Imagining the Future:

Memory, Simulation and Beliefs about Covid

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

April 1, 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 a number of 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 or simulate the novel risk. In the model, greater exposure to related experiences enhances risk perception by making the risk easier to imagine, but dampens risk perception by interfering with recall of other experiences that may feed imagination. The model accounts for our initial findings, but also connects average overestimation of unlikely risks with strong disagreement: people exposed to many interfering experiences underestimate risk and are less sensitive to related experiences. We find empirical support for these and other predictions using our survey data on respondents' Covid as well as non-Covid past experiences.

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 priors and likelihoods come from? An alternative approach is to use personal experiences, as opposed to statistical data (Malmendier and Nagel 2011). But for novel risks, there may be few, if any, closely related personal experiences to draw on to form beliefs. How do people form beliefs in such cases?

We propose that, when forming their beliefs, people recall past experiences – including those from different domains – and use them to imagine the novel risk. Memory-based imagination, which psychologists call "simulation", is known to be central for thinking about the future (Dougherty et al. 1997; Brown et al. 2000). It is related to reasoning by analogy, and entails retrieving and recombining experiences stored in memory (Carroll 1978, Schacter 2007, 2018). We offer a model in which people simulate the future based on recalled experiences but, critically, not all experiences come to mind because memory is spontaneous and selective on the basis of similarity, frequency and interference, in the spirit of Kahana (2012) and Bordalo et al. (2021).

Consider an investor assessing Amazon in its early years. Thinking about it as a "bookstore" would help simulate failure, given the decline of that sector. Thinking about it as an "internet firm" would instead help imagine success, based on innovators like Microsoft. Although most people have access to both angles, the two are not equally available in memory. Investors with more life experiences may arguably focus on "bookstore" more so than new investors, because they are comparatively more exposed to traditional sellers. Experienced investors may then be too pessimistic and inexperienced ones too optimistic, even when they have the same data about Amazon. Through the mechanisms of selective retrieval and simulation, experiences from different domains can matter more than domain specific information and create strong disagreement.

Our model is motivated by a large dataset of beliefs about Covid-related risks that we collected from a large sample of Americans 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 two surprising facts (Bordalo et al. 2020), which we confirmed in subsequent waves. First, 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, while the elderly underestimate it. Second, respondents are more pessimistic if they experienced non-Covid health adversities such as own or a family member's past hospitalization. The effect is quantitatively large, similar to that of moving from the bottom to the top tercile of cumulative Covid deaths in the respondent's state. This is puzzling: one might have expected exposure to other health risks (and surviving them) to lower one's estimated Covid lethality.

To explain these facts, in Section 3 we present a model in which people form beliefs by automatically retrieving from memory Covid and non-Covid past experiences and using them to imagine or "simulate" Covid deaths. Simulation is easier with an experience more similar to a Covid death. We show that non-Covid experiences exert an ambiguous effect on beliefs about Covid. On the one hand, these experiences may help simulate Covid deaths, for instance because they are also severe diseases. On the other hand, recall of these experiences may interfere with recall of more relevant experiences, reducing the ability to imagine deaths from Covid.

Going back to our motivating facts, this trade-off implies that it is a priori ambiguous whether exposure to non-Covid health adversities raises or reduces pessimism. It however predicts that exposure to much less similar adversities such as the non-health ones (e.g. working in a dangerous job or experiencing poverty), should reduce Covid pessimism. The latter experiences, in fact, should interfere more with imagining Covid risks. Section 4 tests this prediction using data on health and non-health adversities from surveys 2 and 3. Consistent with our model, people who experienced more health adversities are more pessimistic about Covid (simulation), while those exposed to more non-health adversities are more optimistic (interference). Non-health adversities

reduce Covid pessimism in absolute terms, and strongly so: going from minimum to maximum non-health adversities undoes the pessimism from moving from zero to maximal local Covid deaths.

Interference is also consistent with the striking age gradient: the database of the elderly is flooded by interfering non-Covid experiences, which reduce simulation of Covid deaths and hence pessimism. This mechanism has an additional implication: due to stronger interference, the beliefs of the elderly should be less sensitive than those of the young to any specific past experience. The data is consistent with this prediction, as well as with a broader implication of interference: exposure to bad health experiences renders respondents less sensitive to local pandemic conditions.

As a final step, we connect memory effects to a respondent's general tendency to overestimate unlikely events. We asked survey respondents to estimate the share of Americans who have red hair. Respondents who overestimate this share are also more pessimistic about Covid. In our model this correlation obtains if people who overestimate red haired Americans rely more on simulation than others, so they overestimate unlikely events in both risky and non-risky domains. The Covid beliefs of respondents who exhibit stronger overestimation of red-haired Americans should then be more sensitive to all experiences, both those that increase pessimism and those that reduce it. This is another new prediction of our model, and it finds empirical support in our data.

Our approach unifies an average tendency to overestimate unlikely risks with strong disagreement among different people. Existing models of overestimation of unlikely events (Kahneman and Tversky 1979) either neglect the opposite phenomenon of underestimation, or attribute it to noise or uncertainty (e.g. Enke and Graeber 2022, Kaw et al. 2020). These models cannot explain why a certain group of people, such as the elderly, should predictably underestimate an unlikely risk. Our model delivers both average overestimation of rare events due to selective retrieval and simulation of rare Covid experiences, and systematic underestimation by subjects who are exposed to many interfering experiences, such as the elderly.

Disagreement could also be explained by experience effects (Malmendier and Nagel 2011, 2016, Malmendier 2021). This approach however has a hard time explaining average

overestimation of an unlikely event such as Covid's lethality early in the pandemic, when the experiences of it are few, or underestimation by the elderly, who saw more fatalities than the young. Moreover, as the pandemic developed, the reactions to the same Covid shock diverged, with some people greatly overestimating and others underestimating its risks. We explain this phenomenon via selective retrieval of different non-Covid experiences, such as past health and non-health adversities. We show that the effect on fostering or inhibiting the imagination of Covid risks can be large, to the point of swamping standard, domain specific experience effects. Selective retrieval and simulation throw new light on the effect of past experiences on beliefs, with new predictions.

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.

Our model builds on the "similarity plus interference" setting of Bordalo et al. (2021). We continue the program of unifying different belief biases based on selective memory. Theoretically, our innovation is to allow for 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 world event, rather than on the 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 stylized 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, stratifying our sample 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 stratification, question order, payments, and quality controls.[2]

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

---

[2] Some of the questions asked were not described here, as they are not directly related to beliefs about Covid risks, either through the lens of our model or of our predictor model (see Sections 3 and 4).

[3] Specifically, 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.

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,[4] but it turns out to more generally proxy 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, though this is not the focus of our analysis.

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.
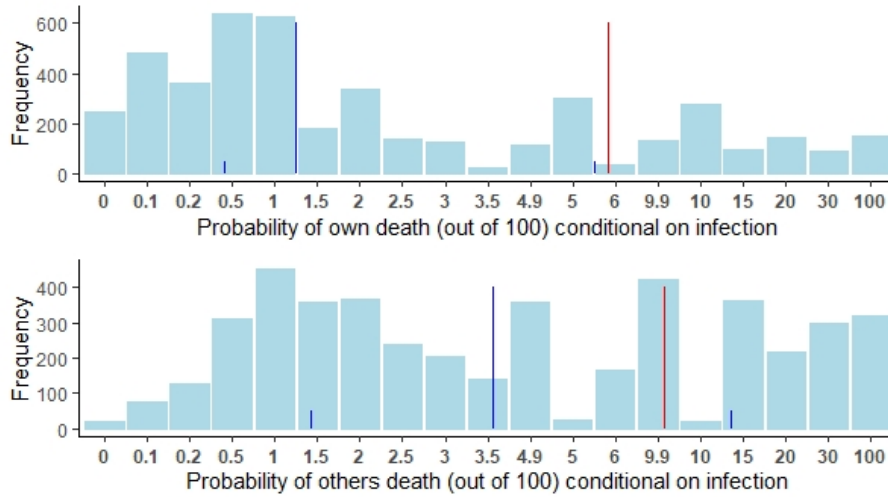
Figure 1

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.

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 the controls that exhibit the highest explanatory power from all measures collected in our survey and from a battery of proxies for state-level covid severity.[6] The selection criterion picks three demographics other than 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.[7] 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 (see footnote 3). 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. Level is the cumulative number of deaths for Covid in the respondent's 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. The controls are the remaining selected variables (Income, Black and Rural for Column 1, Income, Black, Rural and Asian for Column 2).

Dependent variable:

---

[5] Accessible from the New York Times counts, https://www.nytimes.com/interactive/2021/us/covid-cases.html.
[6] Details of variable selection are in Appendix D. After presenting the model, we introduce theoretically justified regressors. When we test the model, controlling for "theoretically non-justified" variables selected based on their strong explanatory power helps assess the robustness of theoretical predictions. The selection method we use takes all possible regressions including all combinations of control variables and it outputs the one which minimizes an information criterion. We employ different information criteria to identify the subset of our predictors, to obtain robust results.
[7] 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.

| | Risk of Own death | Risk of Others death |
	(1)	(2)
Age	-0.131***	-0.236***
	(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 current personal non-Covid health conditions and recent personal non-Covid hospitalization 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). In fact, even a non-personal health adversity such as past hospitalization of a family member sharply increases Covid pessimism for *others*.

Fourth, Covid experiences matter. "Level" measures the cumulative number of deaths in a state at maximal weekly case growth. It proxies for viral severity at peak transmission. The effect of the peak deaths fades over time: if the peak occurred longer ago (so "Days" are higher), pessimism is lower. A Bayesian would expect respondents learning from local dynamics to 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, infections or population do not reliably affect beliefs. As a result, 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 – as we already mentioned – the effect is weak and disappears when controls are added, so political affiliation never gets selected as a predictor of beliefs (as we show later, political affiliation is instead an important determinant of behaviour). 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 experience effects (Malmendier and Nagel 2011), 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%. 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.[8]
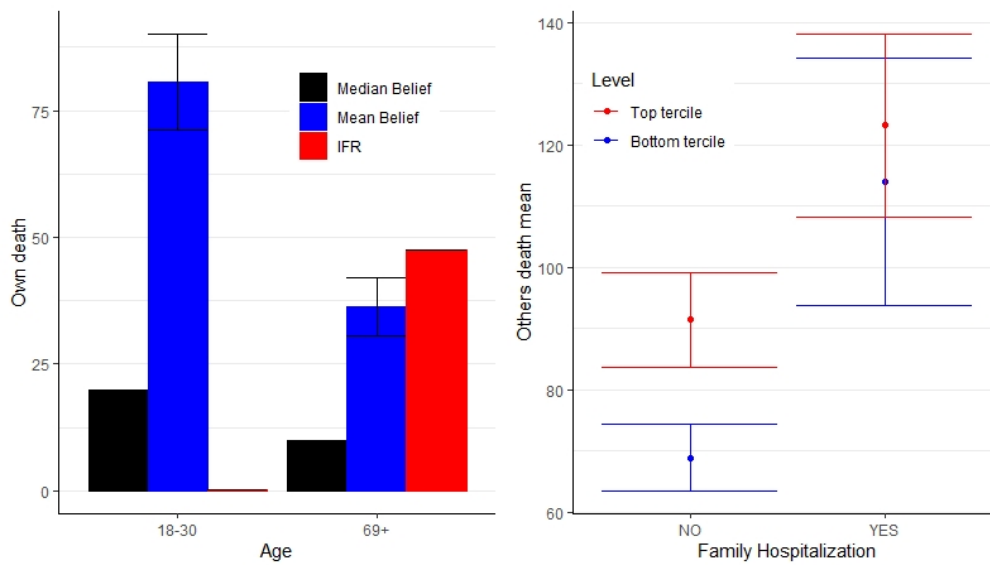
Figure 2

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

---

[8] 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.

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.*

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 the stock market, affect beliefs about that same setting, but not about a similar and even correlated setting such as the bond market (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, given 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*, experiences that are similar enough to Covid deaths or that frequently occur in the respondent's database compete for retrieval, consistent with the well-established roles of similarity, frequency and interference in memory research (Kahana 2012). The experiences that come to mind are then used to imagine Covid death. Recalling experiences that help imagination boosts Covid pessimism, experiences that do not help boost optimism.

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 in 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*.[9] Parameter θ thus captures the DM's reliance on simulation. 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

---

[9] 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.

maximal, equal to 1, when $u = v$. A Covid death is very similar to one from SARS, less similar to 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 exerienced in the remote past are less similar to current ones because they occurred in a different context (Kahana 2012). Our analysis relies on general intuitions about similarity, which can however be formalized using a features-based similarity function.

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 experiences of Covid deaths may be interfered with by instances of other respiratory diseases that come to mind because the latter have high similarity $S(u)$. Crucially, events that frequently occur in the database can be recalled and interfere with Covid experiences even if they are fairly dissimilar from Covid deaths, because their summed similarity in the denominator of (2) is high. Common diseases such as heart attacks or lethal events such as 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.[10] 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, in the sense that 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, and it is easier to imagine a Covid death based on deaths from SARS than based on those from a heart attack because SARS is more similar to the target. However, even experiences that are not very similar to the target may help simulate it: the experience of someone dying in a hospital may help imagine a Covid death.

---

[10] For example, recall from a target list of words suffers intrusions from other lists studied at the same time, particularly for words that are semantically related to the target list, resulting in lower likelihood of retrieval and longer response times (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:[11]

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

Equation (3) describes memory-based beliefs. To see its implications, partition the database E into three sets: i) Covid deaths D, ii) Covid survivals $\overline{D}$, and iii) non-Covid $\overline{C}$. The set $C = D \cup \overline{D}$ 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 ($\sigma(e) = 1$ for $e \in D$), while simulation fails based on other experiences ($\sigma(e) = 0$ for $e \in \overline{D}$). Suppose in addition that similarity is also "narrow": the similarity of Covid experiences to "Covid deaths" is maximal ($S(e) = 1$ for $e \in C$), that of non-Covid experiences is nil ($S(e) = 0$ for $e \in \overline{C}$). Then, the memory-based estimate is frequentist:

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

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

This, however, is a knife edge case. First, similarity is not narrow. Covid experiences share at least some features with non-Covid ones, such as other diseases or adversities. This tends to raise the denominator of Equation (4), promoting underestimation. To see this starkly, suppose that the similarity function is constant. Then, in Equation (3) the average experienced-based estimate is $\hat{\pi}_E = \mathbb{E}(\sigma|E) = \Pr(D|E)$, which is the relative frequency of Covid death experiences in E.

---

[11] 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.

Someone who has few Covid experiences underestimates *FATALITY*. Intuitively, recall of Covid deaths is interfered with by that of frequent experiences that are bad at simulating Covid deaths.

Second, simulation is also not narrow, and is in fact much broader than standard experience effects. Seeing images of Covid patients laying in ICU beds encourages simulation even absent any Covid deaths. Experiences with disease can help simulate Covid deaths, promoting overestimation. To see this starkly, if simulation is a constant σ the experience based estimate in (4) is $\hat{\pi}_E = \sigma$. The DM then overestimates Covid *FATALITY* as long as $\sigma > \pi$, that is, if the objective risk π is low. Simulation promotes overestimation because even experiences that should only be counted in the denominator of (4), mild Covid cases, are used to simulate Covid death, raising the numerator.

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 and both the simulation and similarity functions are narrow. The average belief is then frequentist and corresponds to the statistical benchmark $\hat{\pi} = \pi$. Keeping the Covid database unbiased, suppose that both simulation and similarity become broader: Covid deaths can be simulated using other experiences, $\sigma(e) = \tilde{\sigma} > 0$ for all $e \in \overline{D}$, 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, $|D|/|C| = \pi$. 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 $\hat{\pi}_E > \pi$ if and only if $\pi < \pi^*$. In this case, FATALITY increases in the DM's reliance on simulation, $\frac{\partial \hat{\pi}}{\partial \theta} > 0$.*

Irrelevant experiences exert two conflicting effects. On the one hand, they foster simulation of Covid deaths, which boosts $\hat{\pi}_E$. On the other hand, they interfere with recall of Covid death experiences, which reduces $\hat{\pi}_E$. If Covid deaths are rare, even in an unbiased database there are few Covid death experiences that can be interfered with. As a result, simulation dominates. Many Covid deaths can be simulated based on the numerous non-lethal Covid experiences or on other health adversities, boosting pessimism. People put positive probability on the occurrence of events they had never seen, provided they are somewhat similar to their own 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 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 message of Section 2 is disagreement. Our model features two main sources of disagreement. The first is "noise", which arises because recall from the database is random. Random recall cannot however explain systematic belief differences among groups, say young and old, so we do not explore it here. Second, disagreement can arise because different people have different experience databases E, leading to different simulation and interference. We next show that individual differences in the experience database E can account for the role of different experiences in Table 1 and yield new predictions that we test using data from survey 2 and 3. This includes relevant experiences such as Covid deaths, captured by "Level", as well as irrelevant experiences such as non-Covid health adversities and even non-health adversities.

4. Disagreement: Theory and Empirical Tests

4.1 The effects of relevant and irrelevant experiences

To analyse the role of different databases E, take a subset E_i of experiences sharing certain 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. We obtain the following result.

Proposition 2 *Experience E_i is a source of Covid pessimism, in the sense that increasing its numerosity $|E_i|$ increases FATALITY, $\frac{\partial \hat{\pi}}{\partial |E_i|} > 0$, if and only if $\hat{\pi}_{E_i} > \hat{\pi}_E$. In particular, adding a single experience e to the database E increases FATALITY if and only if:*

$$\pi_e = \sigma(e) > \hat{\pi}_E. \tag{6}$$

Increasing the frequency of experiences E_i boosts pessimism if and only if *FATALITY* estimated using *only* these experiences is higher than *FATALITY* estimated using the entire database E, namely when $\hat{\pi}_{E_i} > \hat{\pi}_E$. This occurs when the specific subset of experiences E_i is on average better for simulation and creates little interference than the full set of experiences E the DM had.

In Equation (6), a specific experience acts as a source of pessimism if its similarity to Covid deaths, and hence its simulation potential, is high compared to other experiences in the database (captured by $\hat{\pi}_E$). One immediate implication is that exposure to an experience that coincides with the target event should boost pessimism. This accounts for the standard experience effect in Table 1, whereby the cumulative level of Covid deaths at peak, "Level", acts as a source of pessimism. "Level" in fact proxies for Covid death experiences, $E_i = D$. Such experiences help simulate Covid deaths better than other experiences in E, formally $\hat{\pi}_D > \hat{\pi}_E$. Thus, exposure to more lethal Covid experiences translates to greater pessimism, consistent with Table 1.[12]

---

[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.

Critically, Proposition 2 implies that even irrelevant experiences shape beliefs. Equation (6) shows that for such experiences memory effects are a priori ambiguous: whether a certain experience is a source of Covid pessimism or optimism depends on its similarity to a Covid death compared to the average experience in the database E. Consider the role of non-Covid health adversities, such as past hospitalizations of self and others, or of current health problems. Equation (6) implies that such experiences encourage Covid pessimism if they are similar enough to Covid deaths, i.e., if $\sigma(e)$ is high enough that they boost simulation compared to interference they create. The fact that in Table 1 these experiences increase Covid pessimism suggests that on average across our subjects their simulation potential $\sigma(e)$ is higher than the interference they create for $\hat{\pi}_E$.

A testable prediction of our model is that exposure to experiences that are sharply less similar to Covid deaths than these non-Covid health adversities should promote Covid optimism, because Equation (6) is violated for them. This is an instance of a more general principle that follows from Equation (6): if an experience e acts as a source of pessimism, then experiences that are more similar to Covid deaths than e should also act as sources of pessimism, while experiences that are less similar to Covid deaths than e should be sources of lesser pessimism or even optimism.

We test this prediction, using data from Surveys 2 and 3. First, in these surveys we measure direct Covid experiences. In particular, we ask whether the respondent had Covid or not. This experience is similar to the target event because it entails the same disease, but different because it might have been mild and certainly non-lethal. Still, it could be used for simulating deaths.

In surveys 2 and 3 we also asked whether respondent experienced a range of adversities in life. We construct an index of "Health Adversities", which measures if the respondent suffered from a serious illness or injury, and an index of "Non-Health Adversities" that measures if 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 serious injury, illness or untimely death of a loved one. These are all severe adversities, and our model makes a clear prediction for them: "Health Adversities", being similar to a disease like Covid (and to the health adversities and conditions

22

proxies used in Table 1) should promote pessimism more than than "Non-Health Adversities". The latter proxy for non-health risks that are much less similar to Covid death. In particular, Non-Health Adversities are sufficiently different from bad medical conditions and hence from Covid that they may even violate Equation (6), creating strong interference and leading to Covid optimism.[13]

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).

	Dependent variable:	
	Others death	
	(1)	(2)
Had Covid		0.441***
		(0.167)
Health adversities		0.047**
		(0.019)
Non health adversities		-0.039***
		(0.015)
Subj. adversity		0.043**
		(0.019)
No. health cond.	0.029**	0.012
	(0.013)	(0.017)
Hosp (self.)	0.218***	0.157**

---

[13] We include "serious injury, illness or untimely death of a loved one" among non-health adversities because from the point of view of the respondent adversity comes from the loss of the loved person. Furthermore, the cause of loss may include but is not restricted to health problems (e.g. it could be due to violence or accidents). The results of Table 2 go through if we omit this variable, see Appendix C. Another possible measure of adversity is recent unemployment; however, this is a less extreme experience than experiencing poverty, and it may be correlated with working in an industry affected by Covid (e.g. hospitality) which may introduce confounds in beliefs about Covid risks.

	(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, non-Covid health adversities boost pessimism. The negative past health experiences measured in surveys 2 and 3 exert the same directional effect as the negative recent health experiences in Table 1, consistent with our model.

Interestingly, having had Covid also acts as a source of pessimism.[14] Our model explains this result with simulation: surviving Covid is sufficiently similar to a Covid death, due to the fact that it entails the same disease, that it helps simulate lethal events, acting as a source of pessimism. A priori, this is quite surprising, for one may have expected Covid survivors to be more optimistic than people who did not get Covid. In the early stages of the pandemic, however, when Covid infections were few and there was significant uncertainty, getting Covid may invite simulation of *FATALITY*, boosting pessimism. Going back to Table 1, this result can help explain why the number of infections in a state or its population do not reduce pessimism, even after controlling for

---

[14] We also measure indirect Covid experiences by asking whether the respondent knows someone who had Covid, someone who was hospitalized for Covid, or someone who died from Covid. 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.

"Level" deaths. People seeing many Covid infections have many experiences that help them simulate *FATALITY*, so they insufficiently discount for the fact that most infections are mild or that their numerosity is due to a higher population. Availability of such simulation material promotes pessimism, causing a departure from the frequentist benchmark in Equation (4).

Crucially, Table 2 shows that having had experiences of Non-Health Adversities goes in the opposite direction, acting as a source of Covid optimism. We explain this finding with 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. In line with Equation (6), this reduces the ability to simulate Covid deaths, fostering optimism.

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.

This result yields two important messages. First, it shows that simulation and interference can produce systematic underestimation and overestimation of unlikely events, shedding light on the disagreement documented in Section 2. Second, it shows that irrelevant experiences can offset standard domain specific experience effects, creating optimism even in highly affected locations.

4.2 The Age Gradient

A key finding in Table 1 is that the elderly are much less pessimistic about Covid risks than the young, to the point that the belief differences are dramatically counterfactual. To see how our model accounts for the age gradient, we can rewrite Equation (3) as follows:

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

where $C \subset E$ is the subset of Covid experiences and $\overline{C}$ is the subset of non-Covid ones. Because Covid is a new risk, older respondents have a larger non-Covid database $\overline{C}$ for given Covid experiences C. Two consequences follow.

First, because Covid experiences are more effective at simulating Covid deaths than non-Covid experiences, formally $\hat{\pi}_{\overline{C}} < \hat{\pi}_E$, Proposition 2 directly implies that older respondent should be less pessimistic than younger respondents. This simple fact already yields the striking age gradient: the database of older people is flooded with many non-Covid experiences. The elderly suffer from stronger interference when trying to simulate Covid deaths.

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).[15] 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 consequence of Equation (7) is that the elderly should be less sensitive to Covid news. As $|\overline{C}|$ gets larger, the marginal effect of Covid experiences becomes small, formally $\hat{\pi}_{\overline{C}} \approx \hat{\pi}_E$. This illustrates a more general principle: In our model, whether an experience is a source

---

[15] 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.

of optimism and pessimism, and the extent to which it moves beliefs in either direction, is not an absolute property of the experience itself. It depends on the database to which that experience is added. This principle yields two new testable implications.

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

$$\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|} = K_{ij} \left[\left(\hat{\pi}_E - \hat{\pi}_{E_i} \right) + \left(\hat{\pi}_E - \hat{\pi}_{E_j} \right) \right], \tag{8}$$

where $K_{ij} > 0$. This yields the following predictions:

i)   *The beliefs of the elderly should be less sensitive to each experience E_i. Formally for $E_j = |\overline{C}|$, when $|\overline{C}|$ is sufficiently large that $\hat{\pi}_{\overline{C}} \approx \hat{\pi}_E$, then $\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |\overline{C}|} = -\frac{\partial \hat{\pi}}{\partial |E_i|}$.*

ii)   *if E_i and E_j are sources of pessimism, $\hat{\pi}_{E_i}, \hat{\pi}_{E_j} > \hat{\pi}_E$, then higher $|E_j|$ dampens the marginal effect of $|E_i|$ on beliefs, $\frac{\partial^2 \hat{\pi}}{\partial |E_i| \partial |E_j|} < 0$.*

Prediction i) says that the beliefs of the elderly should be less sensitive to any specific experience, be it Covid or non-Covid, and be it a source of pessimism or of optimism. The elderly face strong interference, which reduces the marginal impact of any specific experience.

Prediction ii) says that different sources of pessimism should interfere with each other, mutually dampening their marginal effect on beliefs (the same is true for sources optimism). 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 can test these predictions. First, we test the prediction that the beliefs of the elderly are less sensitive to any specific set of experiences (point i)). Second, we test whether increasing exposure to a source of pessimism interferes with other sources of pessimism (point ii).

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 diminishes over time, which is another prediction of point ii) above, by adding age squared to the regression of Table 2 (it should have a positive coefficient).
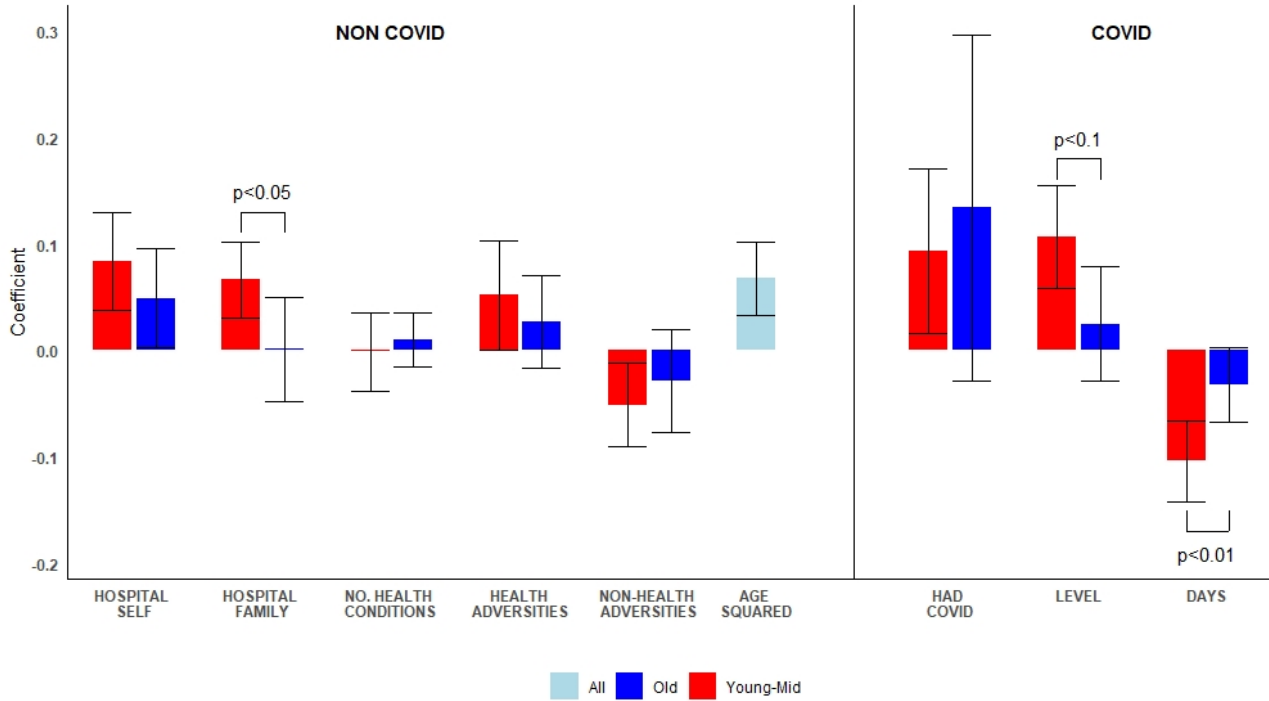
Figure 3.
The figure reports the coefficients obtained by estimating the equations for beliefs of others death in Tables 1 and 2 in the first two terciles of age (18-61) and in the top tercile (62+). Coefficients for variables available in all waves (hospital self, hospital family, no. health conditions, 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 our predictions, 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. 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 are not just insensitive to sources of pessimism, and hence more optimistic. They are less sensitive across the board, which in our

28

model comes from their difficulty of recalling any specific source, due to interference from many other experiences. An F-test for the null hypothesis that the coefficients are identical across the age groups is rejected.[16] Also consistent with the model, the coefficient of age squared is positive.

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. Contrary to our predictions, the elderly who had Covid are more pessimistic than the young, but the effect is noisy because Covid cases in our sample are rare, especially for the elderly. Overall, the interference associated with age seems to also modulate standard experience effects, causing a dampened reaction of beliefs to Covid related events.

In a Bayesian world, older people might react less to news because they have a longer history of data, so they 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.[17] A more important 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.

We next test, following Proposition 3 point ii), whether different past experiences that act as sources of pessimism interfere with each other. Such sources include two Covid experiences ("Level" and "had Covid"), three past non-Covid health adversities ("own hospital", "serious injury" and "serious illness"), and non-Covid health adversity of the respondent's close contacts ("family hospital"). We focus on the cross interference between the local severity of Covid, "Level", and the other past health adversities. This allows to evaluate the importance of memory interference relative to standard domain specific experience effects.

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

---

[16] 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.
[17] For instance, people in the age group 72+ underestimate own lethality by 2.5%, those in age group 65-71 by 1.7%.

with a degree of severity of the 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.
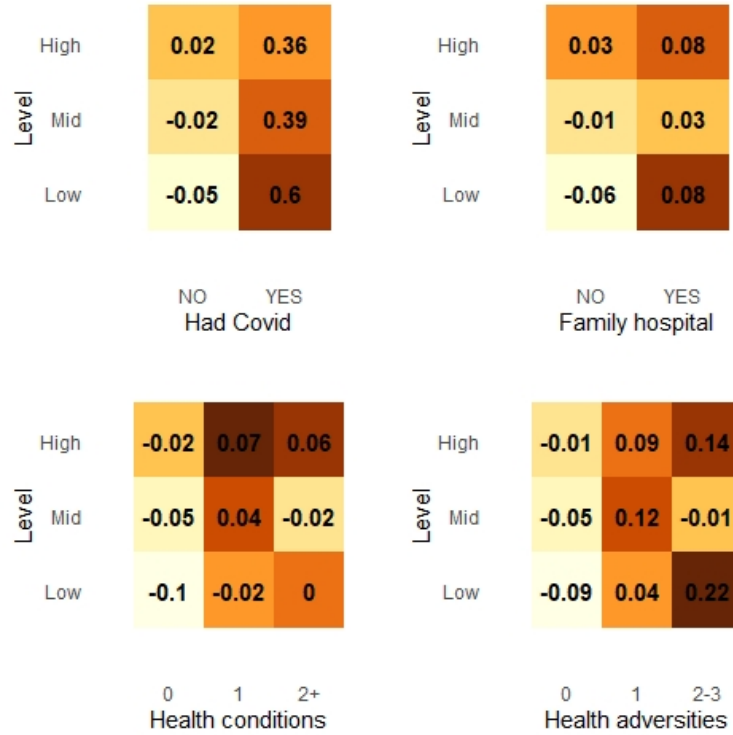
Figure 4.

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.

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 of a standard deviation in beliefs. For respondents who have had Covid, the same change in "Level" is actually associated with a reduction in pessimism. Having had Covid strongly interferes with the local lethality experiences measured by "Level". By far the most drastic 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. This effect captures interference from "Level" to own Covid experiences.

The upper right panel illustrate 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, having a family member hospitalized in a state in the bottom "Level" tercile strongly boosts pessimism, increasing it by 0.14 standard deviations. On the other hand, having a family member hospitalized and strong viral transmission (top tercile of "Level") has a smaller impact on pessimism, which is interference from "Level" to own non-Covid health adversities.

Interference also holds for 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 northwest and southeast cells tend to have darker colors, 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. 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 interaction 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 our model, whereby the marginal impact of an adversity drops when other adversities are added to the database.[18]

---

[18] 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

Overall, the data support a key feature of memory based beliefs: different people react differently to the same experience due to different, even if irrelevant, pre-existing experiences in their databases. An elderly person, facing a lot of interference, may end up underestimating an unlikely new risk. 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. Beliefs can overreact not by being excessively sensitive to local conditions, as commonly assumed, but by being excessively sensitive to irrelevant past experiences that are useful for simulation.

4.3 Red Haired Americans and Reliance on Simulation θ

Our model suggests an interpretation of the estimated share of red-haired Americans as a proxy for reliance on simulation θ. An alternative view is that "red hair" captures a person's higher cognitive uncertainty or noisier numerical perception, which then causes more severe overestimation of rare events. Our model yields a prediction that helps tell these hypotheses apart.

Proposition 4 *The beliefs of people who rely more on simulation should be more sensitive to their Covid as well as non-Covid experiences. Formally, for any given factor X driving beliefs:*

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

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, $\frac{\partial \hat{\pi}_E}{\partial X} > 0$, and disproportionally optimistic if they experience more sources of optimism, $\frac{\partial \hat{\pi}_E}{\partial X} < 0$. Simulation creates a link between stronger average overestimation and higher (rather than lower) weight attached to memory based signals.

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.

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.
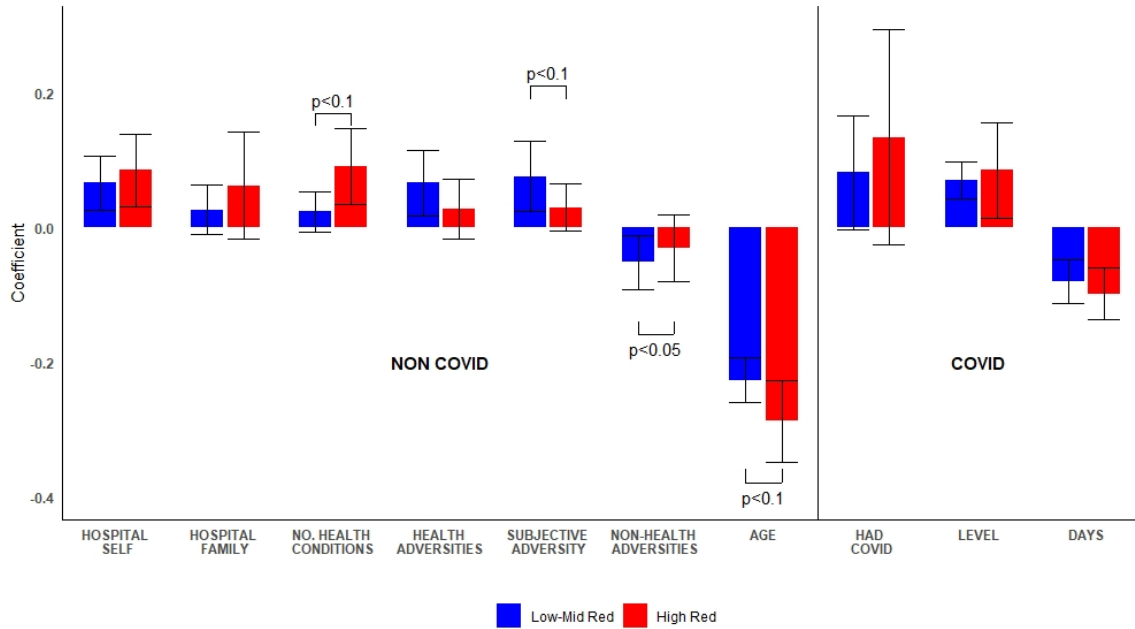
Figure 5

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

33

identical across the red hair groups is rejected.[19] Thus, high red hair does not proxy for a mechanical tendency to report high numbers or adjust estimates upwards. It also does not proxy for greater uncertainty and hence greater insensitivity to information. If anything, Figure 5 indicates that high red hair respondents are more sensitive to experiences. Consistent with our model, higher red hair seems to reflect a stronger weighting of selected past experiences that are similar to the target event, suggesting stronger reliance on simulation.

Our analysis of the age gradient and of the "red hair" proxy throws light on the connection between average overestimation and strong disagreement in Figure 1. Age dampens the impact of specific experiences on beliefs. "Red hair" increases it. Note that red hair and age are almost orthogonal to each other: the correlation between these two covariates is only -0.09.

Our model offers a way to think about their separate roles. Old age means having a database E that is populated by many non-Covid experiences, which create a lot of interference. Based on our previous analysis, this implies that older people should be on average less pessimistic but they should also disagree less, due to the reduced impact of their differential individual experiences. Red hair captures, given a database E, a subject's reliance on simulation θ to form beliefs. Based on our previous analysis, this implies that high "red hair" respondents should on average be more pessimistic and disagree more, due to the heightened impact of their differential experiences.

Figure 6 assesses how consensus overestimation and belief heterogeneity vary with red hair and Age. 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.

---

[19] 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.
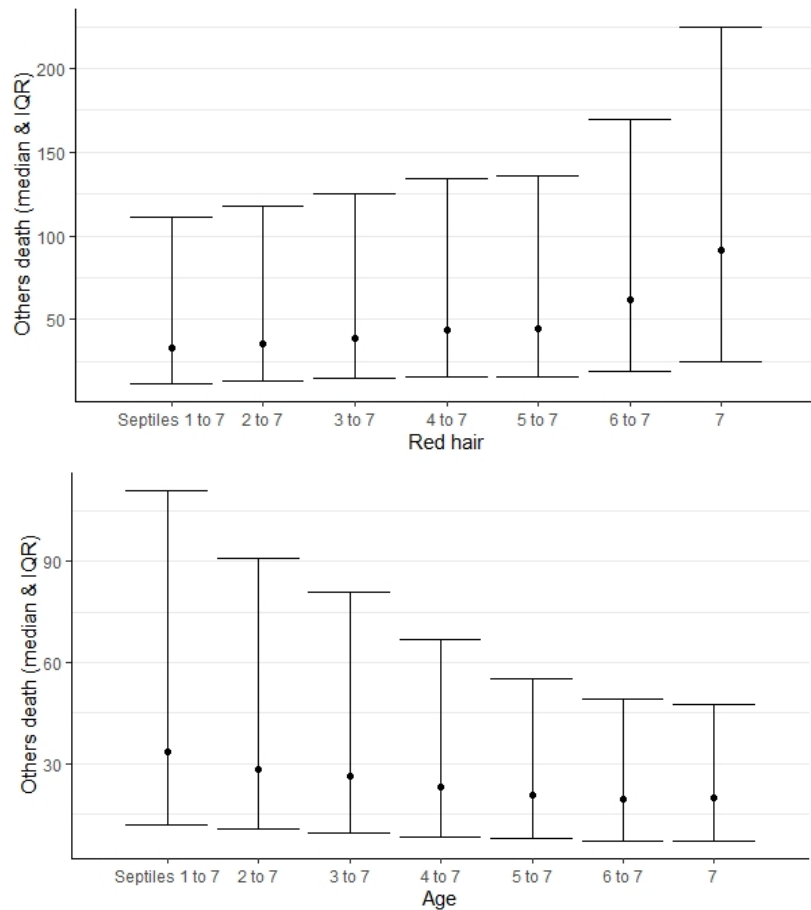
Figure 6.

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

5. Memory, Beliefs and Behavior

Do memory-based beliefs 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.[20]

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 the IV regressions 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 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 ($F \gg 10$ 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:

---

[20] 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.

	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 decided to measure them in surveys 2 and 3, including non-health related ones. The picture that emerges from our research is one in which beliefs about a domain such as Covid depend on a broad range of past experiences, including those from other domains. These experiences, both relevant and irrelevant, affect beliefs because, on the one hand, they provide material to simulate the future and, on the other hand, they interfere with recall of other experiences that are even 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 is also due to systematic underestimation of unlikely events in specific groups, such as the elderly, and to the persistence of these belief differences despite the common experience with even major events.

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 that are new to many people, and in which simulation based on past experiences likely plays a role. We suspect that even in fairly familiar domains simulation and interference can affect beliefs. Our model offers 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 characterize which experiences come to mind and their simulation potential. The tendency to overestimate unlikely events would capture reliance on simulation. In this way, researchers can put even more structure on important memory effects. This could also allow them to unveil new facts, such as the tendency of people from different 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 of many economic decisions. When deciding on a college major or whether to take a new job, a person could greatly benefit from having socially close role models (Conlon and Patel 2022). These are similar and hence foster simulation of success much more than socially distant "artificial" role models or statistical information. A voter assessing a redistributive policy may either selectively retrieve hard working poor, and support it, or free riders, and oppose it. 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. The car rental company Avis successfully advertised itself as an underdog by claiming that, as number two, it tries harder. 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 interfere 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 financial asset based on the thoughts that come to mind. Growing neurological 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 and markets with new models and new data.

References

Abdellaoui, Mohammed, Aurélien Baillon, Laetitia Placido, and Peter P. 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.*

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

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

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

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

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

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

Bordalo, Pedro, Conlon, John J., 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

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 Hoffrage, U. 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, I. and A. 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, Kumaran, Dharshan, Vann, Seralynne D. and Maguire, Eleanor A. 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, G., Witten, D., Hastie, T., and Tibshirani, R. 2013. *An introduction to statistical learning*. New York: Springer

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

Jenkins, J. G., and Dallenbach, K. M. 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.

Kolodner, J. L. 1983. "Maintaining organization in a dynamic long-term memory." *Cognitive science*, *7*(4), 243-280.

Kuchler, T, and Zafar, B. 2019. "Personal Experiences and Expectations about Aggregate Outcomes." The Journal of Finance, 74: 2491-2542

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, Bohm, Vanessa, Ferraro, Simone, Stein, George, and Seljak, Uros. 2021. "Estimating COVID-19 mortality in Italy early in the COVID-19 pandemic." *Nature Communications*, 12(2729).

Mullainathan, Sendhil, Schwartzstein, Joshua, & Shleifer, Andrei 2008. "Coarse thinking and persuasion". *The 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 W., Joel Hellewell, Christopher I. 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

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 $\hat{\pi}_E = \frac{|D|}{|C|} = \pi$. 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:

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

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

$$\frac{|D|}{|C|} = \pi < \pi^* \equiv \frac{\tilde{\sigma}|\overline{D}|}{\tilde{S}|\overline{C}|} + \tilde{\sigma}.$$

Moreover, if non-lethal Covid experiences $\overline{D}$ 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,

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

which implies:

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

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:

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

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 $|E_j \cap E_{-i}|$. After some algebra, one finds that this is equal to:

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

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,

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

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

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

WAVES 2 AND 3 SURVEYS

---

[21] 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.

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

Additional Questions

Waves 2 and 3 feature additional questions focused on personal experiences and activities. These questions are placed after the questions that appeared on the original survey, allowing for cleaner comparisons of answers to the original questions across survey waves.[22]

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

---

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

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

New Treatment Variation

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

Information Experiment

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

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

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

Implementation

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

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

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

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 state, at the time of maximum weekly growth in the state.

55

- Days since Peak (referred to as Peak, also) is the number of days since the peak of cases in the state.

- 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**

58

		(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.

	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 A6

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:

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

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

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 $R^{\wedge}2$ is defined as $\text{adj}R^2 = 1 - \frac{RSS/(n-d-1)}{TSS/(n-1)}$ where TSS is the total sum of squares. The best model is the one which maximizes the adjusted R^2. Finally, we can use m-fold cross-validation. This proceeds as follows: i) divide the sample of n observation in into m non-overlapping groups (folds), each containing around $\frac{n}{m}$ observations; ii) for each $z \in \{1,2,\dots,m\}$ treat fold z as a validation set, fit the model on the remaining folds and compute the mean squared error, MSE_z pertaining to the withheld validation set z; iii) compute $CV_m = \frac{1}{m}\sum_{z=1}^{m} MSE_z$. We will then choose the model with the lowest cross-validation error. What is the best criterion to use is an issue which goes beyond the scope of this discussion. We can refer the reader to Ding et al. (2018). To give a sense of this discussion, in figure A1 we show a comparison of the four decision criteria, applied to the choice of the best model to predict the number of times the respondent had gone out in the period before the survey (table 3 column 1).
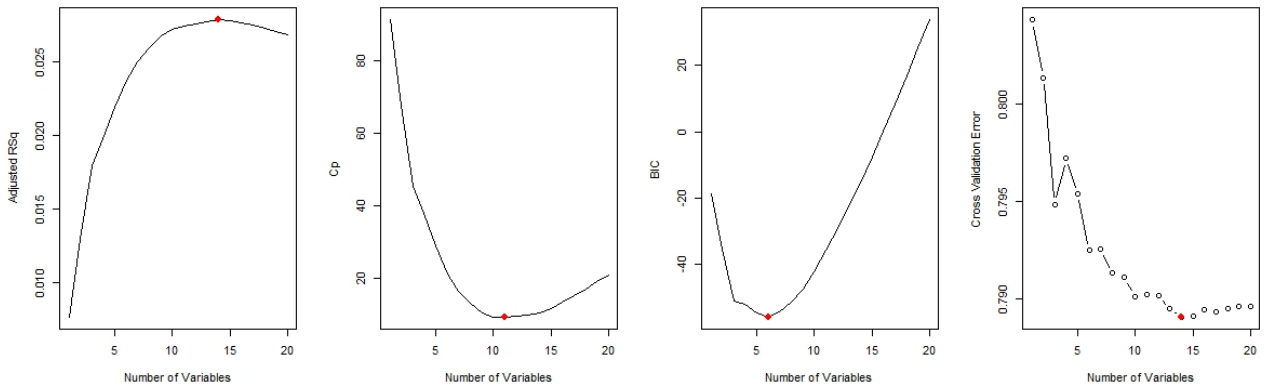
Figure A1

Adjusted R^2, 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.[23] 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.

---

[23] 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.

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 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.[24]

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);

---

[24] 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.

2) We perform some minimal supervision on model selection. We select the model that contains the most robust predictors across the four types of beliefs. This leads to the inclusion of the days since the weekly cases growth peak, and the level of cases in the state of the respondent at the time of maximum weekly growth of cases;[25]

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

---

[25] 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.