Imagining the Future:

Memory, Simulation and Beliefs about Covid

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

June 26, 2022

Abstract

How do people form beliefs about novel risks, with which they have little or no direct experience? We address this question using a 2020 US survey of beliefs about the lethality of Covid. The survey reveals several surprising findings, including most dramatically that the elderly underestimate their own risks, while the young hugely overestimate them. To shed light on the evidence, we present a model in which people selectively and automatically recall past experiences, including those from other domains, and use them to imagine (simulate) the novel risk. In the model, an experience increases perceived risk by making that risk easier to imagine, but decreases perceived risk by interfering with recall of other experiences that may feed imagination. The model accounts for our initial findings, but also yields new predictions based on how non-Covid experiences should shape beliefs about Covid. The model connects average overestimation of unlikely risks with strong disagreement: people with many interfering experiences underestimate risk and are less sensitive to direct experiences with Covid (such as local disease dynamics). We find empirical support for these and other predictions using our survey data on respondents' Covid and non-Covid past experiences.

[1] The authors are from Oxford University, Bocconi University, Harvard University, Bocconi University, and Harvard University, respectively. 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, Thomas Graeber, Spencer Kwon, Dev Patel, Josh Schwartzstein, and Jesse Shapiro for helpful comments.

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?

We argue that in these contexts selective memory plays a key role. When thinking about a new domain, people recall past experiences, including those from different domains, and use them to imagine the novel risk. Memory-based imagination or "simulation," is known to be central for thinking about the future (Dougherty et al. 1997; Brown et al. 2000)[2]. It is related to reasoning by analogy, and entails retrieving and recombining experiences stored in memory (Carroll 1978, Schacter 2007, 2018, Biderman, Bakkour, and Shohamy 2020). Critically, because simulation uses recalled material, beliefs are shaped by well-established regularities of selective recall, specifically similarity, frequency, and interference, in the spirit of Kahana (2012) and Bordalo et al. (2021).

To illustrate, consider how people would assess the following situations. When Amazon was getting started, was it a successor to conventional bookstores, a rapidly declining sector, or the future giant of e-commerce? Is Tesla today part of the unexciting car industry, or of the emerging electric-battery future? Was Donald Trump in 2016 a maverick ready to shake up Washington, or a risk to democracy? In all these instances, we would answer these questions by retrieving from memory both relevant and irrelevant experiences that help us imagine the alternative descriptions. In thinking about

[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.

Tesla, one person may think of competition in the car industry, even with electric cars, the threat of unionization, and the transition to bikes and work from home; another may think about the glorious future of driverless cars and other – perhaps irrelevant – Elon Musk ideas. In thinking about Trump, some imagined stale Washington and populist democracy; others were shocked by his right wing rhetoric. Different people have different ideas come to mind based on their past experiences and on what they do and do not retrieve from memory, and often disagree. We model this process and use it to understand beliefs about Covid risks in the early stages of the pandemic.

Our model is motivated by data on beliefs about Covid risks that we collected from a large sample of U.S. residents in three waves: in May of 2020, two months after the pandemic had started in the US, in July 2020 and November/December 2020. From the first wave we documented three facts (Bordalo et al. 2020), which we confirmed in subsequent waves. First, people who estimate a higher share of Americans who have red hair are also sharply more pessimistic about Covid, pointing to a person-specific tendency to overestimate unlikely events across domains. Second, there is a striking age gradient: older people are *less* pessimistic about Covid's lethality than younger people. The young overestimate the probability they die if infected with Covid, the elderly underestimate it. Why would this group exhibit systematic underestimation? Second, people are more pessimistic if they experienced non-Covid health adversities such as own or a family member's recent hospitalization. This holds when assessing the risk of Covid death for themselves but also for others, who would have not had such experiences. This effect is quantitatively large and puzzling on two counts. First, one might have expected exposure to other health risks (and surviving them) to lower one's estimated Covid lethality. Second, why would own experiences shape beliefs about others?

To explain these facts, Section 3 presents a model in which people form beliefs by spontaneously retrieving from memory either statistical information about Covid's lethality they may have seen in the media, or past experiences. These experiences, which may or may not be related to Covid, are then used to simulate Covid deaths. Simulation is easier with an experience more similar to a Covid death. People differ in the extent to which they rely on experiences (as opposed to statistics) and in the frequency with which different experiences are stored in their memory database.

In the model, beliefs have three properties. First, simulation encourages an overestimation of unlikely novel risks such as Covid lethality, because even events that have not been experienced can be imagined on the basis of past experiences. This tendency coexists with possibly vast disagreement due to different experience databases.

Second, exposure to an experience affects beliefs due to a trade-off between simulation and interference. This trade-off implies that experiences that are similar to a Covid death – most immediately past Covid deaths but also experiences with other severe diseases – help simulate Covid lethality, acting as sources of Covid pessimism. In contrast, less similar but possibly frequent events such as non-health adversities the person may have experienced (e.g. working in a dangerous job), interfere: they may come to mind and block the recall of experiences that are better suited to simulate Covid deaths. As a result, they act as sources of Covid optimism.

Third, due to interference, the effect of a given Covid or non-Covid experience is not the same across people, but depends on the other experiences in the database. In particular, exposure to a source of Covid pessimism, such as a non-Covid health adversity, should diminish the effect of exposure to another source of pessimism such as direct experiences of Covid death. This occurs because recall of one type of experience blocks recall of the other: people worry about one thing at the time. Even Covid-related experiences interfere with each other; information does not fully aggregate.

In Section 4 we test these predictions using data on health and non-health adversities from surveys 2 and 3. We show that people who experienced more health adversities are more pessimistic about Covid, consistent with simulation based on similar experiences. In contrast, people exposed to more non-health adversities are more optimistic, consistent with interference from less similar experiences. We also show that exposure to a specific source of pessimism such as a non-Covid health adversity weakens the effect of exposure to Covid deaths, and vice-versa. In shaping beliefs, experiences interfere with each other.

These findings offer a unified explanation for the age gradient and the red hair estimate. In our model, older age stands for more experiences, and thus more interference: the database of the elderly is flooded by non-Covid experiences and adversities (which they have survived). These interfere with the retrieval of more similar Covid experiences and reduce simulation based on them. For the young, Covid is completely novel and scary, and faces little interference. In turn, a respondent's "red hair" estimate stands for their reliance on experiences (as opposed to statistics) and hence on simulation, which creates overestimation of unlikely events across domains. Based on this interpretation, our model makes two new predictions. First, due to stronger interference, the beliefs of the elderly should be less sensitive than those of the young to any specific past experience. Second, people who overestimate red haired Americans should be more sensitive to all experiences, both those that increase pessimism and those that reduce it. The data supports both predictions.

A vast body of social science research has documented the effect of past experiences for beliefs and decisions (e.g., Weinstein 1989). Insightful work in economics links individual experiences to insurance demand (Kuhnreuther 1978) and IPO investing (Kaustia and Knupfer 2009), political experiences to the demand for redistribution (Alesina and Fuchs-Schundeln 2007), and macroeconomic experiences to stock market participation and inflation expectations (Malmendier and Nagel 2011, 2016). These experience effects are "domain specific" (Malmendier 2021): they affect beliefs about the domain they concern directly. In our model, these domain specific experience effects are a natural by-product of basic regularities of human memory. Crucially, our model implies that irrelevant experiences also matter, in two key ways. First, due to simulation, the irrelevant experiences that are most similar to the current domain help imagine it. Second, due to interference, some irrelevant experiences can block recall of relevant (or similar) ones, moving beliefs in the opposite direction. Selective memory implies that the effect of an experience depends on other experiences in the database. This explains why, when facing the same novel Covid shock, different people formed radically different beliefs based on their non-Covid experiences.

Our model unifies an average tendency to overestimate unlikely risks with strong disagreement among people. Models of overestimation of unlikely events, such as Kahneman and Tversky (1979) either neglect the possibility of underestimation, or attribute it to noise or uncertainty (Enke and Graeber 2022, Kaw et al. 2020). These models cannot explain why a group of people, the elderly in our case, should predictably underestimate an unlikely risk. In our model, the tendency to overestimate rare events is not driven by a mechanical adjustment but by simulation and interference. The latter, in particular, entails heterogeneity of beliefs across groups – including the underestimation of rare events – driven by systematic differences in their databases.

Work on attitudes toward Covid focuses on the media and political affiliation (e.g. Allcott et al 2020, Bursztyn et al 2021). We measure political views and media consumption in surveys 2 and 3. Like the earlier work, we find that these help explain behaviour and policy preferences, but leave the belief patterns we focus on unexplained. We thus focus on cognitive factors in our analysis.

We continue the program of unifying different belief biases based on selective memory. Theoretically, our innovation is to incorporate simulation from memory and for (differential) reliance on past experiences. Simulation is consistent with the "analogical" reasoning of case-based decision theory (Gilboa and Schmeidler 1995). Crucially, in our model analogical mechanisms operate under the constraints of human memory, which is spontaneous and subject to interference from irrelevant events. We document these effects in belief formation about a major event, rather than in abstract laboratory experiments, as in Bordalo et al. (2021b), Enke et al. (2020), and Andre et al. (2021).

Our paper introduces into economic models simulation from memory, representations of the future based on both relevant and irrelevant experiences that spontaneously come to mind. We did not hypothesize that simulation is at work before running the survey. Rather, we ran the survey to find basic facts about Covid beliefs, and obtained surprising results, such as the pessimism of the young and the optimism of the old. We then developed the theory and tested its additional predictions as a way to explain the puzzling data.

2. The survey and the main facts

2.1 The survey

We ran three surveys, in May, July and November/December 2020, collecting a total of 4525 responses. We partnered with Qualtrics to collect the data, imposing sample quotas to ensure ample representation across age, race, gender, region, and income. Each survey consists of several blocks of questions measuring beliefs, experiences, demographics, and preferences and behaviour. Appendix B reports the survey instruments and details about sample requirements and quotas, question order, payments, and quality controls.

Beliefs about Covid-19 Risks. Our key outcome variable of interest is the believed Covid fatality rate (*FATALITY*) for the general US population, for which there are clear benchmarks. We elicit this belief in terms of the distribution of *FATALITY* along three demographics: age, race, and gender. 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." Respondents must assess, within each category, how many of these 1000 people will die from Covid. For age, participants consider 1,000 Americans in each of three groups: under 40 years old, between 40 and 69 years old, and 70 and older. For the race category, they consider 1,000 White, Black, Asian, and Latinx. For the gender category, they consider 1,000 men and women. Our measure of believed fatality risk for others averages these 9 estimates for each individual. We equally weight groups, but results are very similar if we weight by the share of Americans in each category. [3]

We also ask respondents to think about 1,000 people "very similar to you (in terms of age, gender, race, socioeconomic status, zip code, health status, etc.) who will contract Covid-19 in the next 9 weeks." We then ask "of these 1,000 people, how many do you believe will pass away due to Covid-19?" The answer measures respondents' beliefs about *FATALITY* for themselves. It reflects

[3] Specifically, this is the average of three estimates: one averaging beliefs for males and females, another averaging beliefs for three age groups (0-39; 40-69; 70+), and yet another averaging beliefs for four race groups (White; African-American; Asian-American; Latinx-American).

person-specific pessimism and vulnerability to Covid. We also elicit, using the same wording, beliefs about the number of Covid hospitalizations, conditional on infection, and the number of Covid infections for people like themselves. Appendix C reports the main patterns obtained for these outcomes, which are qualitatively similar, but in our main analysis we focus on *FATALITY*.

Experiences. The second block of questions measures experienced adversity. In all survey waves we asked whether respondents – and separately, a family member – have been hospitalized for non-Covid related reasons in the last year. Given the explanatory power of these measures in survey 1, in waves 2 and 3 we added an array of new measures. We asked participants to assess on a 1 – 7 scale the extent to which they agree with the statement: "Over the course of my life, I've experienced significant adversity." We then follow-up with questions about specific experiences: 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 participants whether they have had Covid, and about indirect experiences, namely whether they know someone who had Covid, was hospitalized with Covid, or died from Covid.

Sociodemographic Characteristics. At the beginning of the survey, to obtain a stratified sample, all participants report: year of birth, gender, race (White, Black, Asian, Latino/a), approximate annual household income, and region of the country where they live (Northeast, South, Midwest, West). At the end of the survey we also collect data on the respondents' health experiences, asking whether they have been diagnosed with conditions believed (at the time) to increase vulnerability to Covid: diabetes, heart disease, lung disease, hypertension, obesity, cancer, or another serious immunocompromising condition. We also ask about whether they have been unemployed in the last nine weeks, their state of residence, whether the current place of residence is urban, suburban, or rural; educational attainment; and whether they live with children or the elderly.

The red hair question. 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 was included as a quality control and to familiarize respondents with the question format,[4] but it more generally proxies for one's tendency to overestimate a cued rare event. As such it plays an important role in our analysis.

Preferences and Behavior. We also ask respondents about their behavioural responses to the pandemic and their policy preferences. We ask how soon they believe "stay at home" measures should be lifted, and whether they would resume their normal activities if these measures were lifted today. We ask about avoidance of emergency medical care, and whether they have avoided filling prescriptions, doctor's appointments, or other forms of medical care in the last few weeks. In waves 2 and 3 we ask 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). We also ask participants their political preferences (Republican vs Democrat) and their consumption of news about Covid. This is not our main focus, but we analyze behaviour and politics in Section 5.

2.2 Basic Facts

We document the basic patterns in the data and the puzzles that emerge from them. Figure 1 reports the frequency distribution of estimated *FATALITY* for self and others, restricting to the participants who reported an estimate below 1000 (i.e. below 100%). The vertical blue and red bars report the median and the mean, respectively. The small blue bars mark the interquartile range.

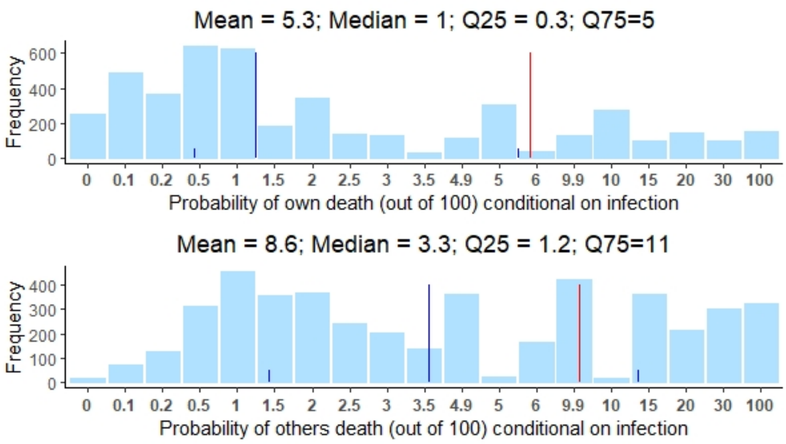

[4] Only participants who estimated that fewer than 1,000 out of 1,000 Americans had red hair could continue in the survey. In addition, participants' answer to the "out of 10,000" question had to be 10 times their answer to the "out of 1,000" question in order to continue in the survey. Other quality controls are described in Appendix B.

Figure 1

The top (resp. bottom) panel reports the distribution of *FATALITY* estimates for self (resp. for others), namely the estimated the number of people, out of 1000 people like self (resp. for others), infected with Covid who will die in the next 9 weeks. For beliefs about others, 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 as described in footnote 3. Ticks on the x-axis refer to the upper limit of the interval.

Two facts stand out. First, there is a systematic overestimation of *FATALITY* from Covid, especially when thinking about others. Median estimates for self and others are at 1% and 3.3%, respectively, mean estimates are at 5.3% and 8.6%. Conventional scientific estimates of *FATALITY* at that time were about 0.68% (Meyerowitz-Katz and Merone 2020). Modal estimates, at about 1%, are quite close to this benchmark, suggesting that many subjects are well calibrated.

Second, there is large dispersion in individual estimates. The interquartile range of believed risks for self is [0.3%, 5%]. This range may not reflect disagreement but rather differential individual vulnerability based on age, health conditions, etc. Large disagreement is however evident in believed risks for others, with a [1.2%, 11%] interquartile range. Disagreement, in the form of a large mass of very pessimistic subjects, is responsible for the average overestimation of this risk.

Where do average pessimism and disagreement come from? In survey 1 (Bordalo et al. 2020) we documented an important role for: i) a respondent's tendency to overestimate rare events, as proxied by the estimated share of red haired Americans, ii) experienced health adversities as measured by personal health conditions and non-Covid hospitalizations, and iii) demographics such as race, income, and especially the respondent's age. Another plausible source of pessimism is the severity of local pandemic conditions. Due to limited variation, we could not reliably assess this factor in the first wave, but we could in waves 2 and 3. We use publicly available state-level data to compute the level of deaths and infections in the respondent's state at the time of taking the survey, their recent weekly growth, their level and growth rates at the time the growth hits its peak, and the days that have passed since the peak.[5] Table B.1 in Appendix C describes these covariates.

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

Table 1 assesses the explanatory power of these factors in all three waves. To assess the robustness of our findings, we use in this and other tables standard methods (Guyon and Elisseeff, 2003; James et al., 2013, see Appendix D for details) to select controls from our entire dataset. We estimate all possible regressions, including all combinations of control variables, and select the specification that fares better in minimizing different information criteria. Details of this procedure are in Appendix D. After presenting the model, we introduce theoretically justified regressors but keep the statistically selected controls to make sure that our theoretical predictions are robust.

The selection criterion picks three demographics besides age: income, race and whether the respondent lives in a rural area. Because these are not tightly interpretable in our theory, we omit them from the tables.[6] Column (1) reports a multivariate regression for beliefs about own *FATALITY*, column (2) reports beliefs about others. Except for dummy variables, all covariates are standardized to render coefficients comparable.

Table 1

The dependent variables are *FATALITY* estimates for self and others, as defined in the text. All variables are standardized except for dummy variables (Hosp self; Hosp fam; Black; Asian; Rural). Red hair is the belief of the respondent about the share of Americans with red-hair. State 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 $t-7$). Days since 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. No. of health conditions takes values from 0 to 7 and counts the number of health conditions of the respondent among the following: 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. The controls are the remaining selected variables (Income, dummy for being Black and dummy for living in a Rural area for Column 1, Income, Black, Rural and dummy for being Asian for Column 2). The number of observations may differ across Columns because sample truncation (e.g. removing subjects who give estimates of death above 1000) is done at the regression level.

	Dependent variable:	
	Risk of Own death	Risk of Others death
	(1)	(2)
Age	-0.131***	-0.236***

[6] Income is a source of optimism; being black, living in a rural area, or being Asian are sources of pessimism (the latter only for others). These results may be interpreted as reflecting experiences, but they may also have other explanations.

	(0.019)	(0.015)
Red hair	0.163***	0.155***
	(0.032)	(0.019)
State Level	0.037**	0.073***
	(0.015)	(0.014)
Days since Peak	-0.057***	-0.084***
	(0.013)	(0.015)
No. health cond.	0.090***	0.032***
	(0.015)	(0.011)
Hosp (self.)	0.245***	0.231***
	(0.078)	(0.062)
Hosp (fam.)		0.093***
		(0.036)
Constant	-0.084***	-0.103***
	(0.022)	(0.022)
Controls	YES	YES
Observations	4,514	4,477
Adjusted R^2	0.071	0.120
Note:		*p<0.1;**p<0.05;***p<0.01 Clustered standard errors at state level

The key findings of survey 1 are robust. First, there is a striking age effect: older people are sharply less pessimistic about Covid risks for both self and others. This result holds despite widespread awareness of the lethality of Covid for the elderly in waves 2 and 3. Second, greater estimated share of Americans with red hair is associated with greater Covid pessimism. Third, current and past non-Covid health adversities raise pessimism. The fact that non-Covid health conditions and hospitalizations increase pessimism about self in column (1) may simply reflect greater vulnerability to Covid by sick respondents. Remarkably, though, these same proxies also raise respondents' pessimism about others in column (2).

Fourth, and this is a new finding relative to Survey 1, Covid experiences matter. "Level" measures the cumulative number of deaths in a state at maximal weekly case growth. Respondents exposed to more severe local Covid conditions, higher "Level", are more pessimistic. This effect fades over time: if the peak occurred longer ago (so “Days” are higher), pessimism is lower. Bayesian belief formation would require that respondents learning from local conditions estimate *FATALITY* by dividing the number of Covid deaths by the number of Covid infections in their state (or by the state’s population as a rough proxy for the latter). However, while more deaths (higher “Level”) boost pessimism, the number of infections or population does not reliably affect beliefs, so infections are not selected by our method. We later argue that our model can account for this fact.

In surveys 2 and 3 we also measured respondents’ political affiliation. Left-wing respondents are a bit more pessimistic about *FATALITY* than right wing ones but the effect is weak and disappears when controls are added, so political affiliation never gets selected as a predictor of beliefs (as we show in Section 5, political affiliation is instead an important determinant of policy views). Our results are robust to including political affiliation in the regressions.

What do these findings tell us about theories of belief formation? The role of “Level” is consistent with standard domain-specific “experience effects” (Malmendier 2021), for it stresses the influence of local Covid death experiences on beliefs and their gradual fading over time. The role of the “red hair” proxy is consistent with a general insensitivity to objective probabilities, and hence a tendency to overestimate unlikely events (Kahneman and Tversky 1979). Such a tendency may be stronger for specific respondents, perhaps because they are more uncertain (Enke and Graeber 2022), or they have noisier numerical perception (Kaw et al. 2020). These effects could be amplified by the ambiguity about Covid risks prevailing in 2020 (Abdellaoui et al 2011).

At the same time, Table 1 raises two key challenges to existing theories. The first is the striking age gradient. As shown in Figure 2 below, the 18-30 age group reports a mean *FATALITY* for self of 8% (median 2%). This is a huge overestimation compared to the true COVID fatality rate for this group, which is 0.01%. On the other hand, the 69+ age group reports a mean *FATALITY* for self of 3.6% (median 1%). This is a substantial underestimation compared to the true infection fatality rate for this group, which is 4.6% (Levin et al. 2020). The elderly underestimate their own risk, contrary to a general tendency to overestimate unlikely events. The age gradient is so strong that it produces the strikingly counterfactual finding that the young believe that their own *FATALITY* is higher than what the elderly believe for themselves. The fact that disagreement in Figure 1 may be due to systematic over- and underestimation of probabilities is challenging for standard theories.[7]

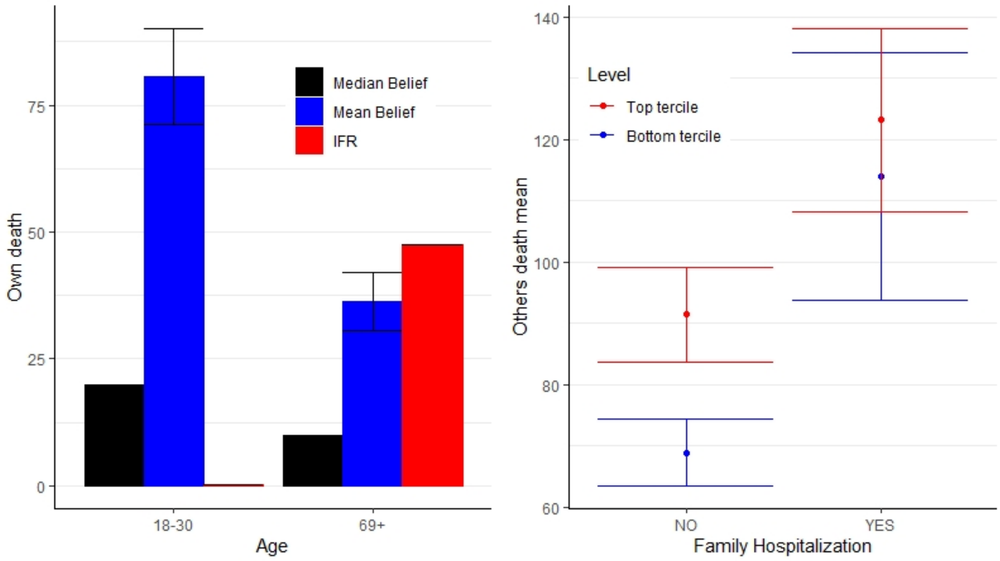

Figure 2

The left panel reports median and mean estimates of *FATALITY* (self) in the lowest and in the highest quintiles of age. 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). The right panel reports estimated *FATALITY* (others) with 95% confidence intervals. Data are split based on the respondent having had a family member hospitalized in the last year (not for Covid) and being in a State in the bottom or top tercile of Covid deaths.

The second challenge raised by Figure 2 concerns non-Covid health adversities. Bad personal health naturally affects beliefs about oneself, but personal and vicarious non-Covid health adversities also raise pessimism for risk facing *others*. In Figure 2, the effect of non-Covid hospitalizations of a family member (results are similar for self-hospitalization) is economically larger and statistically indistinguishable from that of moving from few local Covid deaths (bottom tercile of “Level”) to many (top tercile of “Level”). Having a family member hospitalized for a non-Covid reason dramatically raises pessimism about risks facing *others*.

[7] Heimer et al. (2019) also find that the young are overly pessimistic about their life expectancy while the old are overly optimistic, a fact they explain by the tendency of the young to focus on unlikely causes of death and that of the old to focus on likely diseases. This cannot explain our findings because here the young and the old focus on the same disease.

This finding is puzzling on two counts. First, it cannot be reconciled with standard experience effects, which are domain specific. In this approach, experiences in one setting, such as stocks, affect beliefs in that same setting, but not in similar and even correlated settings such as bonds (Malmendier 2021). Second, it is conceptually not obvious why having experienced non-Covid health adversities should be associated with more pessimism about Covid *FATALITY*, as opposed to encouraging a more "relaxed" attitude toward Covid. These experiences may in fact cause respondents to think that there are so many other health risks (that they survived!).

To shed light on these facts, we present a model of belief formation based on the psychology of memory. When thinking about *FATALITY*, people try to imagine Covid deaths using material retrieved from memory. Experiences in the memory database that are similar enough to Covid deaths or that occur frequently enough compete for retrieval, consistent with the well-established roles of similarity, frequency and interference in memory research (Kahana 2012). If – in the process of recall – a person retrieves experiences that help imagine Covid deaths, then he is pessimistic. If this person retrieves experiences that do not help imagination, he is optimistic.

Imagination based on episodic memory, which psychologists call "simulation", is known to be central for thinking about the future and to form beliefs (Dougherty et al. 1997; Brown et al. 2000). People use past experiences to simulate new ones (Hassabis et al. 2007a,b, Schacter et al. 2012), and the ease of memory-based simulation increases with the similarity between the events (Woltz and Gardner 2015). Events that are easier to simulate are judged to be more likely (Dougherty et al. 1997, Kahneman and Tversky 1981). Simulation is especially important for thinking about new shocks such as Covid, for which people might have few direct experiences.

Our model accounts for the findings documented in this section but also yields new predictions, which we test using the richer measurement of experiences in survey 2 and 3.

3. The model

The Decision Maker (DM) has a database that contains two types of information. The first type is statistical, captured by an estimate π of Covid's *FATALITY*, acquired through news or experts. When our surveys were conducted, the prevalent value of π was on the order of 1-2%, which we take to represent the "correct" assessment and for simplicity to be the same across people.

The second kind of information is a set E containing the DM's episodic memories. These are the DM's life experiences, pertaining to oneself, one's social circle, but also learned from the media. Some experiences concern Covid cases, fatal and non-fatal. Other experiences concern non-Covid health problems, some of high risk (heart attacks), others not (flu). Still other experiences are non-health adversities, such as working in a dangerous occupation or experiencing personal, financial, or other problems. E differs across DMs because of their different life experiences.

The DM assesses *FATALITY* by randomly sampling his database. When thinking about the event of death from Covid, with probability $1-\theta$ the DM samples the statistic π and reports its value. With probability θ the DM samples experiences in E and uses the recalled data to simulate death from Covid. The easier it is to do so, the higher the estimated *FATALITY*.[8] Parameter θ thus captures the DM's reliance on experience. We next formalize recall from E and simulation.

3.1 Recall and Simulation

In line with memory research (Kahana 2012), sampling from E is shaped by similarity and interference: experiences more similar to the cue "death from Covid" are more likely to be retrieved, and recall of these experiences inhibits recall of less similar ones.

Formally, a symmetric function $S(u,v): E \times E \to [0,1]$ measures the similarity between experiences u and v in the database. It increases in the number of features shared by u and v, and is maximal, equal to 1, when $u = v$. A Covid death is very similar to one from SARS, less similar to

[8] As in Bordalo et al. (2021), we can view belief formation as a process whereby the DM draws T samples, each of which contains a statistic or an experience, and the beliefs in Equation (3) are an average across these samples. Compared to Bordalo et al (2021), the novelties here are to allow for simulation (and in particular for differential reliance of beliefs on simulation), and to study belief heterogeneity due to different databases E.

one from a heart attack, and least similar to a death from homicide. Indeed, Covid and SARS are lethal respiratory diseases; heart attacks are not respiratory, and homicides are not diseases. Relative to non-lethal events, a Covid death is most similar to non-fatal Covid, then to infectious or respiratory illnesses (flu or pneumonia), and finally to non-health problems. Similarity also captures recency: Covid deaths experienced further in the past are less similar to very recent ones because they occurred in a different context (Kahana 2012). Our theoretical analysis relies on general intuitions about similarity, which can however be formalized using a features-based similarity function. In the empirical work, we elicit similarity through a survey.

An event such as "Covid death" describes a set of experiences in E sharing two features: 1) they are Covid infections, and 2) they are lethal. We define the similarity between two sets $A \subset E$ and $B \subset E$ as the average pairwise similarity of their elements,

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

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

Based on Equation (1), define $S(e) \equiv S(e, Covid\ death)$ as the similarity between experience e and the event-cue "Covid death".

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

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

From the numerator of (2), experience $e \in E$ is sampled more frequently when it is more similar to a Covid death. When thinking about the probability of dying from Covid, due to similarity we are likely to recall Covid deaths in the news or those of acquaintances.

The denominator of (2) captures interference: all experiences $u \in E$ compete for retrieval, and thus may inhibit recall of e. Interference depends on similarity and frequency. Interference in recalling e is particularly strong from experiences that are similar to the cue. Thoughts of Covid deaths may be interfered with by the recall of other respiratory diseases because the latter have high similarity $S(u)$. But events that frequently occur in the database can be recalled and interfere with Covid deaths even if they are fairly dissimilar from them, because their summed similarity in the denominator of (2) is high. Heart attacks or car accidents may come to mind. People with a larger database find it harder to recall a specific experience e due to many interfering experiences.

Interference is a well-established phenomenon in memory research (e.g., Jenkins and Dallenbach 1924; McGeoch 1932; Underwood 1957). It reflects the fact that we cannot fully control what we recall.[9] Interference inhibits the recall of Covid memories, causing even irrelevant memories to influence beliefs. This will play a key role in producing belief heterogeneity.

If the DM samples personal experience $e \in E$, he is able to imagine a Covid death according to the following formalization of simulation.

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

As in Kahneman and Tversky (1981), simulation is easier when the input is more similar to the target, namely when the two have more features in common. It is easier for the DM to imagine a Covid death based on experienced Covid deaths than based on deaths from SARS, because the former are more similar to the target. Yet, SARS is sufficiently similar that it arguably also helps simulate Covid death. Even less similar experiences can work: seeing someone die in a hospital from a non-infectious disease may help simulate Covid deaths. In general, simulation may weight the features of an experience differently than similarity. For instance, deadly diseases may be dissimilar but especially effective at simulating Covid death, while the flu is similar in many respects but because it is not lethal, it may be poor at simulating Covid death. Here we abstract from this possibility.

[9] 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).

When sampling E, the DM recalls experience $e \in E$ with probability $r(e)$, and uses it to successfully simulate a Covid death with probability $\sigma(e)$. On average, then, the share of simulated Covid deaths across all recalled experiences is given by:[10]

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

Equation (3) describes memory-based beliefs. To see its implications, partition the database E into three sets: i) Covid deaths D_C, ii) Covid survivals S_C, and iii) non-Covid $\overline{C}$. The set $C = D_C \cup S_C$ of lethal and non-lethal Covid experiences is the "relevant" domain specific information.

As a benchmark, suppose that the simulation function is "narrow": the DM perfectly simulates future Covid deaths based on experienced Covid deaths, while simulation fails based on other experiences ($\sigma(e) = 1$ for $e \in D_C$ and $\sigma(e) = 0$ for $e \in E/D_C$). Suppose in addition that similarity is also "narrow": the similarity of Covid experiences to "Covid deaths" is maximal, that of non-Covid experiences to "Covid deaths" is nil ($S(e) = 1$ for $e \in C$ and $S(e) = 0$ for $e \in \overline{C}$). In this knife edge case, the memory-based estimate is frequentist:

$$\hat{\pi}_E = \frac{|D_C|}{|C|}. \tag{4}$$

If the "Covid database" is unbiased, so the relative numerosity of Covid deaths and survivals coincides with that in reality, the average experience-based estimate is identical to the estimate π based on statistical information.

In reality, though, similarity is not narrow. Covid experiences share features with non-Covid ones, such as other diseases or adversities. This tends to raise the denominator of Equation (3). If – at one extreme – similarity were constant, the experience-based estimate $\hat{\pi}_E$ would equal the relative frequency of Covid death experiences in the database E (i.e., $\hat{\pi}_E = \Pr(D_C|E)$). This would create a strong underestimation of a new shock such as Covid, due to interference from many experiences that

[10] Equivalently, every retrieved experience gives rise to a simulation either of a Covid death, with probability $\sigma(e)$, or to itself, with probability $1 - \sigma(e)$. Then $\hat{\pi}_E$ is the share of simulations that produced Covid deaths.

are bad at simulating Covid deaths. But simulation is often also not narrow, and is in fact much broader than standard experience effects. Seeing images of Covid patients laying in ICU beds, or even rough Covid cases, encourages simulation even absent any Covid deaths, as do other experiences with disease. This tends to raise the numerator of Equation (3), promoting overerestimation.

3.2 Memory Based Beliefs

To see the implications of our model, recall how beliefs are formed. With probability $(1-\theta)$ the DM samples statistical information and reports π as his assessed *FATALITY*. With probability θ he samples personal experiences E and uses simulations to estimate *FATALITY*. In a population with a common database E and reliance on simulation θ, the average assessment is:

$$\hat{\pi} = (1-\theta)\pi + \theta\hat{\pi}_E, \tag{5}$$

which combines the statistical "truth" π with the experience-based estimate $\hat{\pi}_E$. *FATALITY* is overestimated on average when $\hat{\pi}_E > \pi$ and underestimated otherwise.

To see when over and underestimation prevail, suppose that the Covid database E is unbiased. If both the simulation and similarity functions are narrow the average belief is frequentist and corresponds to the statistical benchmark $\hat{\pi} = \pi$. Suppose however that both simulation and similarity are somewhat broad: Covid deaths can be simulated using other experiences, $\sigma(e) = \tilde{\sigma} > 0$ for all $e \in E/D_C$, and non-Covid experiences are somewhat similar to Covid deaths, $S(e) = \tilde{S} > 0$ for all $e \in \overline{C}$. We then get the following result (all proofs are in Appendix A):

Proposition 1 *Suppose that the Covid database is unbiased,* $|D_C|/|C| = \pi$. *If irrelevant experiences are recalled and used to simulate Covid deaths,* $\tilde{S}, \tilde{\sigma} > 0$, *there is* $\pi^* \equiv \pi^*(\tilde{S}, \tilde{\sigma})$ *such that FATALITY is overestimated if and only if its true value is low enough, namely* $\hat{\pi} > \pi$ *if and only if* $\pi < \pi^*$. *If* $\pi < \pi^*$, *FATALITY increases in the DM's reliance on experience,* $\partial\hat{\pi}/\partial\theta > 0$.

Irrelevant experiences exert two conflicting effects. On the one hand, they foster simulation of Covid deaths, which boosts $\hat{\pi}$. On the other hand, they interfere with recall of Covid death experiences, which reduces $\hat{\pi}$. If Covid deaths are rare, in an unbiased database there are few Covid death experiences that can be interfered with. Thus, Covid deaths are simulated based on numerous non-lethal Covid experiences or on other health adversities, causing overestimation. People put positive probability on events they had never seen, provided they are similar to their experience.

This mechanism helps explain two key findings in Section 2. It can account for the overestimation of *FATALITY* in Figure 1 by both the average and median respondent. It also suggests an interpretation of the "red hair" variable as a proxy for the DM's reliance on simulation θ. DMs with higher θ should have a greater tendency to overestimate not only *FATALITY* but also other unlikely events, such as the share of red haired Americans, consistent with Table 1.

The second key finding of Section 2 is that both relevant and irrelevant experiences shape disagreement. In our model, disagreement arises when people have different databases E, through the interaction between simulation and interference.[11] To analyse this interaction, take a subset E_i of experiences sharing some features (e.g. non-Covid adversities), and suppose that we increase its numerosity $|E_i|$, while keeping constant its similarity $S(E_i)$ to the target event. Our model produces two key mechanisms of disagreement, described in Propositions 2 and 3.

Proposition 2 *Increasing the numerosity of the subset* E_i *increases FATALITY,* $\partial\hat{\pi}/\partial|E_i| > 0$, *if and only if* $\hat{\pi}_{E_i} > \hat{\pi}_E$; *that is, if and only if estimated FATALITY is higher when using only* E_i *than when using the full database* E. *In particular, adding a single experience* e *to* E *increases FATALITY if and only if* e *is sufficiently similar to Covid death compared to an average member of* E, $\sigma(e) > \hat{\pi}_E$.

A tradeoff emerges in the model: increasing exposure to an experience boosts Covid pessimism by providing material that helps simulating Covid deaths, but dampens pessimism by interfering with recall of other experiences that may be more effective at simulation. Whether an experience is a source of Covid pessimism or optimism depends on the balance of these effects.

[11] In our model disagreement is also due to "noise", namely to random recall from the database. Random recall cannot however explain systematic belief differences among groups, say young and old, so we do not explore it here.

At one extreme, this tradeoff accounts for narrow domain-specific experience effects. Exposure to local Covid deaths (higher "Level" in Table 1), should boost Covid pessimism: these experiences are maximally similar to the target event and hence maximally suitable for simulating it.[12] But this tradeoff also accounts for the effect of non-Covid experiences. At intermediate degrees of similarity, it may explain why somewhat similar, yet domain irrelevant, non-Covid health adversities such as hospitalization of self and others boost Covid pessimism in Table 1: due to similarity along the health dimension, they help simulate severe Covid cases. According to Proposition 2, such domain irrelevant simulation is at work if the database contains many experiences that are worse at simulating Covid death than the irrelevant health adversities, i.e., if $\hat{\pi}_E$ is low.

Crucially, Proposition 2 also implies that – at the other extreme – exposure to dissimilar experiences that are bad at simulating Covid deaths should reduce Covid pessimism. For instance, being exposed to many adversities not due to personal poor health should reduce Covid pessimism. When thinking about Covid deaths, these adverse experiences come to mind and interfere with recall of better simulation material. In Section 4 we test this prediction using the measurement of experiences from surveys 2 and 3.

Simulation and interference also yield predictions for the interaction between experiences.

Proposition 3 *Increasing the numerosity* $|E_j|$ *of the set of experiences* E_j *influences the marginal effect of increasing the numerosity* $|E_i|$ *of other experiences* E_i *as follows:*

$$\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|} = K_{ij}\left[\left(\hat{\pi}_E - \hat{\pi}_{E_i}\right) + \left(\hat{\pi}_E - \hat{\pi}_{E_j}\right)\right], \quad K_{ij} > 0. \tag{6}$$

Because different experiences compete for retrieval, they interfere with each other. As a result, the marginal impact of an experience depends on other experiences the DM has lived. This is key: the effect of a specific experience is not absolute, it depends on other experiences in the database.

[12] Recency of an experience also facilitates its retrieval, by increasing its similarity to the present moment (Kahana 2012), so all else equal if Covid experiences are more recent the DM is more pessimistic (see the Appendix A for a proof). This mechanism captures the recency effect of "Days" in Table 1.

This principle yields many predictions which we come back to in Section 4. One prediction, directly related to Table 1, says that non-Covid hospitalization of a family member should dampen the pessimism produced by a rise in local Covid deaths, and vice-versa. Intuitively, the experience of hospitalization interferes with recall of news when making probabilistic judgments, and vice-versa. Because it relies on retrieval from memory, the impact of domain specific experiences depends on the entire structure of the database.

Overall, selective memory places a rich structure on how past experiences and information affect a person's beliefs about an event, and on the connection between average belief bias (in our case overestimation) and disagreement in a group of people. From Equation (5), these effects are due to two characteristics: the tendency to rely on experience θ, and the ability to imagine the event based on experience. The latter is captured by the composition of the memory database E, which affects simulation and interference and hence $\hat{\pi}_E$. These characteristics θ, E yield many new predictions, which we now test.

4. Disagreement: Empirical Tests

We first at forces driving the simulation of Covid deaths on the basis of the database E. Section 4.1 tests the trade-off between simulation and interference (Proposition 2): past exposure to adverse health experiences should increase Covid pessimism, while exposure to adverse experiences that are dissimilar from Covid, such as those not due to poor health, should reduce Covid pessimism. Section 4.2 examines interference across experiences (Proposition 3): greater exposure to one source of pessimism, say having a family member hospitalized, should dampen the impact of another source of pessimism, say more severe local Covid conditions, and vice-versa.

We next show that these mechanisms account for the age gradient and for the red hair effect. In Section 4.3 we show that age can be viewed as a proxy for stronger interference in E, namely lower ability to imagine Covid death. Here Proposition 3 makes the prediction, which we test, that the beliefs about Covid of the elderly should be *less* sensitive to any experience they have had. In Section

4.4 we go back to the interpretation of higher red hair estimate as a stronger tendency θ to rely on experience. We test one prediction following from Propositions 1 and 2: the beliefs about Covid of respondents estimating a higher share of red hair Americans should be *more* sensitive to any experience they have had. Finally, we show that, in line with our model, respondents' overestimation of *FATALITY* and their disagreement are systematically connected to their age and red hair estimates.

4.1 The trade-off between simulation and interference

We test this tradeoff by using the finer measurement of past personal adversities from Surveys 2 and 3. Among the new experiences measured in these surveys, two groups of adversities are relevant here. The first consists of health adversities the respondent had in the past, in the form of a "serious illness" or "a serious injury". We construct an index of "Health Adversities" as the sum of these two dummies. Based on the role of non-Covid health adversities in Table 1, we expect these Health Adversities to increase pessimism.

The second group consists of adversities that are related to the respondent having had life difficulties for non-health reasons. These measure 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. We construct an index of "Non-Health Adversities" as the sum of these four dummies. "Non-Health Adversities" are intuitively less similar than "Health Adversities" to Covid risks, because the former do not capture problematic health conditions the person suffered from. As a result, Proposition 2 predicts that such "Non-Health Adversities" should interfere more with simulating Covid deaths, reducing Covid pessimism (or increasing pessimism less than do "Health Adversities").

In surveys 2 and 3 we also measure direct Covid experiences. In particular, we ask whether the respondent "Had Covid." Again, based on the results of Table 1, which suggest that non-lethal personal health conditions help simulate Covid deaths, we expect this experience to boost pessimism.

Intuitively, having had Covid (and perhaps still having to recover from it due to long Covid) is an adverse health condition even more similar to Covid death, helping simulate the latter.

Table 2 tests these predictions. Column (1) reports the regression for *FATALITY* in Table 1, column (2), estimated in waves 2 and 3. In column (2) we add the dummy for whether the respondent Had Covid as well as past "Health Adversities" and "Non-Health Adversities". We also add our "Subjective Adversity" measure, which captures perceived adverse experiences.

Table 2

The dependent variable is *FATALITY* estimates for others, as defined in the text (see footnote 3). All variables, except for 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. Subjective adversity is the rate of agreement with the sentence "Over the course of my life, I've experienced significant adversity." The controls are the remaining selected variables (Income, Black, Asian and Rural). The number of observations may differ across Columns because sample truncation (e.g. removing subjects who give estimates of death above 1000) is done at the regression level.

	Dependent variable:	
	Others death	
	(1)	(2)
Health adversities		0.047^{**}
		(0.019)
Non health adversities		-0.039^{***}
		(0.015)
Had Covid		0.441^{***}
		(0.167)
Subj. adversity		0.043^{**}
		(0.019)
No. health cond.	0.029^{**}	0.012
	(0.013)	(0.017)
Hosp (self.)	0.218^{***}	0.157^{**}
	(0.078)	(0.073)
Hosp (fam.)	0.061	0.058
	(0.045)	(0.044)
Level	0.061^{***}	0.059^{***}
	(0.023)	(0.023)

Days	-0.098***	-0.097***
	(0.024)	(0.023)
Red hair	0.169***	0.165***
	(0.033)	(0.033)
Age	-0.227***	-0.212***
	(0.017)	(0.021)
Constant	-0.114***	-0.128***
	(0.026)	(0.030)
Controls	YES	YES
Observations	2,972	2,953
Adjusted R^2	0.119	0.133
Note:		*p<0.1;**p<0.05;***p<0.01
		Clustered standard errors at state level

Consistent with Table 1 and with Proposition 2, experiencing non-Covid "Health Adversities" boosts pessimism. Crucially, and also consistent with our model, experiencing "Non Health Adversities" goes in the opposite direction, acting as a source of Covid optimism. In our model, this is due to interference: having gone through a bumpy life, characterized by risks related to one's occupation, poverty, or serious problems of a loved one, makes it easier to retrieve risks different from Covid. This reduces the ability to simulate Covid deaths, fostering optimism.

To complement our judgments of similarity, in May 2022 we asked a diverse sample of U.S. residents to rank eight experiences from our original surveys in terms of subjective similarity to a severe Covid outcome. Full details of this supplementary survey are reported in Appendix B. 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 the two additional adverse experiences in Table 1: having experienced a non-Covid hospitalization and having experienced a family member hospitalization. Lower rank means higher similarity. The average rank (in terms of similarity to a Covid fatality) attached to the two components of "Health Adversities" is 3.4. The average rank attached to the four components of "Non-Health Adversities" is 5.11. Consistent with the estimates in Table 2 and with our model, non-health adversities are on average judged to be less similar to the target event than health adversities. As a result, the former set of experiences should be less good than the latter for simulating Covid deaths, dampening Covid pessimism.[13],[14] A weakness of this similarity measurement is that it was conducted two years after we measured beliefs, and after the pandemic had evolved substantially. Future surveys should jointly measure beliefs, experiences, and similarity.

Quantitatively, the effect of Non-Health Adversities is large. The coefficients in Table 2 imply that moving from zero to four Non-Health Adversities is associated with 25 fewer predicted Covid deaths out of 1000 infected. To increase predicted Covid deaths by the same 25 units, the observed number of cumulative deaths in the state (at the peak of weekly case growth) must go from 0 to 17000. This is a large number, given that the maximum number of cumulative Covid deaths at peak in the data is 15700. That is, an otherwise average person who has experienced maximal Non-Health Adversities and is going through a local Covid peak has the same pessimism as a person unaffected by Non-Health adversities and who is experiencing zero local Covid deaths. The effect of Non-Health Adversities can fully offset the role of rising local Covid deaths.

Consider now the effect of "Had Covid". Consistent with our model, in Table 2 personal exposure to Covid is a source of pessimism.[15] In a "rational" world, one may have expected Covid

[13] Ideally, in line with Proposition 1, one would want to compute the experience based estimate $\hat{\pi}_H$ obtained when using only "Health Adversity" and the estimate $\hat{\pi}_{NH}$ obtained when using only "Non Health Adversities". If as an approximation we assume that similarity linearly declines in the rank and that simulation is equal to similarity, we find that in our sample of respondents the average value of $\hat{\pi}_H$ is 0.56 and the average value of $\hat{\pi}_{NH}$ is 0.44. Consistent with Proposition 1, Covid death is better simulated by the average respondent using health than non-health adversities, $\hat{\pi}_H > \hat{\pi}_{NH}$.

[14] In terms of average rank provided, the rank ordering for individual experiences is: serious illness, loss of loved one, accident or injury, family hospitalization, non-Covid hospitalization, dangerous job, poverty, and military service. The ranking is consistent with our classification, except for "serious injury, illness or untimely death of a loved one" which is ranked above a health adversity such as "serious injury". We still include this proxy into our "Non Health Adversities" dummy because, consistent with our original classification, the actual experience of a loss of a parent or partner can entail severe non-health related consequences. The results of Table 2 go through if we omit this variable from the non-health adversities index (Appendix C).

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

survivors to be more optimistic about *FATALITY* than people who did not catch the virus. As we argued, however, simulation predicts the opposite: experience with this disease, especially if rough, can help imagine less lucky or more vulnerable people dying from it. Interference in Proposition 3 further implies that this effect should be especially strong during early stages of the pandemic, when Covid infections and deaths were few. In Section 4.3 we test this prediction of the model.

Going back to Table 1, the role of "Had Covid" can help explain why the number of infections in a state or its population do not reduce *FATALITY* pessimism, even after controlling for "Level" deaths. People seeing many Covid infections have many experiences that help them simulate *FATALITY*, leading to higher estimates regardless of population size. Availability of such simulation material promotes pessimism, causing a departure from the frequentist benchmark in Equation (4).

In sum, measured experiences, both domain relevant and irrelevant, shape beliefs based on their similarity to Covid death, as predicted by the tradeoff between simulation and interference.

4.2 Interference Across Experiences

The second mechanism for belief heterogeneity is interference across experiences. Two people living the same experience react differently to it based on other experiences stored in their memory database. Proposition 3 yields the following result.

Corollary 1. *if E_i and E_j are sources of pessimism, $\hat{\pi}_{E_i}, \hat{\pi}_{E_j} > \hat{\pi}_E$, then Equation (6) implies that higher $|E_j|$ dampens the marginal effect of $|E_i|$ on beliefs, $\partial^2\hat{\pi}/\partial|E_i|\partial|E_j| < 0$.*

Different sources of pessimism should interfere with each other, mutually dampening their marginal effect on beliefs (the same is true for sources optimism). This prediction is key, for it suggests that the effect of an experience, including a domain-specific one, cannot be studied in isolation: it depends on the entire database. For instance, having had a health problem increases pessimism through simulation, but it also interferes with retrieval of another source of pessimism such as local Covid deaths. People worry about one thing at a time.

We next test for the cross interference between the local severity of Covid, as measured by "Level", and other sources of pessimism in the data: 1) the experience of having had Covid, 2) the three personal non-Covid health adversities ("own hospital", "serious injury" and "serious illness"), and 3) the non-Covid health adversity of the respondent's close contacts ("family hospital").

Figure 4 reports the results. Each panel corresponds to the interaction of "Level" with one of the other past health adversities. In each panel, a bin is identified by a tercile of "Level" combined with a degree of severity of the other health adversity on the horizontal axis. Each bin reports the average Covid pessimism in the corresponding sample, measured by the average residual obtained from regressing *FATALITY* on all regressors of Table 2 except for the two variables that define the panel. Darker colours represent higher assessment of *FATALITY* risk.

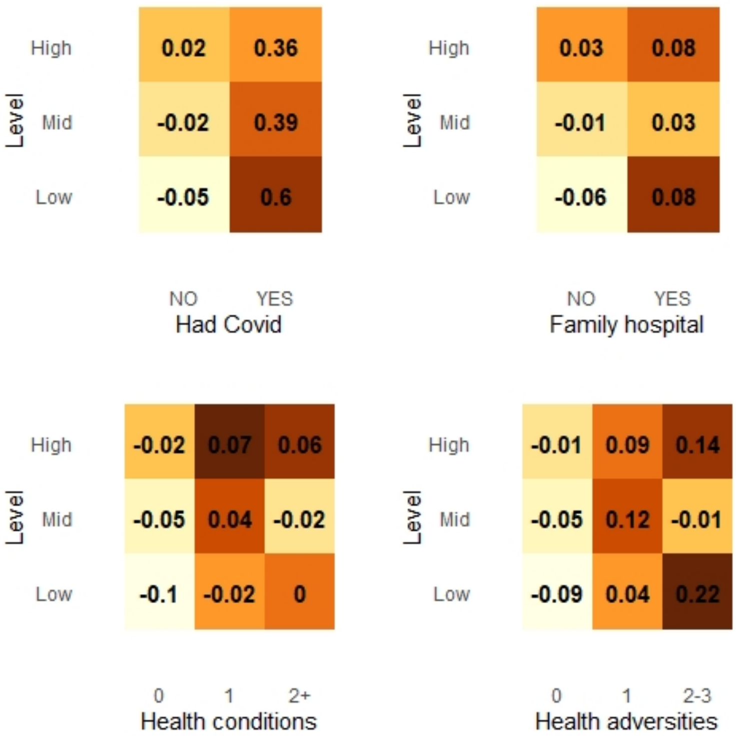

Figure 3.
The Figure reports the residuals of the standardized beliefs of *FATALITY* (for others), estimated by removing from the model in column 2 of Table 2 the variables "Level" and i) "Had Covid" (top left), ii) "Family Hospitalization" (top right), iii) "Number of Health Conditions" (bottom left), and iv) "Health Adversities". Health adversities refer to the sum of serious injury, serious illness, and self hospitalization dummies. Level Low, Mid, High refer to the three terciles of the distribution of State Level deaths for Covid (defined on all waves or on waves 2 & 3, depending on the sample). Reported values are average residuals in each cell. Different colours indicate different average residuals up to the third decimal.

The upper left panel illustrates interference between different Covid experiences. For respondents who have not had Covid, moving from the bottom to the top tercile of “Level” is associated with an increase in pessimism of 0.07 = 0.02-(-0.05) of a standard deviation in beliefs. For respondents who have had Covid, the same change in “Level” is actually associated with a reduction in Covid pessimism. That is, consistent with interference, having had Covid strongly dampens the effect of local lethality experiences measured by “Level”. Consistent with Corollary 1, interference is mutual: “Had Covid” is in fact interfered with by local Covid deaths. The most drastic Covid experience for a respondent is to contract Covid in a state in the bottom “Level” tercile, which is associated with 0.65 standard deviations higher pessimism. Contracting Covid during strong viral transmission (top tercile of “Level”) has a much smaller impact on pessimism.

The upper right panel illustrates interference between Covid and non-Covid experiences, in particular between “Level” and “Family Hospital”. For respondents who have not had a family member hospitalized, moving from the bottom to the top tercile of “Level” is associated with an increase in pessimism of 0.09 standard deviations. For respondents who have had a family hospitalization, the same change in “Level” is actually associated with no increase in pessimism, a strong form of interference of own non-Covid health adversities with “Level”. Own or family hospital experiences boost simulation of Covid, and interfere with local pandemic conditions. Again, interference is mutual, so it also works from “Level” to “Had Covid”: having a family member hospitalized in a state in the bottom “Level” tercile strongly boosts pessimism (by 0.14) while the impact is much smaller when local Covid prevalence is high (top tercile of “Level”).

Interference also holds in the other two panels, which show that higher “Level” reduces the marginal impact of non-Covid health adversities, and higher non-Covid health adversities reduce the marginal impact of “Level”. Visually, the colour gradient is strongest when moving from south west to northwest and southeast, capturing a tendency for a significant Covid or non-Covid health adversity to have a larger marginal impact if it occurs in isolation as opposed to jointly.[16] People worry about one thing at a time.

Overall, the evidence confirms that the effect of an experience is not absolute. It depends on other experiences in the database, and it does so according to the structure of selective recall. When thinking about Covid risks, experiences from other domains come to mind and interfere, dampening the effect of Covid-specific local news.[17]

4.3 The Age Gradient

A key source of belief heterogeneity in Table 1 is age: the elderly are much less pessimistic about Covid risks than the young, in a dramatically counterfactual way. We now show that the tradeoff between simulation and interference can also account for this effect. In particular, the elderly are less able to simulate Covid deaths due to strong interference from other experiences. This explanation is intimately connected to Propositions 2 and 3.

To see the connection with Proposition 2, note that Equation (3) can be rewritten as:

$$\hat{\pi}_E = \frac{\mathbb{E}(\sigma S|C)|C| + \mathbb{E}(\sigma S|\overline{C})|\overline{C}|}{\mathbb{E}(S|C)|C| + \mathbb{E}(S|\overline{C})|\overline{C}|}, \tag{7}$$

where $C \subset E$ is the subset of Covid experiences and $\overline{C}$ is the subset of non-Covid ones.

Obviously, in Line with Tables 1 and 2, Covid experiences are more effective at simulating Covid deaths than non-Covid ones, formally $\hat{\pi}_C > \hat{\pi}_E > \hat{\pi}_{\overline{C}}$. Thus, by Proposition 2, respondents with

[16] In Appendix C we assess interference between all pairs of health adversities (Covid and non-Covid) by running versions of Tables 1 and 2 in which we add the interactions between any two sources of pessimism at the time, and in which we also consider the role of a respondent's current health adversities. The results confirm a broad pattern of interference consistent with the model, whereby the marginal impact of an adversity drops when other adversities are added to the database. To interpret this result, note that the correlation between different Covid and non-Covid health adversities is small. Among the health adversities above, the largest correlation is a 0.16 correlation between "Level" and "Family Hospital". When these variables are orthogonalized with respect to the other controls, their correlation drops to 0.035. These low correlations assuage the concern that the interference detected by the interactive regressions we estimate in Appendix C may be spuriously due to the concave effect of any given health adversity on pessimism.

[17] This also implies that beliefs can overreact not by being excessively sensitive to local conditions, as commonly assumed, but by being excessively sensitive to irrelevant past experiences that activate simulation. A person subject to severe health adversities may become suddenly pessimistic about a new disease such as Covid, but then react little to a growth in local Covid deaths, and the latter may give a misleading impression of underreaction.

a database that is ceteris paribus richer in non-Covid experiences (i.e. having higher $|\overline{C}|$ for given $|C|$) are less able to simulate Covid deaths and thus are less pessimistic. Because Covid is a new shock, the age gradient immediately follows: the database of old people is flooded with many non-Covid experiences. These experience create interference, leading the elderly to be more optimistic.

This account is consistent with memory research, which stresses that the failure to remember specific events is to a large extent caused by a failure of retrieval from the memory database on the basis of cues (Shiffrin 1970).[18] An older person who cannot remember whether they locked the door earlier that day is failing to retrieve the exact event among a vast number of similar events in the past (Wingfield and Kahana 2016). Our model captures interference of this sort. When thinking about Covid deaths older people recall many adversities over the course of their lives, some related to health and some not. These interfere with recalling Covid deaths, promoting optimism.

The second way in which the age gradient arises is interference across different experiences in Proposition 3, which immediately implies the following testable prediction.

Corollary 2. *The beliefs of the elderly should be less sensitive to each experience E_i. In Equation (6), denoting non-Covid experiences as $E_j = \overline{C}$ yields, when $|\overline{C}|$ is sufficiently large that $\hat{\pi}_{\overline{C}} \approx \hat{\pi}_E$:*

$$\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |\overline{C}|} = -\frac{\partial \hat{\pi}}{\partial |E_i|}. \tag{8}$$

If experience E_i is a source of Covid pessimism, $\partial\hat{\pi}/\partial|E_i| > 0$, it will be less so for an older respondent because the latter has a larger database, consisting of many non-Covid experiences $|\overline{C}|$. The latter in fact interfere with any specific experience the person may have had. Analogously, if experience E_i is a source of Covid optimism, $\partial\hat{\pi}/\partial|E_i| < 0$, it will be less so for an older respondent.

[18] There is evidence that over time that memories "physically" degrade, which also causes forgetting. This effect can reduce the size of the database of the elderly compared to what it could have been with no degrading. What we need for our analysis is that such degrading is sufficiently low that the elderly have a larger database of non-Covid experience than the young. Consistent with this, in our data the elderly report having on average experienced a larger number of Health and Non-Health adversities than the young.

To test for the lower sensitivity of the elderly, we estimate separately the specifications of Tables 1 or 2, depending on whether the relevant experience is available for all three waves or not, for older people people (62+) and the rest. Figure 3 reports the estimated coefficients and confidence intervals for non-Covid sources of optimism and pessimism (panel A), and for Covid experiences (panel B), for the elderly (in blue) and the rest (in red). We also assess whether the interference effect of older age exhibits diminishing marginal strength, which is another prediction of Equation (6) above, by adding age squared to the regression of Table 2 (it should have a positive coefficient).

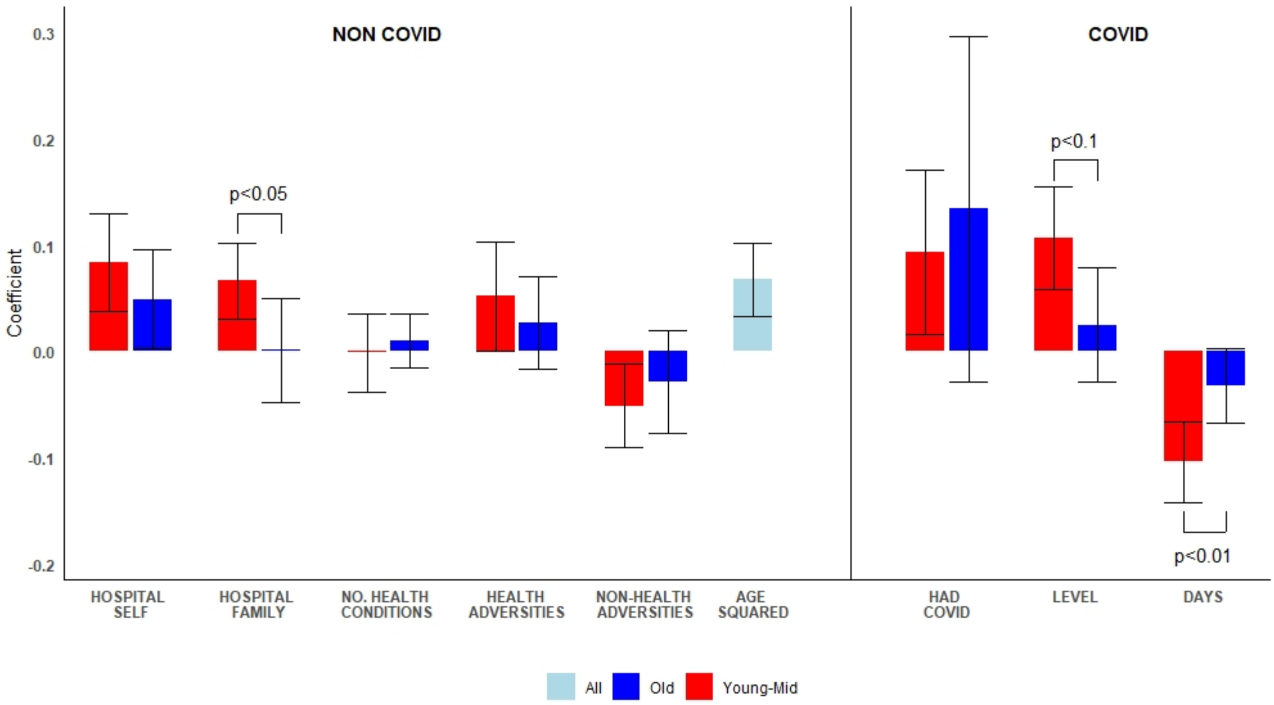

Figure 4.

The figure reports the coefficients obtained by estimating the equations for beliefs of others death in Tables 1 and 2 in the first two terciles of age (18-61) and in the top tercile (62+). Coefficients for variables available in all waves (hospital self, hospital family, no. health conditions, 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. Age squared coefficient is obtained by adding age squared to the model presented in column 2 in Table 1. For the sake of comparability, all variables (including dummies) were standardized.

Consistent with Proposition 1, the elderly's beliefs react less pessimistically to a non-Covid hospitalization of self or a family member, and to health adversities, defined as having had a serious injury or illness in the past. The dampening effect of age also holds for sources of optimism such as non-health adversities: elderly who have experienced poverty or dangerous jobs are less optimistic than younger people who faced the same adversities. The elderly tend to also be less sensitive than the young to Covid experiences: they are not as pessimistic when the level of peak deaths is higher, and their optimism does not rise by much as the peak recedes into the past. Interference modulates standard experience effects. Contrary to Corollary 1, the elderly who had Covid are more pessimistic than the young. This effect, while statistically insignificant, is intuitive: for the elderly Covid is a much more severe disease than for the young, the experience of which must be scary.[19] Also consistent with Proposition 3, the coefficient of age squared is positive.

Overall, the data support the prediction that, due to interference, the elderly are less sensitive to any specific experience. An F-test for the null hypothesis that the coefficients are identical across the age groups is rejected.[20] The elderly are not just insensitive to sources of pessimism, and hence more optimistic. They are less sensitive across the board, which in our model comes from their difficulty of recalling any specific source, due to interference from many other experiences.

In a Bayesian world, older people might react less to news because they have more data, and have less to learn. This would however also imply that as people get older their beliefs should become more accurate, which is not the case in the data. For instance, people in the age group 72+ underestimate own lethality by 2.5%, while those in age group 65-71 do so by 1.7%. A larger problem is that Covid is a new shock, so the elderly and the young should be equally ignorant about it. In our model, the elderly react less to the shock not because they know more, but because their many irrelevant experiences interfere with imagining Covid as a particularly severe mortality risk.

4.4 Red Haired Americans and Reliance on Experience θ

We finally connect the willingness of a respondent to rely on experience, θ, with its ability to use experiences to imagine Covid risks, as proxied also by the specific experiences the person has

[19] This effect arises in our model if having had Covid is more similar to a Covid death for an older respondent. This naturally follows from the similarity function in Equation (1), because the target event "Covid death" is disproportionally composed by the elderly deaths. For simplicity we shut down this effect, which arguably also plays a role in other personal experiences in Figure 3, by considering comparative statics in which we vary the numerosity of experiences E_i while keeping their similarity $S(E_i)$ to Covid death constant across different respondents.

[20] A test on the interaction of age with all variables included in all waves (Table 1, Column 2) gives p = 0.01. A test on the interaction of age with all variables included in waves 2 and 3 (Table 2 Column 2) gives p = 0.00.

lived. Our model makes prediction on how our proxy for θ, the estimated share of red haired Americans, should interact with the determinants of the ability to simulate.

Corollary 3. *The beliefs of people who rely more on experience should be more sensitive to their Covid as well as non-Covid experiences. More broadly, for any experience based factor X:*

$$\frac{\partial \hat{\pi}}{\partial X \partial \theta} = \frac{\partial \hat{\pi}_E}{\partial X}. \tag{9}$$

If "red hair" is a proxy for θ, then respondents who estimate a higher share of red haired Americans should be disproportionally pessimistic if they experience more sources of pessimism, $\partial \hat{\pi}_E / \partial X > 0$, and disproportionally optimistic if they experience more sources of optimism, $\partial \hat{\pi}_E / \partial X < 0$. Simulation creates a link between a respondent's overestimation and a higher weight he attaches to memory based signals. This prediction is inconsistent with the interpretation of our red hair proxy as a general tendency toward insensitivity, due to noise or cognitive uncertainty.

To test this prediction we estimate our baseline specification of Table 1, column 4, but distinguish the top "red hair" tercile from the rest. Figure 5 reports the estimated coefficients and confidence intervals for each one of the relevant covariates in the two "red hair" groups.

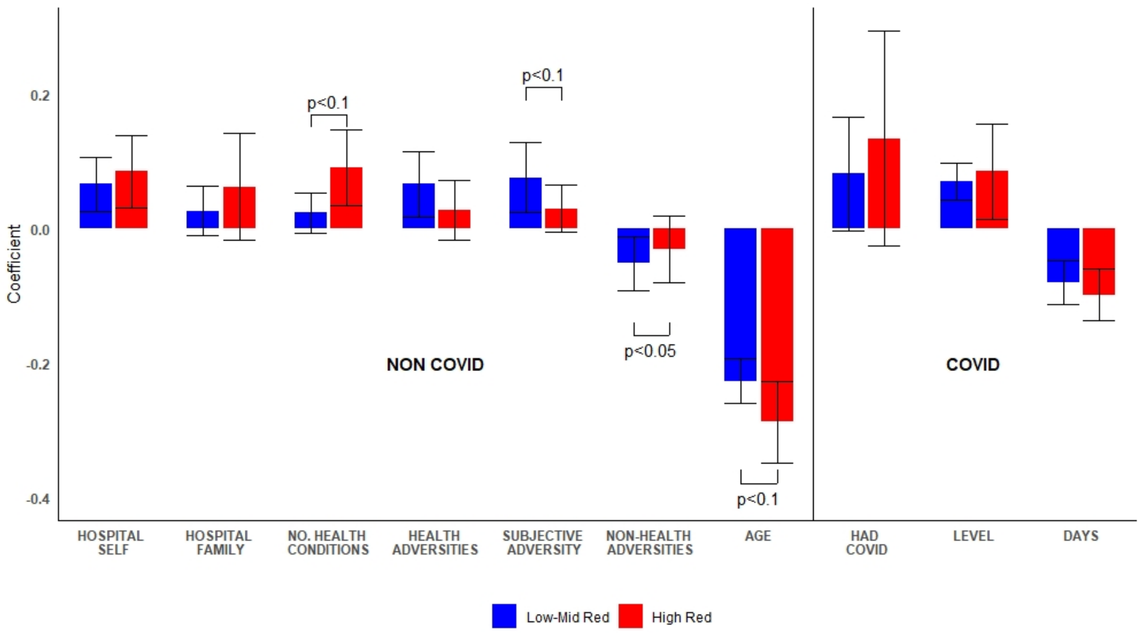

Figure 5

The figure reports the coefficients obtained by estimating the equations for beliefs of others death in Tables 1 and 2 in the first two terciles for red hair estimates (up to 50 out of 1000) and in the top tercile (more than 50).

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, subjective adversities, non-health adversities, had Covid) were obtained by estimating the model from column 2 in Table 2. For the sake of comparability, all variables (including dummies) were standardized.

There is an overall tendency for high "red hair" respondents (in red) to be more sensitive to determinants of pessimism and of optimism than low "red hair" respondents (in blue), consistent with our model. High red hair respondents tend to be more pessimistic than low red hair ones after experiencing non-Covid hospitalization for themselves, a non-Covid hospitalization of a family member, a higher number of heath conditions and subjective adversities, and (directionally) Covid experiences (though no effect is seen in the case of the health adversity proxy).

Crucially, high red hair respondents also react more to factors that promote optimism such as non-health adversities and age. An F-test for the null hypothesis that the coefficients are identical across the red hair groups is rejected.[21] This evidence suggests that the tendency to overestimate rare events is tightly connected to the reliance on personal experiences as opposed to statistical data, lending support to our memory based model.

We conclude by returning to the connection between average overestimation of Covid lethality and disagreement about it, which our model unifies as a consequence of memory based simulation. In our theory, two attributes of a respondent should drive this connection: reliance on experience and ability to simulate Covid lethality, as shaped by interference. People who strongly rely on experience should display stronger overestimation but also larger disagreement based on their different databases. People who face more interference should display weaker average overestimation and less disagreement due to the difficulty of recalling any specific element of their database.

Figure 6 tests for this prediction using our proxy of reliance on experience, the "red hair" answer, and our proxy for interference, a respondent's age. Note that in our data the correlation

[21] A test on the interaction of red hair with all variables included in all waves (Table 1, Column 2) gives p = 0.06. A test on the interaction of red hair with all variables included in waves 2 and 3 (Table 2 Column 2) gives p = 0.03.

between age and red hair estimate is low, equal to -0.09, so these two are largely independent sources of variation. In the top panel, we split our sample in septiles of red hair. In the bottom panel, we split it into septiles of Age. Each panel first reports the median estimate of *FATALITY* and the interquartile range for the full sample, followed by the median beliefs and interquartile ranges of the samples obtained by removing septiles 1 through 6, as indicated in the x-axis.

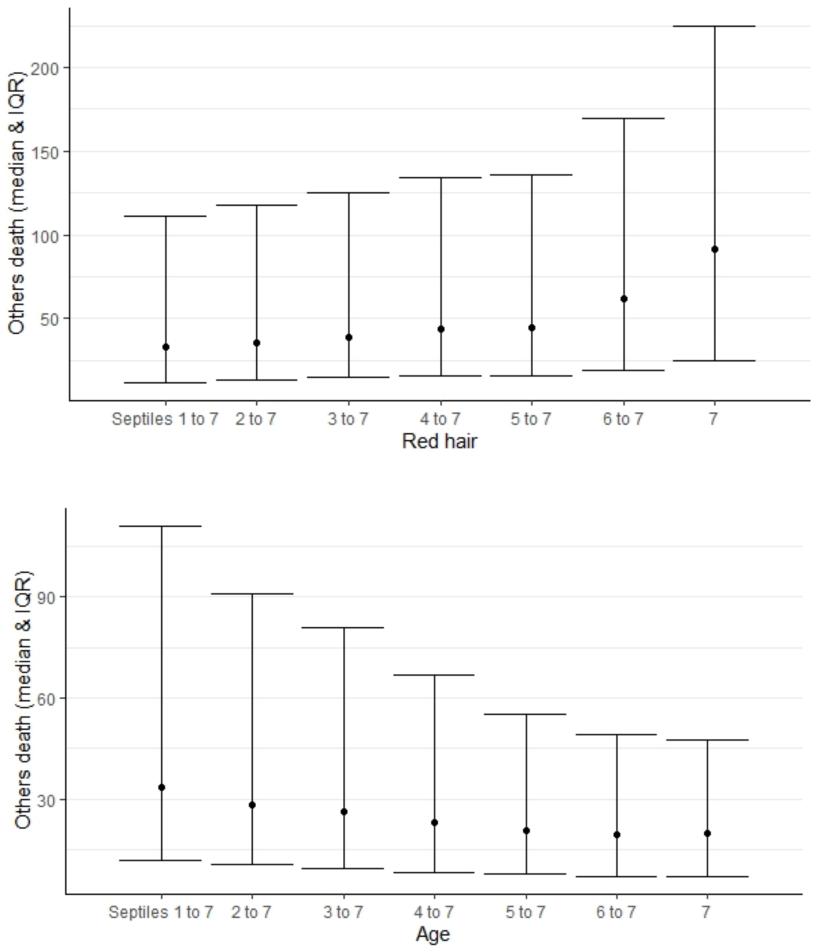

Figure 6.
The Figure plots estimates of *FATALITY* (others) for different ranges of red hair estimates. The top panel reports the median and the inter-quartile range by septiles of red hair estimate, from the whole sample on the left to the last septile only on the right. Bottom panel reports the median and the inter-quartile range by septiles of age, from the whole sample on the left to the last septile only on the right.

Higher septiles of red hair are associated with higher consensus *FATALITY* and substantially higher belief heterogeneity, as measured by the interquartile range. Higher septiles of Age are, in contrast, associated with lower consensus *FATALITY* and substantially lower belief heterogeneity.

Consistent with the model, consensus over/underestimation and disagreement are systematically predictable by the distribution of age and reliance on experience for judgments.

5. Memory, Beliefs and Behavior

Existing work on the pandemic has stressed the importance of political beliefs to shape 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?

In our survey we measured behaviour and attitudes such as: how often respondents leave home for reasons other than work or exercise, whether they have recently forfeited medical care in order to avoid leaving home, and whether they are in favour of 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 more difficult 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.[22]

Table 3 reports the estimates. Relative to the predictors of beliefs from Table 1, we add political affiliation which, while not selected as a predictor of beliefs, is a commonly cited predictor of attitudes towards the pandemic (Bursztyn et al 2020). We omit red hair from Table 3 in columns 2, 4, and 6. We report only "others death" and political views; in Appendix C we report all regression coefficients. 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

[22] 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.

policy decision such as whether to lift the lockdown. Political affiliation instead emerges as a key predictor for the latter, consistent with existing work.

Table 3.

The dependent variables are i) "going out", 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.?", which takes values 1 (never), 2 (once a week), 3 (twice a week), 4 (three or more times a week), ii) "med avoid", 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?", which takes values 1 (Yes, completely), 2 (Somewhat), 3 (Not at all), and iii) "Lift lockdown", the answer to the question "Would you resume your normal activities if lockdown or "stay-at-home" measures were lifted today?", which takes value from 1 (Definitely yes) to 5 (Definitely not). Death others is the estimate *FATALITY* (others), instrumented with the estimated number of red-haired Americans (F >> 10 in all cases). Republican degree is a variable which measures political orientation of the respondent which takes values from 1 (Strongly Democratic) to 7 (Strongly Republican). All variables are standardized and controls include variables which were selected by performing a dependent variable specific model selection algorithm.

	Dependent variable:					
	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^{***}	-0.228^{**}	-0.057^{**}	-0.278^{**}	-0.002	-0.119
	(0.023)	(0.112)	(0.023)	(0.114)	(0.019)	(0.098)
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)
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

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 and others, and that non-Covid health adversities had such a strong impact on Covid pessimism for others. We felt that this had to do with experiences, so we measured them in surveys 2 and 3, including non-health related ones. We discovered that beliefs about a domain such as Covid depend on a broad range of past experiences, including those from very different domains. These experiences, both relevant and irrelevant, affect beliefs because they provide material to simulate the future but also because they interfere with recall of other experiences that might even be better for simulation.

We formalize this process by building on established knowledge about simulation and interference from cognitive sciences. We obtain a range of predictions that help explain our initial puzzle but also many other findings, including the role of non-health past adversities as sources optimism, and the interference between domain relevant and irrelevant experiences. More broadly, the model offers a parsimonious account of the coexistence, frequently encountered in survey data, of consensus overestimation of unlikely events and large disagreement, where the latter also includes systematic underestimation of unlikely events in specific groups, such as the elderly. This role of experiences from other domains also accounts for the persistence of these belief differences despite the common experience in a given domain, as is the case with major events such as Covid.

Here we focused on Covid, but our approach may shed light on beliefs in other domains. Cryptocurrencies, global warming, the war in Ukraine are events new to many people, in which simulation from past experiences likely shapes beliefs. We suspect that even in familiar domains simulation and interference can affect beliefs. Our model delivers new hypotheses to test and new methods to test them. We did not design our survey having the simulation plus interference hypothesis in mind, but future surveys should try to measure the model's key ingredients: the database, meaning the frequency of a broad range of experiences, the similarity of these experiences to the event whose probability is assessed, and the respondents' tendency to overestimate unlikely events across domains. The measurement of similarity and frequency would allow a researcher to discover which experiences come to mind and their simulation potential. The tendency to overestimate unlikely events would capture reliance on experience. Such data would put structure on memory effects in generic domains, and possibly unveil new information, such as the tendency of people from different backgrounds or cultures to make different similarity judgments. In our model, this would translate into recalling different experiences when assessing the same event, creating belief differences.

These mechanisms can improve our understanding not only of beliefs but of many economic decisions. When deciding on a college major or whether to take a new job, people often rely on the experiences of socially close role models (Conlon and Patel 2022, Exley et al. 2022). Such people are similar to the decision maker and hence foster simulation much more than socially distant "artificial" role models or statistical information. In politics, a voter assessing a redistributive policy may selectively retrieve either the hard-working poor, and support it, or free riders, and oppose it. In fact, arguments that "talk past each other" by focusing on different subsets of the data already suggests a role of selective memory, and points to sets of experiences that can interfere with each other.

Critically, memory can explain why decisions often appear highly stable but sometimes display remarkable instability when individuals are purposely presented with different yet largely irrelevant frames. In particular, selective retrieval of past experiences would also help explain why well-crafted 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: it fosters retrieval of experiences that are good for simulating what the persuader is interested in, and interferes with conflicting thoughts.

More generally, memory is a key building block for all of our cognitive activities, so its effect can be far reaching. Even the distinction between beliefs and preferences may be more tenuous than conventionally thought. When we think about 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, which in turn are based on past experiences (Bordalo, Gennaioli, Shleifer 2020). Growing neuroscientific evidence indicates that memory is a critical part of this process (Shadlen and Shohamy 2016). We think that embracing 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. "The rich domain of uncertainty: Source functions and their experimental implementation." *American Economic Review* 101, no. 2 (2011): 695-723.

Alabrese, Eleonora and Fetzer, Thiemo. 2018. "Who is Not Voting for Brexit Anymore?" (2018). *CESifo Working Paper No. 7389.*

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.

Altig, David, Scott Baker, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Scarlet Chen, Steven Davis, Julia Leather, Brent Meyer, Emil Mihaylov, Paul Mizen, Nicholas Parker, Thomas Renault, Pawel Smietanka, and Greg Thwaites. 2020. "Economic Uncertainty Before and During the COVID-19 Pandemic." *Journal of Public Economics*, 191, 104274.

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.

Anderson, Richard. 1974. "Substance recall of sentences." *Quarterly Journal of Experimental Psychology, 26*(3): 530–541.

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

Becker, Sascha, Thiemo Fetzer, Dennis Novy. 2017. "Who voted for Brexit? A comprehensive district-level analysis." *Economic Policy*, 32(92): 601-650

Belot, Michele, Syngjoo Choi, Julian Jamison, Nicholas Papageorge, Egon Tripodi, and Eline van den Broek-Altenburg. 2020. "Six-Country Survey on Covid-19." SSRN wp 13230.

Benjamin, Daniel. 2019. "Errors in Probabilistic Reasoning and Judgment Biases." Handbook of Behavioral Economics, 2.

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, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2016. "Stereotypes." *Quarterly Journal of Economics* 131(4), 1753-1794.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2020. "Older People are Less Pessimistic about the Health Risks of Covid-19" *NBER* wp 27494

Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer. 2021. “Memory and Probability” *NBER wp 29273*

Brown, Norman, Lori Buchanan, and Roberto Cabeza. 2000. "Estimating the frequency of nonevents: The role of recollection failure in false recognition." *Psychonomic Bulletin & Review* 7(4): 684-691.

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2020. “Misinformation During a Pandemic”, NBER wp 27417.

Christoff, Kalina, Alan Gordon, Jonathan Smallwood, Rachelle Smith, and Jonathan Schooler. 2009. “Experience sampling during fMRI reveals default network and executive system contributions to mind wandering.” *Proceedings of the National Academy of Sciences*, 106 (21), 8719-8724.

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, Ben, Frederik Schwerter, and Florian Zimmermann. 2020. “Associative Memory and Belief Formation.” NBER wp 26664

Exley, Christine, Oliver Hauser, Molly Moore, and John-Henry Pezzuto. 2022. “Beliefs about Gender Differences in Social Preferences,” working paper.

Fan, Ying, Yesim Orhun, and Dana Turjeman. 2020. “Heterogeneous Actions, Beliefs, Constraints and Risk Tolerance During the COVID-19 Pandemic.” NBER wp 27211.

Gigerenzer, Gert, and Ulrich Hoffrage. 1995. “How to Improve Bayesian Reasoning Without Instruction: Frequency Formats.” *Psychological review*, *102*(4), 684-704.

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

Gilovich, Thomas. 1981. “Seeing the past in the present: The effect of associations to familiar events on judgments and decisions.” *Journal of Personality and Social Psychology*, *40*(5), 797-808.

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

Hassabis, Demis, and Maguire, Eleanor A. 2007. “Decostructing 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

Heimer, Rawley, Kristian Myrseth, and Raphael Schoenle. 2019. “YOLO: Mortality beliefs and household finance puzzles.” *The Journal of Finance*, *74*(6), 2957-2996.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An introduction to statistical learning*. New York: Springer

Jehiel, Philippe. 2005. “Analogy-based expectation equilibrium.” *Journal of Economic theory*, *123*(2), 81-104.

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

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

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

Kahneman, Daniel, and Tversky, Amos. 1981. “*The simulation heuristic*.” Stanford Univ CA Dept of Psychology.

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

Kolodner, Janet. 1983. “Maintaining organization in a dynamic long-term memory.” *Cognitive science*, *7*(4), 243-280.

Kuchler, Theresa, and Basit Zafar. 2019. “Personal Experiences and Expectations about Aggregate Outcomes.” The Journal of Finance, 74: 2491-2542

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.

Lichtenstein, Sarah, Paul Slovic, Baruch Fischhoff, Mark Layman, and Barbara Combs. 1978 "Judged frequency of lethal events." *Journal of experimental psychology: Human learning and memory,* 4(6), 551.

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.

Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under uncertainty: Heuristics and biases." *Science* 185(4157), 1124-1131.

Tversky, Amos, and Itamar Simonson. 1993. “Context-dependent preferences.” *Management Science*, *39*(10), 1179-1189.

Read, Daniel, George Loewenstein, and Matthew Rabin. 1999. "Choice bracketing." *Journal of Risk and Uncertainty*, 19(1), 171-197.

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

Shiffrin, Richard. 1970. “Forgetting: Trace erosion or retrieval failure?” *Science, 168*(3939), 1601-1603

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, D. J., & Gardner, M. K. 2015. “Semantic priming increases word frequency judgments: Evidence for the role of memory strength in frequency estimation.” *Acta Psychologica*, *160*, 152-160.

Appendix A. Proofs

Proof of Proposition 1 In the normative benchmark in which only Covid deaths can be used to simulate the target event and in which only Covid experiences are recalled to form judgments, the memory based estimate is frequentist, namely $\hat{\pi}_E = \frac{|D_C|}{|C|} = \pi$. If experiences other than Covid deaths can be used to simulate Covid death by factor $\tilde{\sigma}$ and if non Covid experiences can be recalled when thinking about Covid lethality according to similarity $\tilde{S}$, then using Equation (4) we have that:

$$\hat{\pi}_E = \frac{|D_C| + \tilde{\sigma}\left[|S_C| + \tilde{S}\left|\overline{C}\right|\right]}{|C| + \tilde{S}\left|\overline{C}\right|}.$$

It is immediate to find that this is larger than the frequentist estimate if and only if the true ifr is sufficiently low:

$$\frac{|D_C|}{|C|} = \pi < \pi^* \equiv \frac{\tilde{\sigma}|S_C|}{\tilde{S}\left|\overline{C}\right|} + \tilde{\sigma}.$$

Moreover, if non-lethal Covid experiences S_C are more recent, and thus more similar to Covid deaths, then the probability of simulation $\tilde{\sigma}$ is higher. This then implies that, all else equal, $\hat{\pi}_E$ is higher.

Proof of Proposition 2 Partitioning the experience database E into $E_i \subset E$ and $E_{-i} \equiv E \backslash E_i$ and using Equation (4) we obtain that memory based beliefs are equal to:

$$\hat{\pi}_E = \frac{\mathbb{E}_i(\sigma S)|E_i| + \mathbb{E}_{-i}(\sigma S)|E_{-i}|}{\mathbb{E}_i(S)|E_i| + \mathbb{E}_{-i}(S)|E_{-i}|}, \tag{A.1}$$

where $\mathbb{E}_x(.)$ denotes the average in subset E_x. It is immediate to find that:

$$\frac{\partial \hat{\pi}_E}{\partial |E_i|} = \frac{\mathbb{E}_i(\sigma S)\mathbb{E}_{-i}(S)|E_{-i}| - \mathbb{E}_i(S)\mathbb{E}_{-i}(\sigma S)|E_{-i}|}{[\mathbb{E}_i(S)|E_i| + \mathbb{E}_{-i}(S)|E_{-i}|]^2}. \tag{A.2}$$

Rearranging terms this yields:

$$sign\left\{\frac{\partial \hat{\pi}_E}{\partial |E_i|}\right\} = sign\left\{\frac{\mathbb{E}_i(\sigma S)}{\mathbb{E}_i(S)} - \frac{\mathbb{E}_{-i}(\sigma S)}{\mathbb{E}_{-i}(S)}\right\} = sign\left(\hat{\pi}_{E_i} - \hat{\pi}_E\right)$$

Higher frequency of experience E_i increases pessimism if the experience is easier to simulate Covid deaths than the rest. Next, define $S'(e) = s * S(e)$ for $e \in E_i$. Then,

$$\left.\frac{\partial \hat{\pi}_E}{\partial s}\right|_{s=1} = \frac{\mathbb{E}_i(\sigma S)\mathbb{E}_{-i}(S)|E_{-i}||E_i| - \mathbb{E}_i(S)\mathbb{E}_{-i}(\sigma S)|E_{-i}||E_i|}{[\mathbb{E}_i(S)|E_i| + \mathbb{E}_{-i}(S)|E_{-i}|]^2},$$

which implies:

$$sign\left\{\left.\frac{\partial \hat{\pi}_E}{\partial s}\right|_{s=1}\right\} = sign\left\{\frac{\mathbb{E}_i(\sigma S)}{\mathbb{E}_i(S)} - \frac{\mathbb{E}_{-i}(\sigma S)}{\mathbb{E}_{-i}(S)}\right\} = sign\left(\hat{\pi}_{E_i} - \hat{\pi}_E\right).$$

Proof of Proposition 3. To study the cross partial $\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|}$ with respect to a set of experiences $E_j \subset E$ that is non fully overlapping with E_i, $E_j \cap E_{-i} \neq \emptyset$, we can rewrite $(A.2)$ as:

$$\frac{\partial \hat{\pi}}{\partial |E_i|} = \frac{\mathbb{E}_i(\sigma S)\left[\mathbb{E}_{E_j \cap E_{-i}}(S)\left|E_j \cap E_{-i}\right| + \mathbb{E}_{-ij}(S)\left|E_{-ij}\right|\right] - \mathbb{E}_i(S)\left[\mathbb{E}_{E_j \cap E_{-i}}(\sigma S)\left|E_j \cap E_{-i}\right| + \mathbb{E}_{-ij}(\sigma S)\left|E_{-ij}\right|\right]}{\left[\mathbb{E}_i(S)|E_i| + \mathbb{E}_{E_j \cap E_{-i}}(S)\left|E_j \cap E_{-i}\right| + \mathbb{E}_{-ij}(S)\left|E_{-ij}\right|\right]^2}. \tag{A.3}$$

where $E_{-ij} = E \backslash E_i \cup E_j$. Now take the derivative of the above expression with respect to E_j by holding E_i constant, which amounts to taking the derivative with respect to $\left|E_j \cap E_{-i}\right|$. After some algebra, one finds that this is equal to:

$$\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial \left|E_j\right|} = K_{ij}\left[\left(\hat{\pi}_{E_i} - \hat{\pi}_{E_j \cap E_{-i}}\right) - 2\frac{\mathbb{E}_{-i}(S)|E_{-i}|}{\mathbb{E}_E(S)|E|}\left(\hat{\pi}_{E_i} - \hat{\pi}_{E_{-i}}\right)\right],$$

where $K_{ij} > 0$. Exploiting the fact that $\hat{\pi}_E = \left[1 - \frac{\mathbb{E}_{-i}(S)|E_{-i}|}{\mathbb{E}_E(S)|E|}\right]\hat{\pi}_{E_i} + \frac{\mathbb{E}_{-i}(S)|E_{-i}|}{\mathbb{E}_E(S)|E|}\hat{\pi}_{E_{-i}}$ we can write:

$$\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|} = K_{ij}\left[\left(\hat{\pi}_E - \hat{\pi}_{E_i}\right) + \left(\hat{\pi}_E - \hat{\pi}_{E_j \cap E_{-i}}\right)\right],$$

Which implies:

$$sign\left\{\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|}\right\} = sign\left\{\left(\hat{\pi}_E - \hat{\pi}_{E_i}\right) + \left(\hat{\pi}_E - \hat{\pi}_{E_j \cap E_{-i}}\right)\right\}$$

To see the empirical implications, note that we have the following measures of experiences: 1) Covid C, 2) non Covid health H, 3) Non health adversities NH, 4) Age A. There are three cases.

First, if both E_i and E_j boost pessimism, that is $\hat{\pi}_E < \hat{\pi}_{E_i}$ and $\hat{\pi}_E < \hat{\pi}_{E_j \cap E_{-i}}$, then we have $\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|} < 0$. This predicts a negative interaction between C and H. Second, if both E_i and E_j reduce pessimism, that is $\hat{\pi}_E > \hat{\pi}_{E_i}$ and $\hat{\pi}_E > \hat{\pi}_{E_j \cap E_{-i}}$, then we have $\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|} > 0$. This predicts a positive interaction between NH and A. Third, if E_i boosts while E_j reduces pessimism, that is $\hat{\pi}_E < \hat{\pi}_{E_i}$ and $\hat{\pi}_E > \hat{\pi}_{E_j \cap E_{-i}}$, the sign of $\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|}$ is generally ambiguous. Thus, we cannot sign the interaction between C and NH and in principle also the one of C and H with A.

Consider now the age interactions. For old people, $\overline{C}$ is large, so $\hat{\pi}_E \approx \hat{\pi}_{\overline{C}}$ and also $\hat{\pi}_{\overline{C} \cap E_{-i}} \approx \hat{\pi}_{\overline{C}}$. As a result,

$$sign\left\{\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |\overline{C}|}\right\} = sign\left\{\left(\hat{\pi}_E - \hat{\pi}_{E_i}\right) + \left(\hat{\pi}_E - \hat{\pi}_{\overline{C} \cap E_{-i}}\right)\right\} \approx$$
$$sign\{(\hat{\pi}_E - \hat{\pi}_{E_i}) + (\hat{\pi}_E - \hat{\pi}_{\overline{C}})\} \approx$$
$$sign\{(\hat{\pi}_E - \hat{\pi}_{E_i})\} = -\frac{\partial \hat{\pi}_E}{\partial |E_i|}$$

Comparing old people to the younger, the former should react less to any experience.

Proof of Corollary 1. It follows directly from inspection of Equation (6), together with the condition $K_{ij} > 0$ (Proposition 3) that $\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|} < 0$ if $\hat{\pi}_{E_i}, \hat{\pi}_{E_j} > \hat{\pi}_E$.

Proof of Corollary 2. Replacing E_j by the set of non-Covid experiences $\overline{C}$ in Equation (6), and taking the limit of large age so that $\hat{\pi}_{\overline{C}} \approx \hat{\pi}_E$, Equation (6) becomes

$$\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |\overline{C}|} = K_{ij}\left(\hat{\pi}_E - \hat{\pi}_{E_i}\right)$$

with $K_{ij} > 0$, so that the sensitivity of beliefs $\hat{\pi}$ to any set of experiences E_i decreases in $\overline{C}$: if $\hat{\pi}_{E_i} > \hat{\pi}_E$ (E_i is a source of pessimism), then $\hat{\pi}$ becomes less pessimistic as $\overline{C}$ increases, and conversely if $\hat{\pi}_{E_i} < \hat{\pi}_E$.

Proof of Corollary 3. The average belief of people with tendency θ to simulate from memory is given by Equation (5), $\hat{\pi} = (1-\theta)\pi + \theta\hat{\pi}_E$. Since the first term does not depend on experiences X, Equaion (9) follows from inspection, as applied either to the average beliefs of this group, or to the expected belief of a subject characterized by θ.

Appendix B. The Survey

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.[23] The median time taken to complete our survey is approximately 10.5 minutes.

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

[23] 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.

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.[24]

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 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.

[24] 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.

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 hopsitalization 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

In line with Proposition 1, we can compute the experience based estimate $\hat{\pi}_H$ ($\hat{\pi}_{NH}$) obtained when only "Health Adversity" ("Non Health Adversities") are used for simulation. To do so, we assume that i) similarity linearly declines in the rank, that is $S(e) = 1 - \frac{r(e)}{8}$, where $r(e)$ is the average rank of experience e, ii) simulation is equal to similarity, formally $\sigma(e) = S(e)$, and iii) we compute for each respondent who has had at least one health adversity the memory based estimate $\hat{\pi}_H$ based on those, and for each respondent who has had at least one non-health adversity the memory based estimate $\hat{\pi}_{NH}$ based on those. The estimates at point iii) are computed using the assumptions i) and ii) about similarity and simulation from points, and using the average rank of Table B1 as an input. We that that the average value of $\hat{\pi}_H$ in the population is 0.56 and the average value of $\hat{\pi}_{NH}$ is 0.44. Consistent with Proposition 1, the average respondent in the sample has $\hat{\pi}_H > \hat{\pi}_{NH}$.

Appendix C. Summary Statistics and Robustness

In this appendix we present:

1. Summary statistics, correlations, and description of the variables included in our analysis;
2. The full version of tables 1, 2, and 3. 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.
- State Level (commonly referred as Level, also) 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).

- Days since Peak (referred to as Peak, also) 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 output of table 1, in the first two columns. Hence, coefficients for Income, Black, Asian, and Rural are shown. In columns 3 and 4, it presents results for infection and hospitalization beliefs. 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.

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

Income	-0.036**	-0.044***	-0.083***	-0.043**
	(0.016)	(0.016)	(0.013)	(0.019)
Black	0.111**	0.164***		0.084**
	(0.053)	(0.048)		(0.042)
Asian		0.205***		
		(0.060)		
Rural	0.123***	0.068**		0.064*
	(0.033)	(0.030)		(0.035)
Constant	-0.084***	-0.103***	-0.027*	-0.086***
	(0.022)	(0.022)	(0.014)	(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 2, in the first two columns. Column 3, in Table A4, 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.

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.”

	Dependent variable:		
	Others death		
	(1)	(2)	(3)
Had Covid		0.441***	0.446***
		(0.167)	(0.167)
Health adversities		0.047**	0.046**
		(0.019)	(0.019)
Non health adv.		-0.039***	
		(0.015)	

Non health adv. (small)			-0.031*
			(0.016)
Subj. adversity		0.043**	0.041**
		(0.019)	(0.019)
No. health cond.	0.029**	0.012	0.010
	(0.013)	(0.017)	(0.016)
Hosp (self.)	0.218***	0.157**	0.160**
	(0.078)	(0.073)	(0.073)
Hosp (fam.)	0.061	0.058	0.050
	(0.045)	(0.044)	(0.044)
State Level	0.061***	0.059***	0.061***
	(0.023)	(0.023)	(0.023)
Days since Peak	-0.098***	-0.097***	-0.097***
	(0.024)	(0.023)	(0.023)
Red hair	0.169***	0.165***	0.166***
	(0.033)	(0.033)	(0.033)
Age	-0.227***	-0.212***	-0.216***
	(0.017)	(0.021)	(0.020)
Income	-0.035	-0.043*	-0.042*
	(0.024)	(0.023)	(0.022)
Black	0.143***	0.133**	0.136**
	(0.053)	(0.054)	(0.054)
Asian	0.239***	0.249***	0.252***
	(0.089)	(0.092)	(0.091)
Rural	0.108***	0.113**	0.116***
	(0.042)	(0.044)	(0.044)
Constant	-0.114***	-0.128***	-0.129***
	(0.026)	(0.030)	(0.029)
Observations	2,972	2,953	2,953
Adjusted R^2	0.119	0.133	0.132
Note:			*p<0.1; **p<0.05; ***p<0.01
			Clustered standard errors at state level

Table C5 shows the full output of table 3. As we explained in the main text, controls were chosen by performing model selection for each specific dependent variable.

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 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.

	Dependent variable:					
	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***	-0.228**	-0.057**	-0.278**	-0.002	-0.119
	(0.023)	(0.112)	(0.023)	(0.114)	(0.019)	(0.098)
Max weekly growth death	-0.057***	-0.055***				
	(0.014)	(0.014)				
Days since wk death peak	0.044*	0.036				
	(0.023)	(0.023)				
Current level death			-0.019	-0.028		
			(0.023)	(0.020)		
Age	0.065***	0.023	0.227***	0.169***		
	(0.018)	(0.031)	(0.016)	(0.031)		
Age squared			0.065***	0.076***		
			(0.016)	(0.015)		
Female	-0.051***	-0.049**			0.113***	0.115***
	(0.019)	(0.020)			(0.020)	(0.021)
Black					0.026	0.034*
					(0.018)	(0.019)
Asian	-0.071***	-0.062***			0.056***	0.066***
	(0.019)	(0.017)			(0.014)	(0.018)
Rural			-0.102***	-0.089***		
			(0.020)	(0.019)		
Education			-0.092***	-0.093***		
			(0.017)	(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)

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:

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

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 $\sum_{j=1}^{p}I\left(\beta_j\neq 0\right)$.

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: Mallow's C_p, Bayesian information criterion (BIC), and adjusted R^2. Mallow'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 Mallow's C_p or the BIC. The adjusted R^2 is defined as 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).

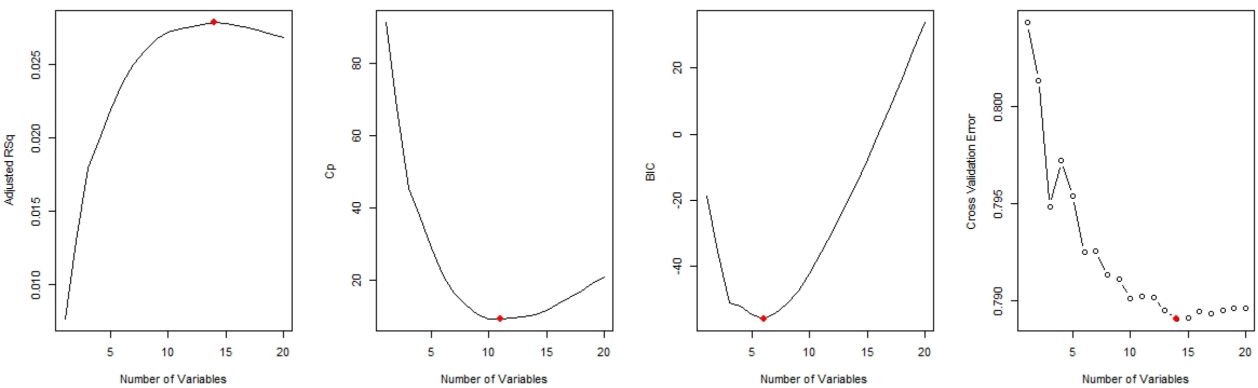

Figure A1

Adjusted R^2, Mallow'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.[25] 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 Mallow's C_p suggests to

[25] 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.

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 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.[26]

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

[26] 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.

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;[27]

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 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

[27] 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.