reports the residuals of the standardized beliefs of FATALITY (for others), estimated by removing from the model of Table 2's Column 2 the variables "Had Covid" and i) "Family Hospitalization", ii) "Number of Health Conditions", and iii) "Health Adversities". Health adversities refer to the sum of serious injury, serious illness, and self

hospitalization dummies. Reported values are average residuals in each cell. Different colours indicate different average residuals up to the third decimal.

Appendix D. Model Selection

The regressions presented in the main text show output models obtained from best subset selection. In our survey, we collect several demographics and ask several behavioral questions, along with beliefs about Covid. This is a typical case where we might want to remove irrelevant predictors. There are two compelling reasons to do that: i) when the number of predictors is high, prediction accuracy of the OLS model will be good but there might be a lot of variability in the least squares fit; ii) interpretability of models which include a lot of predictors is difficult. It is often the case that some or many of the variables used in a multiple regression model are in fact not associated with the response. Including such irrelevant variables leads to unnecessary complexity in the resulting model. By removing these variables—that is, by setting the corresponding coefficient estimates to zero—we can obtain a model that is more easily interpreted. Although in our case the number of observations is much higher than the number of potential covariates (hence variability should not be an issue), we still aim at keeping only the most relevant predictors. To do so, we employ a machine learning algorithm called best subset selection (Guyon and Elisseeff, 2003; James et al., 2013). Other applications of best subset selections in economics include Alabrese and Fetzer (2018) and Becker et al. (2017). The method works as follows: we fit a separate least squares regression for each possible combination of the p predictors. That is, we fit all p models that contain exactly one predictor, all $\binom{p}{2}$ models that contain exactly two predictors, and so forth. We then look at all of the resulting models, with the goal of identifying the one that is best, according to some information criteria. More formally, the algorithm entails the following steps:

- 1) We denote \mathcal{M}_0 the *null model*, containing no covariates;
- 2) For $k \in \{1, 2, \dots, p\}$ we:
 - a) Fit all $\binom{p}{k}$ models containing k covariates;
 - b) Pick the best of these $\binom{p}{k}$ models and denote it \mathcal{M}_k . The best model is the one with the highest R^2 . In every set of models with k covariates, we can compare them by using the R^2 , since the number of covariates is fixed within the set;
- 3) Select the best model, among $\mathcal{M}_0, \dots, \mathcal{M}_p$ using cross-validation or an information criterion (Mallow's C_p , BIC, adjusted R^2).

We can express the best subset selection problem as a nonconvex and combinatorial optimization problem. The objective is to find the optimal s for:

$$\min_{\beta} \sum_i^n \left(y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j \right)^2 \quad \text{subject to} \quad \sum_{j=1}^p I(\beta_j \neq 0) \leq s$$

This requires that the optimal solution involves finding a vector β such that the residual sum of squares is minimized and no more than s coefficients are different from 0. The algorithm presented above (points 1-3) solves this optimization problem for every value of s and then picks among the optimal models for the different values of s . Best subset selection can thus be expressed as a regularized regression with penalization term equal to $\sum_{j=1}^p I(\beta_j \neq 0)$.

In point 3 of our description of the algorithm, we refer to the selection of the best model, among $\mathcal{M}_0, \dots, \mathcal{M}_p$. We will discuss three information criteria: Mallows's C_p , Bayesian information criterion (BIC), and adjusted R^2 . Mallows's C_p is defined as $C_p = \frac{1}{n}(RSS + 2d\hat{\sigma}^2)$, with RSS being the residual sum of squares, d the total number of parameters used and $\hat{\sigma}^2$ is an estimate of the variance of the error ϵ associated with each response measurement. In the case of the linear model with Gaussian errors, C_p is equivalent to the Akaike information criterion (AIC). BIC is defined as $BIC = \frac{1}{n}(RSS + \log(n) d\hat{\sigma}^2)$. The BIC replaces $2d\hat{\sigma}^2$ with $\log(n) d\hat{\sigma}^2$. Since, $\log(n) > 2$ if $n > 7$, the BIC places a heavier penalty on models with many variables and it usually selects smaller models than the C_p . As can be easily guessed, to identify the best model we aim at minimizing either the Mallows's C_p or the BIC. The adjusted R^2 is defined as $\text{adj}R^2 = 1 - \frac{RSS/(n-d-1)}{TSS/(n-1)}$ where TSS is the total sum of squares. The best model is the one which maximizes the adjusted R^2 . Finally, we can use m -fold cross-validation. This proceeds as follows: i) divide the sample of n observation in into m non-overlapping groups (folds), each containing around $\frac{n}{m}$ observations; ii) for each $z \in \{1, 2, \dots, m\}$ treat fold z as a validation set, fit the model on the remaining folds and compute the mean squared error, MSE_z pertaining to the withheld validation set z ; iii) compute $CV_m = \frac{1}{m} \sum_{z=1}^m MSE_z$. We will then choose the model with the lowest cross-validation error. What is the best criterion to use is an issue which goes beyond the scope of this discussion. We can refer the reader to Ding et al. (2018). To give a sense of this discussion, in figure A1 we show a comparison of the four decision criteria, applied to the choice of the best model to predict the number of times the respondent had gone out in the period before the survey (table 3 column 1).




Figure A1

Adjusted R^2 , Mallows's C_p , BIC and cross-validation error to select the best model to describe the propensity to go out. The best model, according to each criterion, is highlighted in red.

The set of potential predictors is the set of demographics and we can see that the BIC selects the regression with 6 covariates, namely age, dummy for female, dummy for Asian, Number of health conditions, family member been hospitalized (not for covid), and population of the state, which we included as controls in table 3.²⁷ Figure A1 offers the perfect insight to reflect on the different information criteria. BIC suggests that the best model is the one with 6 covariates. We have already explained why the BIC tends to select more parsimonious models. In this case both the adj. R^2 and cross-validation suggest to use a 14 covariates model and Mallows's C_p suggests to include 11 covariates. However, we can see that the 6 variable model is very close to the best model for each of the four criteria. This was the principle which guided us in our work. We usually selected the best model, according to the BIC criterion, and verified if this was close to be optimal for the other three.

We now give some more details on how we selected the best model for each of our dependent variables. Tables 1 and 2 report the output of the models we selected to describe beliefs about Covid death. A similar procedure is employed to describe beliefs about Covid infection and hospitalization. We split the variables in 3 sets:

- 1) Set A: state level Covid dynamics. For all the three waves it contains the following variables (for Covid cases or deaths): current level; maximum weekly growth; days since growth peak; current weekly growth; level at the time of maximum growth;
- 2) Set B: personal characteristics and Covid experiences. For all the three waves it contains the following variables: age, gender, ethnicity, region, income, urbanization, employment, a lot

²⁷ Table A4 reports also variables on Covid dynamics, which were the object of a separate variable selection and politics, which was added for theoretical reasons.

of health info on the self and family, state population, the estimated number of red haired Americans;

- 3) Set B': these are additional variables in waves 2 and 3: interactions with family members, several measures of adversities in life, several measures of direct and indirect exposure to Covid; political preferences; several opinions on Covid.

One caveat with best subset selection is that certain variables may be dropped in case they are highly correlated with each other. This is why, in some cases we perform some minimal form of supervision, like for example retaining some predictors which are very relevant according to our memory model, but were not selected by the machine learning algorithm.²⁸

Our model selection consists of the following stages:

- 1) We perform model selection, for each of the 4 dependent variables (Covid infection, hospitalization, and death for self, Covid death for others), in set A of state level Covid dynamics (10 predictors);
- 2) We perform some minimal supervision on model selection. We select the model that contains the most robust predictors across the four types of beliefs. This leads to the inclusion of the days since the weekly cases growth peak, and the level of cases in the state of the respondent at the time of maximum weekly growth of cases;²⁹
- 3) We perform model selection, for each of the 4 dependent variables, in set B and B' of demographics (23 predictors for all waves; 35 predictors for waves 2 and 3);
- 4) We show the resulting models which contain the variables selected in stages 1-3 in table 1;
- 5) Table 2 column 2 contains the best model obtained when performing model selection in waves 2 and 3, plus all the covariates which were selected on all waves (table 1 column 2), even if they were excluded by performing model selection in the last two waves.

A similar procedure is employed to select the best subset of predictors from set B to predict the number of times the respondent had gone out, the tendency to avoid medical appointments, and the preference for lifting lockdown. These are included in table 3. We included political orientation as a

²⁸ For example, health adversities and non health adversities. Each of them had been considered separate potential predictors and serious injury only had been selected. We decided to include them jointly as indices.

²⁹ To give a sense of how our mild supervision worked, best subset selection suggested those two predictors for all but one dependent variable. For beliefs about infection, the best model would have included the maximum weekly growth of cases in the state, instead of the level. The model we picked had negligible differences with the "optimal" one, in terms of prediction accuracy.

control in table 3, since this is believed to be a relevant factor in orienting behavior and policy preference regarding “stay-at-home” measures.

References for Appendix D

Alabrese, E. and Fetzer, T. Who is Not Voting for Brexit Anymore? (2018). CESifo Working Paper No. 7389, Available at SSRN: <https://ssrn.com/abstract=3338743>

Becker, S. O., Fetzer, T., and Novy, D. (2017). Who voted for Brexit? A comprehensive district-level analysis. *Economic Policy*, 32(92): 601-650

Ding, J., Tarokh, V., and Yang, Y. (2018). Model selection techniques: An overview. *IEEE Signal Processing Magazine*, 35(6), 16-34.

Guyon, I. and A. Elisseeff (2003). An Introduction to Variable and Feature Selection. *Journal of Machine Learning Research* 3, 1157–1182

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). An introduction to statistical learning. New York: Springer

Appendix E. The Primed-Recall Experiment

The following pages present the primed-recall experimental materials.

What is the purpose of this research?

This study aims to understand how past experiences inform beliefs about the future.

What can I expect if I take part in this research?

If you take part in this study, you will complete a 15-minute survey. You will be asked questions about your past experiences and about your expectations about the likelihood of different events in the future. You will be asked to provide brief written answers about these events. You will also be asked some basic questions about yourself. If you complete the survey, you will be paid \$3.25. This payment will be made via the Prolific platform within 48 hours of your participation.

Please be aware that there are comprehension checks that you must answer correctly in order to complete the study. Failure to answer these comprehension checks correctly could result in dismissal from the study without pay.

While we believe the risks of this study are minimal, you should be aware that some of the experiences we ask you about may be unpleasant to think about. You may choose to skip questions about past experiences that you are uncomfortable answering.

You will not interact with any other participants in this study.

Your answers will be linked to your Prolific ID at the time of data collection. The research team will remove your Prolific ID from the dataset before analyzing the data or sharing aggregate results.

Please do not provide any other information in your answers to this survey that would allow for others to identify you – this includes names, zip codes, contact information, and other personal details. The research team will attempt to remove any individually identifiable information from your responses before analyzing or sharing your data. This de-identified data will be stored for future research use and may be shared with other researchers or research participants.

What should I know about a research study?

- Whether or not you take part is up to you.
- Your participation is completely voluntary.
- You can choose not to take part.
- You can agree to take part and later change your mind.
- Your decision will not be held against you.
- Your refusal to participate will not result in any consequences or any loss of benefits that you are otherwise entitled to receive.
- You can ask all the questions you want before you decide.

You may not be told everything

As part of this research design, you may not be told everything about the purpose or procedures of this research. There are different versions of this study. While you will be fully informed about the version of this study that you have been assigned to, you will not be informed about different versions of this study that other participants are in.

Who can I talk to?

If you have questions, concerns, or complaints, or think the research has hurt you, talk to the

research team: Katherine Coffman, kcoffman@hbs.edu, 617 495 6538.

Thank you for participating today.

We care about the quality of our survey data and hope to receive the most accurate measures of your beliefs. It is important to us that you provide thoughtful, careful answers to this survey.

Do you commit to providing your thoughtful and careful answers to the questions in this survey?

- I commit to providing thoughtful and careful answers
- I do not commit to providing thoughtful and careful answers
- I can't promise either way

Some of the questions in this survey will ask you to make your best estimate as to the likelihood of different things. To give you some practice and get you used to thinking in these terms, we have a few example questions for you to work through. Note that you are required to answer the following questions in order to proceed.

Example 1: According to the United States Census, **approximately 2 out of every 100 Americans** live in Massachusetts. This is equivalent to approximately 2%.

Using this estimate, tell us **how many out of 5,000 randomly-chosen Americans** live in Massachusetts?

Example 2: Think about **a group of 100 randomly-chosen American cities**. How many of them do you believe receive more than 1ft of snow in a typical winter?

When asked, "Out of 100 randomly-chosen American cities, how many receive more than 1ft of snow in a typical winter?", you entered that you believe $\{q://QID49/ChoiceTextEntryValue/1\}$ out of 100 randomly-chosen American cities receive more than 1ft of snow in a typical winter.

This is equivalent to X%.

Now, using the same estimate you provided above, please tell us, out of **1,000** randomly-chosen American cities, how many receive more than 1ft of snow in a typical winter?

Great! From this point forward, you may choose to skip questions that you feel uncomfortable answering. Please click next to get started on the survey.

[Note: Individuals are randomly-assigned to see one of the three ‘experience prompts’ below, or randomly-assigned to a control treatment that proceeds directly to the next section.]

EXPERIENCE PROMPTS

In this part of the survey, we will ask you about an experience you may have had. We appreciate you taking the time to provide thoughtful, honest answers to these questions. An important reminder before you start:

Please do NOT provide any information in your answers to these questions that would allow others to identify you – this includes, for example, names, zip codes, or contact information.

[IDTHEFT]

Have you, a loved one, or your employer dealt with identity theft, a data breach, stolen credit card information, or a compromised password in the past?

Yes

No → skip to next section

Thinking back on this experience, please respond to the questions below.

- What were the intentions of the perpetrators? If you do not know, please give us your best guess.
- Please briefly describe any damages, losses, or inconvenience you, your loved one, or your employer suffered as a result of this experience.
- What emotions do you associate with thinking back on that experience?

[FINANCE]

Have you or a loved one ever struggled with finances?

Yes

No → skip to next section

- What was the main cause of that struggle?
- What losses, difficulties, or sacrifices do you remember from that time?
- What emotions do you associate with thinking back on that experience?

[LOSS]

Have you lost a loved one to illness?

Yes

No → skip to next section

- Thinking back on this experience, please respond to the questions below.
- Briefly - using no names or other personally identifiable information - describe who the loved one was and what the illness was.
- What losses, difficulties, or sacrifices do you remember from that time?

[For all Experience Prompts]

- Please take the next few minutes to type 3 – 4 sentences about this personal experience. Feel free to write whatever comes to mind as you think about this experience.

- On a scale of 1 – 7, where 1 is not at all and 7 is extremely vividly, how vividly do you remember this personal experience?
-

[LIKELIHOOD ASSESSMENTS – Everyone sees these]

Please take a moment to read the following information about cyberattacks. This information is adapted from IBM's web resources.

Please take a moment to read the following information about cyberattacks. This information is adapted from IBM's web resources.

Cyberattacks are unwelcome attempts to steal, expose, alter, disable or destroy information through unauthorized access to computer systems. They can be launched by individual actors, like hackers, or by criminal organizations or state actors. While not always the case, cyberattacks can be associated with cyber warfare or cyberterrorism.

Cyber attacks target things like financial data, personal or sensitive data, intellectual property, infrastructure, or government departments and agencies. Not all cyberattacks are successful. Many organizations and governments have implemented plans to limit the possibility of disruption or harm. But, if successful, cyberattacks can cause significant damage - data loss or manipulation, major service disruptions, and financial losses.

We are interested in your beliefs about the chances of a significant cyberattack in the future. In particular, please consider the questions below.

What do you believe is the likelihood that **you will be significantly impacted** by a cyberattack over the next 5 years? This impact could come in the form of financial loss or through damage to critical infrastructure, such as power lines, hospitals, banking systems, communication satellites, or manufacturing.

Please indicate on a scale of 0 - 100, where **0 indicates that there is no chance** at all that you will be significantly impacted by a cyberattack and **100 indicates that you will definitely be significantly impacted**.

Now, think of **1,000 people just like you** in the United States.

Out of those 1,000 people, how many do you believe will be significantly impacted by a cyberattack over the next 5 years? Again, this impact could come in the form of financial loss or through damage to critical infrastructure, such as power lines, hospitals, banking systems, communication satellites, or manufacturing.

When answering the previous question, how vividly did you imagine what an infrastructure-disrupting cyber attack would be like?

Please answer on a scale of 1 – 7 where 1 is not vividly at all and 7 is extremely vividly.

Please take a few minutes now to write down 3 – 4 sentences about any thoughts that came to mind when you were answering the previous questions on cyberattacks.

For each of the experiences below, please indicate whether you have had this past experience.









	Yes (1)	No (0)
Have you, a loved one, or your employer dealt with identity theft, a data breach, stolen credit card information, or a compromised password in the past? (Had_ID)	<input type="radio"/>	<input type="radio"/>
Do you have memories of the Sept. 11th, 2001 terrorist attacks in the United States? (Had_Sept11)	<input type="radio"/>	<input type="radio"/>
Have you lost a loved one to serious illness?	<input type="radio"/>	<input type="radio"/>
Have you or a loved one experienced a recent extreme weather event in your area, such as a blizzard, hurricane, tornado, or flood?	<input type="radio"/>	<input type="radio"/>
Have you or a loved one ever struggled with finances?	<input type="radio"/>	<input type="radio"/>
Have you or a loved one been hospitalized recently?	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/>	<input type="radio"/>
Have you or a loved one dealt with addiction?	<input type="radio"/>	<input type="radio"/>
Have you or a loved one had a serious accident or injury recently?	<input type="radio"/>	<input type="radio"/>



For each of the experiences below, please indicate **how similar you believe that experience is to a significant cyberattack** - in particular, a cyberattack that significantly disrupts critical civilian infrastructure, such as power lines, hospitals, banking systems, communication satellites, or manufacturing.

Please answer on a scale of 1 – 7 where **1 is Not at all similar** and **7 is Extremely similar**.

1 2 3 4 5 6 7

Identity theft, a data breach, or having had credit card information or a password stolen ()	
The Sept. 11th terrorist attacks ()	
Losing a loved one to serious illness ()	
An extreme weather event such as a blizzard, hurricane, tornado, or flood ()	
Struggling with finances ()	
Having been hospitalized recently ()	
Dealing with addiction ()	
A serious accident or injury ()	

[DEMOGRAPHICS]

What is your year of birth?

What is your sex?

- Male
- Female
- Non-binary
- Prefer not to say

Choose one or more races that you consider yourself to be:

- White
- Latino or Latina
- Black or African American
- Asian
- Other
- Prefer not to say

Information about income is very important to understand. Would you please give your best guess? Please indicate the answer that includes your entire household income in (previous year) before taxes.

- Less than \$25,000
- \$25,000 to \$49,999
- \$50,000 to \$74,999
- \$75,000 to \$99,999
- \$100,000 to \$124,999
- \$125,000 to \$150,000
- \$150,000 or more

What region best describes where you live?

- Northeast
- South
- Midwest
- West

- Other

What is the highest level of school you have completed or the highest degree you have received?

- ☐ Less than high school degree
- ☐ High school graduate (high school diploma or equivalent including GED)
- ☐ Some college but no degree
- ☐ Associate degree in college (2-year)
- ☐ Bachelor's degree in college (4-year)
- ☐ Master's degree
- ☐ Doctoral or Professional degree (PhD, JD, MD)

In a typical year, approximately how much snow falls in the area you live in? We encourage you to use the internet to look up the answer for your area.

Appendix F. Supplemental Results for Primed Recall Experiment

Table F1.
The Impact of Primed and Lived Experiences on Cyberattack Estimates, ITT

	OLS Predicting Index of Cyberattack Estimates		
	High Quality Only		
	1	2	3
<i>ID theft</i> prime	0.15** (0.061)	0.15** (0.060)	0.12* (0.072)
<i>Finance</i> prime	0.14** (0.062)	0.15** (0.060)	0.19** (0.077)
<i>Loss</i> prime	0.057 (0.061)	0.054 (0.060)	0.14* (0.078)
Below Median Similarity x <i>ID theft</i> prime			0.058 (0.085)
Below Median Similarity x <i>Finance</i> prime			-0.085 (0.086)
Below Median Similarity x <i>Loss</i> prime			-0.15* (0.086)
Had <i>ID theft</i>		0.46*** (0.050)	0.46*** (0.050)
Had <i>finance</i>		0.22*** (0.064)	0.21*** (0.064)
Had <i>loss</i>		0.092* (0.047)	0.086* (0.047)
<i>Snow</i>	0.16*** (0.023)	0.17*** (0.022)	0.17*** (0.022)
Controls	Y	Y	Y
Observations	2090	2083	2083
Adjusted R-squared	0.049	0.102	0.103

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.

Table F2.
The Impact of Primed and Lived Experiences on Cyberattack Estimates,
No High Quality Restriction

	OLS Predicting Index of Cyberattack Estimates		
	ITT		TOT
	1	2	3
<i>ID theft</i> prime	0.18*** (0.051)	0.26*** (0.058)	0.24*** (0.072)
<i>Finance</i> prime	0.14*** (0.050)	0.14*** (0.053)	0.26*** (0.066)
<i>Loss</i> prime	0.12** (0.050)	0.17*** (0.058)	0.33*** (0.073)
Below Median Similarity x <i>ID theft</i> prime			0.044 (0.087)
Below Median Similarity x <i>Finance</i> prime			-0.23*** (0.076)
Below Median Similarity x <i>Loss</i> prime			-0.29*** (0.087)
Had <i>ID theft</i>		0.36*** (0.050)	0.36*** (0.050)
Had <i>finance</i>		0.17** (0.070)	0.17** (0.069)
Had <i>loss</i>		0.071 (0.049)	0.068 (0.049)
<i>Snow</i>	0.16*** (0.018)	0.18*** (0.020)	0.18*** (0.020)
Controls	Y	Y	Y
Observations	2976	2424	2424
Adjusted R-squared	0.052	0.090	0.096

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.

Table F3.
The Impact of Primed and Lived Experiences, Individual Cyberattack Estimates

	OLS Predicting Cyberattack Estimate for Self (z-score)			OLS Predicting Cyberattack Estimate for People Like You (z-score)		
	High Quality Only					
	ITT	TOT		ITT	TOT	
	1	2	3	4	5	6
<i>ID theft</i> prime	0.22***	0.32***	0.31***	0.058	0.092	0.054
	(0.061)	(0.068)	(0.085)	(0.060)	(0.068)	(0.085)
<i>Finance</i> prime	0.19***	0.18***	0.31***	0.079	0.051	0.092
	(0.061)	(0.065)	(0.082)	(0.061)	(0.065)	(0.082)
<i>Loss</i> prime	0.11*	0.15**	0.29***	0.00048	0.038	0.11
	(0.061)	(0.070)	(0.091)	(0.060)	(0.070)	(0.091)
Below Median Similarity x <i>ID theft</i> prime			0.027			0.077
			(0.10)			(0.10)
Below Median Similarity x <i>Finance</i> prime			-0.23**			-0.074
			(0.093)			(0.093)
Below Median Similarity x <i>Loss</i> prime			-0.24**			-0.14
			(0.11)			(0.11)
Had <i>ID theft</i>		0.34***	0.34***		0.38***	0.38***
		(0.063)	(0.063)		(0.063)	(0.063)
Had <i>finance</i>		0.25***	0.25***		0.28***	0.28***
		(0.083)	(0.083)		(0.083)	(0.083)
Had <i>loss</i>		0.050	0.045		0.0052	0.0035
		(0.059)	(0.059)		(0.059)	(0.059)
<i>Snow</i>	0.11***	0.13***	0.13***	0.19***	0.21***	0.21***
	(0.022)	(0.025)	(0.025)	(0.022)	(0.025)	(0.025)
Controls	Y	Y	Y	Y	Y	Y
Observations	2097	1709	1709	2091	1704	1704
Adjusted R-squared	0.047	0.083	0.088	0.045	0.079	0.079

Notes: * denotes $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Controls include age, sex, race/ethnicity, education, income, region. *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.

Table F4.
Similarity and Vividness
 OLS Predicting **Index of**
Cyberattack Estimates OLS Predicting **Reported**
Vividness of Cyberattack
High Quality Only, Treatment-on-Treated, Omitting Control Group

	1	2	3	4
Perceived Similarity of Primed Experience	0.056*** (0.013)	-0.034 (0.051)	0.060*** (0.012)	-0.015 (0.046)
Reported Vividness of Primed Experience	0.062*** (0.021)	0.00061 (0.039)	0.19*** (0.019)	0.14*** (0.036)
Similarity x Vividness		0.016* (0.0087)		0.014* (0.0080)
<i>Snow</i>	0.17*** (0.031)	0.17*** (0.031)	0.034 (0.029)	0.031 (0.029)
Controls	Y	Y	Y	Y
Observations	1172	1172	1161	1161
Adjusted R-squared	0.060	0.062	0.114	0.115

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. The reported vividness of cyberattack measure in Columns 3 and 4 is the z-score of the individual's self-reported vividness with which they imagined a cyberattack. The reported vividness of primed experience measure is the individual's self-reported vividness with which they imagined the primed experience on a 1 – 7 scale. *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.

Table F5.
Simulation and Vividness
 OLS Predicting **Index of Cyberattack Estimates**
High Quality Only

Reported Vividness of Cyberattack	0.28*** (0.022)
<i>Snow</i>	0.15*** (0.022)
Controls	Y
Observations	2014
Adjusted R-squared	0.118

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. The reported vividness of cyberattack measure is the z-score of the individual's self-reported vividness with which they imagined a cyberattack. *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.

Table F6.
Similarity and Responsiveness to Experiences, ITT
 OLS Predicting Index of Cyberattack Estimates
 High Quality Only, Intent-to-Treat

	Pooled		Below Median <i>Snow</i>		Above Median <i>Snow</i>	
	1	2	3	4	5	6
$\bar{S}_i(E_i)$, Total Similarity of Lived Experiences	0.19*** (0.025)	0.19*** (0.025)	0.15*** (0.035)	0.15*** (0.035)	0.23*** (0.037)	0.24*** (0.037)
$S_i(e_p)$, Similarity of Primed Experience	0.14*** (0.025)	0.20*** (0.029)	0.16*** (0.034)	0.19*** (0.040)	0.13*** (0.036)	0.21*** (0.042)
$\bar{S}_i(E_i) \times S_i(e_p)$		-0.061*** (0.018)		-0.039 (0.026)		-0.084*** (0.024)
<i>Snow</i>	0.16*** (0.022)	0.16*** (0.022)	0.36*** (0.080)	0.36*** (0.080)	0.087** (0.044)	0.099** (0.044)
Controls	Y	Y	Y	Y	Y	Y
Observations	2088	2088	1061	1061	1027	1027
Adjusted R-squared	0.114	0.119	0.086	0.087	0.117	0.127

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.

Table F7.
Similarity and Responsiveness to Experiences, No Restriction to High Quality
 OLS Predicting Index of Cyberattack Estimates

	Pooled		Treatment-on-Treated		Above Median	
	1	2	Below Median		<i>Snow</i>	
			3	4	5	6
$\bar{S}_i(E_i)$, Total Similarity of Lived Experiences	0.17*** (0.023)	0.18*** (0.023)	0.13*** (0.031)	0.14*** (0.031)	0.22*** (0.033)	0.23*** (0.033)
$S_i(e_p)$, Similarity of Primed Experience	0.13*** (0.022)	0.21*** (0.027)	0.16*** (0.031)	0.22*** (0.038)	0.11*** (0.033)	0.20*** (0.040)
$\bar{S}_i(E_i) \times S_i(e_p)$		-0.069*** (0.016)		-0.056** (0.023)		-0.082*** (0.022)
<i>Snow</i>	0.18*** (0.020)	0.18*** (0.020)	0.29*** (0.073)	0.30*** (0.073)	0.12*** (0.040)	0.12*** (0.040)
Controls	Y	Y	Y	Y	Y	Y
Observations	2427	2427	1267	1267	1160	1160
Adjusted R-squared	0.112	0.119	0.085	0.088	0.100	0.111

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.

Table F8.
Similarity and Responsiveness to Experiences, Individual Outcome Measures

	OLS Predicting Z-Score of Cyberattack Estimate for Self High Quality Only, Treatment-on-Treated					
	Pooled		Below Median <i>Snow</i>		Above Median <i>Snow</i>	
	1	2	3	4	5	6
$\bar{S}_i(E_i)$, Total Similarity of Lived Experiences	0.18*** (0.028)	0.18*** (0.028)	0.14*** (0.039)	0.14*** (0.039)	0.21*** (0.039)	0.23*** (0.040)
$S_i(e_p)$, Similarity of Primed Experience	0.17*** (0.027)	0.22*** (0.033)	0.21*** (0.038)	0.24*** (0.047)	0.14*** (0.039)	0.21*** (0.048)
$\bar{S}_i(E_i) \times S_i(e_p)$		-0.051*** (0.020)		-0.029 (0.030)		-0.071*** (0.027)
<i>Snow</i>	0.13*** (0.025)	0.13*** (0.025)	0.22** (0.089)	0.22** (0.089)	0.045 (0.049)	0.050 (0.049)
Controls	Y	Y	Y	Y	Y	Y
Observations	1712	1712	873	873	839	839
Adjusted R-squared	0.111	0.114	0.093	0.093	0.108	0.115

Notes: * denotes $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Controls include age, sex, race/ethnicity, education, income, region. $\bar{S}_i(E_i)$ 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 1 ft of snow in a typical year.

Table F8b. Similarity and Responsiveness to Experiences, Individual Outcome Measures

OLS Predicting Z-Score of Cyberattack Estimate for People Like You High Quality Only, Treatment-on-Treated						
	Pooled		Below Median <i>Snow</i>		Above Median <i>Snow</i>	
	1	2	3	4	5	6
$\bar{S}_i(E_i)$, Total Similarity of Lived Experiences	0.17*** (0.028)	0.18*** (0.028)	0.11*** (0.039)	0.12*** (0.039)	0.22*** (0.040)	0.23*** (0.040)
$S_i(e_p)$, Similarity of Primed Experience	0.070** (0.028)	0.12*** (0.034)	0.093** (0.038)	0.13*** (0.047)	0.052 (0.040)	0.13*** (0.049)
$\bar{S}_i(E_i) \times S_i(e_p)$		-0.057*** (0.020)		-0.043 (0.030)		-0.076*** (0.027)
<i>Snow</i>	0.20*** (0.025)	0.20*** (0.025)	0.43*** (0.089)	0.43*** (0.089)	0.12** (0.050)	0.12** (0.050)
Controls	Y	Y	Y	Y	Y	Y
Observations	1707	1707	869	869	838	838
Adjusted R-squared	0.083	0.087	0.050	0.052	0.076	0.084

Notes: * denotes $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Controls include age, sex, race/ethnicity, education, income, region. $\bar{S}_i(E_i)$ 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.

Table 9. The Role of Similarity in Experience Effects

	I	II
ID Theft Prime	0.15** (0.060)	0.13** (0.059)
Financial Prime	0.15** (0.060)	0.14** (0.060)
Lost Loved One Prime	0.054 (0.060)	0.035 (0.059)
Had ID Theft	0.46*** (0.050)	0.13 (0.12)
Had Financial Struggles	0.22*** (0.064)	0.064 (0.077)
Had Loss of Loved One	0.092* (0.047)	-0.063 (0.058)
Weighted ID Theft		0.053***

		(0.017)
Weighted Financial Struggles		0.048***
		(0.014)
Weighted Loss of Loved One		0.072***
		(0.017)
Snow	0.17***	0.16***
	(0.022)	(0.022)
Controls	Y	Y
Observations	2083	2080
Adjusted R-squared	0.102	0.124