NBER WORKING PAPER SERIES

WHICH BELIEFS?
BEHAVIOR-PREDICTIVE BELIEFS ARE INCONSISTENT WITH INFORMATION-BASED BELIEFS: EVIDENCE FROM COVID-19

Ori Heffetz
Guy Ishai

Working Paper 29452
http://www.nber.org/papers/w29452

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2021

We thank Ned Augenblick for substantially contributing (including financially) to this project in its early stages, Tal Asif and Lev Maresca for excellent research assistance, Steven Woloshin for helpful comments and data sharing, and Hunt Allcott, Michele Belot, Wändi Bruine de Bruin, Bnaya Dreyfuss, Ofer Glicksohn, Aharon Haver, Charles Manski, Ted O'Donoghue, Alex Rees-Jones, Germán Reyes, Paul Windschitl and participants in Cornell's BERG, HUJI, Purdue, TAU Coller, the 36th IECA Conference, and USC's CESR COVID-19 Work In Progress Conference for helpful comments. We are grateful to the Cornell Center for Social Sciences (CCSS) for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Which Beliefs? Behavior-Predictive Beliefs are Inconsistent with Information-Based Beliefs: Evidence from COVID-19
Ori Heffetz and Guy Ishai
NBER Working Paper No. 29452
November 2021
JEL No. D83,D84,D91,I12

ABSTRACT

We investigate the relationship between (a) official information on COVID-19 infection and death case counts; (b) beliefs about such case counts, at present and in the future; (c) beliefs about average infection chance—in principle, directly calculable from (b); and (d) self-reported health-protective behavior. We elicit (b), (c), and (d) with a daily online survey in the US from March to August 2020 (N   13,900).

We have three main findings: (1) beliefs elicited as infection case counts are closely related to present and future official case-count information; however (2) beliefs elicited as risk perceptions —i.e., the chance to get infected—are inconsistent with those case-count beliefs, even when mathematically, they should be identical; notably, (3) it is the latter—the risk perceptions—that are significantly better predictors of reported behavior than the former.

Together, these findings suggest that researchers and policymakers, who increasingly engage in direct elicitation and communication of numeric measures of uncertainty, may get very different outcomes, depending on which measures they use. We discuss potential implications for public communication of health-risk information.

Ori Heffetz
S.C. Johnson Graduate School of Management
Cornell University
324 Sage Hall
Ithaca, NY 14853
and The Hebrew University of Jerusalem
and also NBER
oh33@cornell.edu

Guy Ishai
Hebrew University of Jerusalem
Bogen Family Department of Economics
and Federmann Center for the Study of Rationality
guy.ishai@mail.huji.ac.il

An online appendix is available at http://www.nber.org/data-appendix/w29452

To what extent are beliefs under uncertainty based on available information? To what extent do they predict actions? These questions, always at the core of economics, become questions of life and death during pandemics. Contributing to the literature that studies beliefs by direct elicitation (recently reviewed by Manski 2018), we present new survey evidence that their answers may change dramatically depending on *which* beliefs are studied, i.e., on *how* they are elicited.

Our survey takes advantage of the unusual ubiquity and unique standardization, early in the pandemic, of communicated official information about the spread of COVID-19. For several months in 2020, confirmed local daily infection and death counts appeared saliently and frequently in the media and on official websites, and were closely monitored and discussed by the public. By comparing individuals' reported beliefs against such official information benchmarks, we investigate the extent to which beliefs are information-based; and by examining the correlations of these beliefs with reported behavior, we investigate the extent to which beliefs predict actions.

We have three main findings: (1) beliefs elicited as infection *case* counts are closely related to present and future official case-count information; however (2) beliefs elicited as *risk* perceptions—i.e., the chance to get infected—are inconsistent with those case-count beliefs, *even when mathematically, they should be identical*; notably, (3) it is the latter—the risk perceptions—that are significantly better predictors of reported behavior than the former. Together, these findings suggest that researchers and policymakers, who increasingly engage in direct elicitation and communication of numeric measures of uncertainty, may get very different outcomes, depending on *which* measures they use.

Section 1 describes our data. We use a daily online survey of US adults on Amazon MTurk, from March 24, 2020 to August 24, 2020 ($N = 13,880$), to elicit a set of COVID-related beliefs and reported behaviors. We merge the survey data with daily official state-level infection and death case counts. Figure 1 (next page) summarizes our main results, which we investigate in detail in Section 2. All quantities are reported as percentages on log scales.

Figure 1: Information, Beliefs and Behavior (March–August 2020)

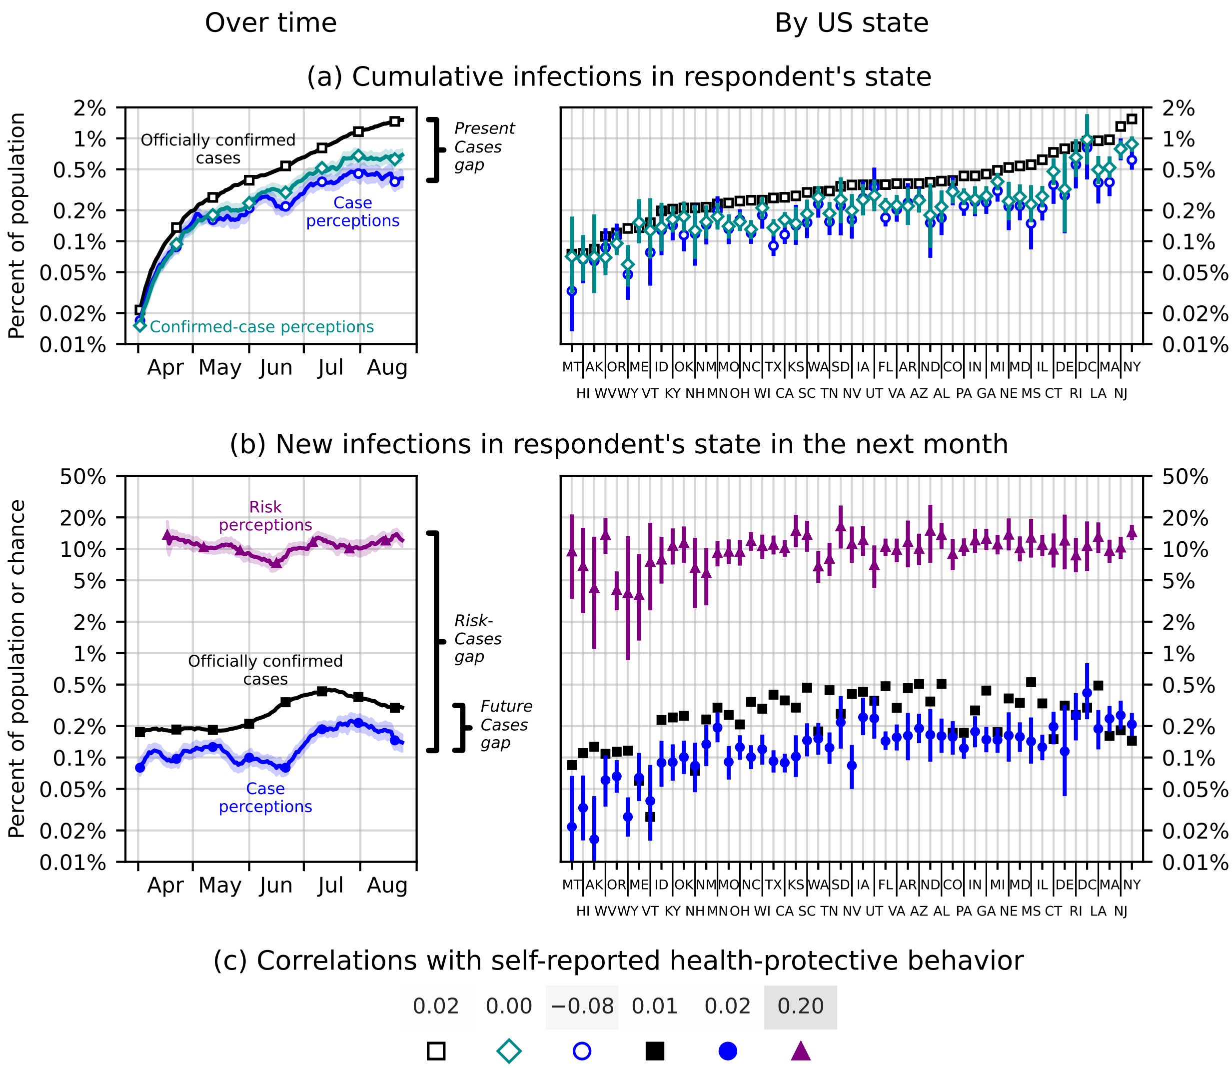

Legend: (quantities defined at the survey-response level, referring to respondent's state population)

Official data:

□ Cumulative percent of population confirmed as infected as of today (the observation day).

■ Future percent of population newly confirmed as infected during the next month (30 days).

Survey data:

◇ Perceived cumulative percent of population confirmed as infected as of today.

○ Perceived cumulative percent of population (actually) infected as of today.

● Predicted percent of population to get infected during the next month.

▲ Predicted average chance of a person (from the state) to get infected during the next month.

Notes: Panels (a), (b): All quantities are represented on log-scale vertical axes as percent of state population or percent chance. Each observation is transformed using $\log\left(\left[1+\frac{x}{100}\cdot(\text{state pop.})\right]/\left[1+\text{state pop.}\right]\cdot 100\right)$, the transformed observations are averaged, and averages are exponentiated (Appendix D.1 shows that all results are robust to this log transformation). Over time: 10-day moving (weighted) averages. Light-colored areas and error bars: bootstrapped 95% confidence intervals. Panel (c): Pearson correlations between the log-transformed percentages (by color/shape) and the number (0–9) of reported health-protective private behaviors (e.g., washing hands more often).

We first investigate the relationship between beliefs about the number of COVID-19 cases and official case-count information. We elicit beliefs about current and future numbers of infections in each respondent's US-state of residence. Panel (a) compares beliefs about current cumulative cases (blue hollow circles) to the officially reported numbers (black hollow squares). Panel (b) compares predictions about *new* cases in the next month (blue solid circles) to official new-case numbers (black solid squares); these two new-case measures are constructed by subtracting, respectively, believed and official present cumulative cases from predicted and future-reported cumulative cases.[1]

In a perfect-measurement, perfect-information benchmark, circles and squares should coincide in panel (a) and furthermore, at the individual-observation level, be perfectly correlated. Assuming rational expectations, circles and squares should also coincide *on average* in panel (b), with no predictions regarding correlations.

Finding 1. *Case perceptions—perceptions and predictions of numbers of infection cases—generally follow official numbers, and somewhat understate them.*

While case perceptions in panels (a) and (b) understate (and possibly lag behind) the official numbers by 49 and 52 percent on average—henceforth, *Present Cases gap* and *Future Cases gap*, respectively—they remain in the same order of magnitude, and they generally follow the official-information time and state trends. Furthermore, case perceptions in panel (a) are moderately correlated with official numbers: $r = 0.41$.[2]

We show this finding's robustness across respondent subpopulations and (randomized) question order in Section 2.1. We also show there that elicited perceptions of current *officially confirmed* cases "reported by the authorities" (cyan diamonds in panel (a)) are close to the above perceptions of current *actual* cases (circles)—suggesting belief in neither under- nor

[1] We use the terms "beliefs" and "perceptions" interchangeably, and sometimes use the term "predictions" for beliefs that concern future outcomes.

[2] Figures for elicited and official *death* numbers, rather than *infection* numbers, are reported in Appendix C.1. They show an even tighter relation between official information and perceptions: the Present Cases and Future Cases gaps are, respectively, 32 and 14 percent understatements; the correlation between perceived current deaths and information is 0.54.

over-detection/reporting of COVID cases. Perceived cases (circles) are on average 17 percent lower than perceived confirmed cases (diamonds), a robustly small gap with unstable sign; the two are strongly correlated: $r = 0.68$.

Second, we investigate the relation between *case* perceptions and *risk* perceptions. We elicit a set of perceptions of medical and economic risk. To cleanly compare these risk perceptions—elicited as probabilities—and the above case perceptions—elicited as case counts—the set also includes risk perceptions that are in principle *mathematically equivalent* to the case perceptions represented by the solid circles in panel (b): the perceived state-average infection risk in the next 30 days (purple solid triangles). In an internally-consistent-beliefs benchmark, solid circles and triangles should be identical and, at the individual level, be perfectly correlated.[3]

Finding 2. *Risk perceptions are inconsistent with their in-principle-mathematically-equivalent case perceptions, and overstate them by orders of magnitude on average.*

Risk perceptions (triangles) are 79 times higher on average than case perceptions (circles)—henceforth, *Risk-Cases gap*. The correlation between triangles and circles, $r = 0.13$, is also surprisingly far below the perfect-correlation benchmark.

Notably, we find this Risk-Cases gap in spite of a survey design that arguably facilitates consistent reporting of case and risk perceptions. For example, the survey interface accepts case perceptions as either an absolute number (of people) or a percent (of the population), simultaneously translating between the two formats and displaying both as the respondent types in one or the other; risk perceptions are entered as a percent (chance) (for screenshots, see Figure 2 on page 12). This gap remains large both over time and across states; in Section 2.2 we find that demographic characteristics and question order can explain only a

[3] The above statement is true under the assumption that when asked about perceived/predicted cases (solid circles), respondents report *averages*, rather than other measures of central tendency. (If person i in a state with population N gets infected with probability p_i, then the state-average infection probability (triangles) is $\frac{1}{N}\sum_{i=1}^{N} p_i$. This expression equals the *average* realized fraction of infected people in the population.)

small fraction of it.[4]

Finally, to investigate the relation between beliefs and behavior, we elicit self-reports of risk-mitigating actions. Beginning in June 2020, we asked respondents about adopted health-protective behaviors, such as washing hands more frequently or avoiding crowds (for full details on the evolution of our survey design, see Section 1.1).

Finding 3. *Self-reported health-protective behavior is moderately positively correlated with risk perceptions and uncorrelated with case perceptions.*

Of all the perceptions elicited in our survey, the number of behaviors reported as adopted is best predicted by—i.e., it is most strongly correlated with—the above risk perceptions (triangles; $r = 0.20$), while having close to zero correlation with the (in principle mathematically identical) case perceptions (solid circles). Panel (c) reports these correlations (see Appendix C.8 for full correlation tables).

This finding also generally holds within respondent subpopulations and is little affected by controlling for demographic variables (including state and day fixed effects) and other specification variations. In terms of magnitudes, the bottom risk-perceptions decile report average perceived infection risk of 0.2 percent—which is rather close to official benchmarks—while engaging with 4.0 protective behaviors on average. In comparison, the top decile report perceived risk of 66 percent—which is wildly unrealistic—while engaging with 5.8 protective behaviors. This large difference in behavior, of 0.8 standard deviations, could make a big difference in a pandemic.

Our survey design and sample size allow us to conduct many other robustness checks, investigate subpopulations, and report a rich set of additional findings. We briefly summarize them throughout Section 2, with details relegated to the appendix.

Why do we find risk perceptions essentially unrelated, in both levels and correlations, to

[4] Together, our first and second main findings also imply that elicited risk perceptions (triangles in Figure 1) overstate the official information benchmark (solid squares) by orders of magnitude. In isolation, this gap could in principle be rationalized as reflecting a belief in massive under-detection/reporting of actual cases, but that would be inconsistent with the finding that actual and confirmed cases (hollow circles and diamonds) are perceived to be relatively close.

case perceptions (and to official benchmarks)? While theoretically equivalent, case and risk perceptions are elicited using different survey questions. Their differences must therefore be related to differences in elicitation details, which we explore in Section 3. First, exploiting our dual-format interface, we find that differences in response format (counts versus percentages) are likely important. Second, using an additional survey (conducted in early 2021, $N = 1,530$), we find that a case-perceptions question that we modify to be more similar to a risk-perceptions question in wording, response format, and structure, elicits perceptions that look strikingly similar to risk perceptions: unrelated to official benchmarks but correlated with behavior. We cautiously conclude that in our context, elicited beliefs may depend more on a question's wording, response format and structure—the question's "look and feel"—than on the underlying mathematical concept the question asks about.

We discuss related literature, implications of our findings, and open questions in Section 4. We first show that our first and second main findings appear consistent with a psychological literature investigating deviations between perceptions about probabilities and relative frequencies (e.g., Gigerenzer and Hoffrage 1995). We then relate our third main finding to a literature investigating the relation between health-risk beliefs and protective behavior (e.g., Brewer et al. 2007). To the best of our knowledge, we are the first to compare the association with behavior of beliefs elicited as probabilities versus as relative frequencies—two elicitation forms routinely used in economic studies and thus of special interest to economists. Moreover, we provide an "all-included" investigation, within a single study, of the very different associations these beliefs have both with information benchmarks—something that a vast psychological literature has documented—and with reported behavior—a novel contribution.

Finally, we discuss a potential implication for public communication of risk: depending on policy goals, policymakers may want to reconsider the case-count language that was so prominently used early in the pandemic. To demonstrate our point, we focus on one much-discussed application: partisan differences in beliefs and in risk-mitigating behaviors in the US during COVID-19. A growing body of recent work (e.g., Allcott et al. 2020, Barrios and Hochberg 2020, Bruine de Bruin et al. 2020, and Fan et al. 2020) finds that relative to Republicans, Democrats consume different news, perceive COVID as riskier, and engage in more social distancing—suggesting that information interventions may be effective in reducing such differences. We replicate past findings, but also find that *within* each political group, behavior is still much more strongly correlated with risk perceptions—which appear out of touch with reality—than with either case perceptions or official case information. Hence, to the extent that our correlations imply causation, policies improving *case-count* communication may have limited behavioral effects, while policies directly targeting risk perceptions, perhaps through directly communicating infection *chances* or *population percentages*, may be more effective.[5]

We conclude in Section 5, discussing broader implications as well as future directions. The disconnect we find between differently elicited beliefs calls into question researchers' ability to easily and reliably elicit beliefs using standard survey questions. However, viewing our findings as mainly demonstrating measurement issues in belief elicitation misses the bigger picture. Returning to the motivating questions we opened with, we ultimately find in our data a weak relation not only between differently elicited beliefs but, importantly, between observable objects: people's *information*—their "input"—and *behavior* (though self-reported)—their "output." Our study, which compares beliefs with both information and reported behavior in a single, real-world, high-stakes setting sheds light on *this* disconnect as well. Echoing vast literatures in psychology (reviewed in Section 4), our results may call into question the idea, still the standard among economists, that beliefs—the connecting link between information and behavior—should be modeled as a single object.

[5] However, our finding that those with highest risk perceptions, who appear to grossly overstate actual risk, engage in more protective behaviors, may create a dilemma for policymakers, because it may imply that such public panic can also have desired behavioral implications. Correcting risk perceptions may thus overall reduce protective behavior.

1 Data

We use an online survey to elicit perceptions and reported protective behaviors, and a public data source to retrieve counts of COVID-19 confirmed cases and deaths.

1.1 Online Survey

Survey Design. Table 1 shows shortened versions of the survey questions and summarizes details regarding question order and response format. (See Appendix A.1 for all screenshots and Appendix B for all response distributions.) The survey consists of six modules, A–F. They are preceded by an entry question that elicits current US state of residence (or DC), and are followed by a final screen that includes demographic and exit questions. The six modules' internal order is, with equal probabilities, *F*–A–(B↔C)–D–E or A–(B↔C)–*F*–D–E or A–(B↔C)–D–E–*F*, meaning that (i) the reported-behavior module F can be first, in the middle, or last; (ii) B–C and C–B are equally likely; and (iii) the other modules' order is fixed.[6] Section 2 and Appendix D.2 investigate order effects and find that no single order substantially changes our main results. In the rest of the paper we therefore pool the data across all survey orders.

As the table conveys, the survey structure generally encourages respondents to think about the elicited constructs as related to one another. As a consequence, our results should be viewed as an upper bound on the strength of the relations between information, different types of perceptions, and behavior.[7]

Finally, to aid readers of our paper in judging how compelling our evidence is, we provide information on the evolution of our survey design. We started collecting data early in the

[6] A small subsample was given an order A–(B↔C)–*E*–D (before F was added to the survey), to test the effect of the distance between modules B and E on the Risk-Cases gap (for results, see Section 2.2; for more details, see Appendix A.2; for randomization balance tests, see Appendix A.4).

[7] We do not incentivize respondents to report accurate beliefs. Rather, we ask them to "answer truthfully" and, when eliciting perceptions about publicly available information (in module A), to "answer without looking up the information." Respondents therefore have no incentive to "cheat" by looking up these numbers, and we find no evidence that they do (for example, only 13 percent of module-A responses are within 5 percent of the official counts). Importantly, such "cheating" would have only affected the Present-Cases-gap part of our first main finding, but neither its Future-Cases-gap part nor our second and third main findings.

Table 1: Survey Design

Module		Question		Timing	Format	Order	Comments
Confirmed-case perceptions	A	How many people in *[state]* have been **reported** by the authorities as [**infected** [A1] / **dead** [A2]] due to the coronavirus *[timing]*?		as of today	# or %	A1–A2 (fixed order)	Page 1.
Case perceptions	B + C	How many people in *[state]* will have [been **infected** [B]] / [**died** [C]] due to the coronavirus *[timing]*? (number may differ from the one reported by the authorities)	(1)	as of today	# or %	2×2: B–C or C–B, (1)–(2)–(3) or (3)–(2)–(1)	B (C) on page 2 and C (B) on page 3.
			(2)	a week from now	# or %		
			(3)	a month from now	# or %		
Risk perceptions + Predicted well-being	D1	What is the chance that **you or your immediate family** will suffer **bad medical outcomes** due to the coronavirus *[timing]*?		in the next month	%	D1–D2–D3 or D3–D1–D2	Page 4.
	D2	What is the chance that **you or your immediate family** will **lose your jobs or run out of money** due to the coronavirus *[timing]*?		in the next month	%		
	D3	What is the anticipated **well-being** of **you and your immediate family** *[timing]*?		in the next month	0–100		
	E1	What is the chance that **you** will get **infected** *[timing]*?		in the next month	%	E1–E2 or E2–E1	Page 5 (pages 4–5 in a complementary version). E1 and E2 asked since days 15, 25 respectively.
	E2	What is the **average chance** of a person in *[state]* to get **infected** *[timing]*?		in the next month	%		
Self-reported health-protective behavior	F	Which of the following have you done *[timing]* to **keep yourself safe from coronavirus**? - 9 private-domain questions about hand-washing frequency, cleaning habits and cautious touching / breathing habits. - 3 public-domain questions about going out and meeting others.	(1) – (12)	in the last week	Yes / No	F1–F12 (fixed order)	Page 6 (before page 1 or between 3 and 4 in complementary versions). Asked since day 88.

Notes: Design details of main survey modules. *[state]*: US state of residence (self-reported in survey intro). Modules A–C are answered using a dual-format (# or %) interface; see Figure 2. Question order is sometimes randomized within modules; see Order column. For full survey text and screenshots, including behavior and demographic questions, see Appendix A.1. For details about complementary survey versions, see Appendix A.2.

pandemic, using modules A–D, to investigate general relations between elicited beliefs and the newly ubiquitous, uniquely standardized, official COVID-19 case-count information. We immediately noticed a striking gap between respondents' general health-risk estimates in question D1 and predicted infection rates in module B. To investigate it, starting on the 15th day of data collection, we added module E (which always appears to respondents only *after* the original modules A–D). At first it contained only a personal-infection-risk question, E1. As the gap remained, ten days later we added a state-average-infection-risk question, E2, to rule out several potential explanations (e.g., private health-risk information). Finally, two months later we added module F, with self-reported-behavior questions, to investigate the relation between actions and the beliefs in modules A–E. The table's Comments column reports the timing of added modules and questions.

Response Format. Figure 2 reproduces the survey's perceptions-elicitation screens. Panel (a) shows that case perceptions (A–C) are elicited using a dual-format interface. Respondents are asked about absolute case numbers ("How many people..."), but can choose to enter either an absolute number or a percent of the state population. As they type in, their response is simultaneously translated into the other format and saliently displayed in both formats. Panel (b) shows that risk perceptions (D–E) are elicited using simple textboxes for entering percent chance. This combination of interfaces encourages respondents to recognize (i) the equivalence between number and percent in the case-perceptions questions and (ii) the link between case perceptions and risk perceptions, as both are displayed (and possibly also entered) in percent.[8]

Respondent Population. Data were collected daily for 5 months, from March 24 to August 24, 2020 (154 days in total). The daily mean number of responses is 92 (SD = 26),

[8]For completeness, we mention that two other elicitation formats are used in the survey: anticipated well-being (D3) is elicited with a simple textbox for entering a number from 0 to 100; protective behaviors (F) are elicited using a set of Yes/No checkboxes.

Figure 2: Screenshots of Perceptions Elicitation Interfaces

(a) Case-perceptions questions

(b) Risk-perceptions question

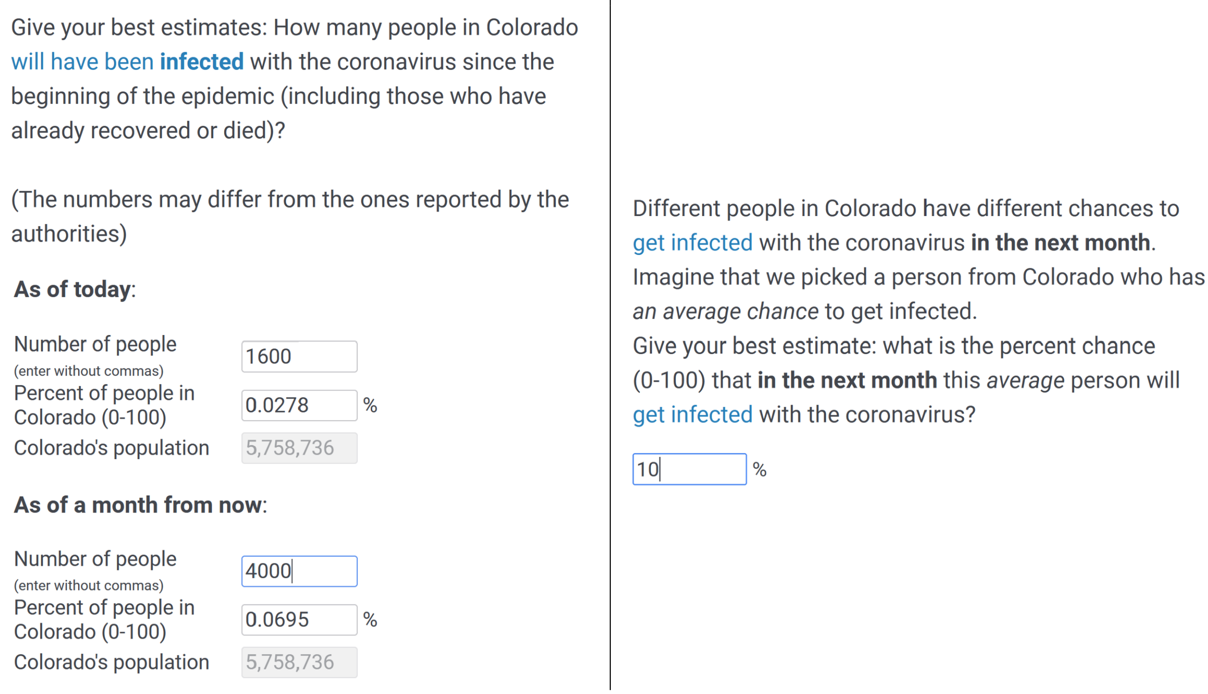

Notes: Panel (a): perceived cumulative number/percent of state-level COVID-19 infection cases as of today (question B1) and a month from now (B3); the difference (B3 − B1), new infections in the next month, is used to construct the case perceptions (solid circles) in Figure 1. A similar dual-format elicitation interface—simultaneously translating numbers to percentages and vice versa—is used throughout modules A–C of the survey. Panel (b): perceived state-average infection risk in the next month (E2); risk perceptions (triangles) in Figure 1. A similar elicitation format, using percent chance, is used throughout modules D–E.

and median survey duration is 5 minutes.[9] We recruited respondents on Amazon MTurk by posting on the platform, each day typically around noon ET, a task paying \$0.70. We set a minimum MTurk-experience criterion and screened out those from outside the US, following the protocol suggested by Kennedy et al. (2018), to minimize low-quality responses. For more details, see Appendix A.3.

From March to June 1, 2020, 6,327 (unique) respondents completed the survey. As we observed sign-up slowing down, starting on June 2 we allowed past respondents (as of June

[9]March 30 and June 12 had very few responses due to a human error in publishing the survey. Excluding those dates, the minimum number of observations per day is 48 and the maximum is 241. For a histogram and distribution of daily responses, see Appendix A.4.

2) to participate one more time; we collected 7,840 additional responses. In Section 2 we use the resulting partial panel to further show the robustness of our findings.[10]

The sample is not US-representative; it is younger, more educated, and more liberal-leaning. However, it is fairly broad, heterogeneous, and representative of all US states. For descriptive statistics, see Appendix B.

Raw, Full and Main Samples. Our raw sample includes 14,167 full survey responses. We exclude observations that (1) managed to bypass the single-participation restriction and respond more than once before or after the restriction reset (210 observations, 1.5 percent), (2) report percent values below 0 or above 100 in at least one question (65 obs., 0.5 perc.), or (3) report an age below 18 (14 obs., 0.1 perc.). This generates our *full sample* of 13,880 observations. Additionally, since our main analysis focuses on respondents' predictions of *new* infections in the next month—constructed as the difference between their predicted *cumulative* infections a month from now (question B3) and today (B1)—we exclude 724 responses (5.2 percent of the full sample) with a negative difference. To verify that this exclusion does not drive our results, we also conduct several versions of our main analysis on the full sample and, as we report in Appendix D.3, none of the results we examine is meaningfully impacted. The resulting *main sample* consists of 13,156 responses by 10,538 unique respondents, of whom 2,618 (25 percent) responded twice (once before and once after June 2, 2020).

[10]Difficulty to reach a respondent is an often unobserved characteristic that may affect survey results (Heffetz and Rabin 2013). In our context, respondents who participate twice (a quarter of all participants) may be easier to reach than those who respond only once; they may also be more experienced in (or bored by) answering our survey in their second time. However, we find that our main results are similar across one-time and two-time participants both prior to and after June 2 (for details, see Appendix C.10).

1.2 Official Case-Count Data

We use publicly available data from *The New York Times*.[11] Each state×date record includes the cumulative numbers of officially confirmed COVID-19 infections and deaths announced by that day midnight ET. We match each survey response with the infection and death records from the relevant state and date, ET.[12] Past published numbers are sometimes updated later on, but we only use the originally published numbers, since they are the ones that were available when our respondents answered the survey.[13]

2 Results

Our main results, reported in Figure 1, were summarized in the introduction. This section analyzes them in more detail, and summarizes additional results and robustness checks fully reported in the appendix.

2.1 Relation Between Information and Perceptions

Infection and death cases, month- and week-forward predictions. Our first main finding is that respondents' case perceptions are closely related to official numbers. Figure 1 shows that perceptions of confirmed infections at present (hollow circles), and predictions of new infections in the next 30 days (solid circles), are generally close to official figures, understating them by 49 percent (Present Cases gap) and 52 percent (Future Cases gap) on average, respectively. These results are not unique to the perceptions plotted in Figure 1. Appendix C.1 shows that predictions of new infections in the next 7 days are understated by 42 percent; and that perceptions of cumulative confirmed deaths and predictions of new

[11] https://github.com/nytimes/covid-19-data. (In earlier versions of our analysis we got essentially the same results using data provided by The COVID Tracking Project at The Atlantic, at https://covidtracking.com/, but the Project stopped reporting data in March 2021.)

[12] Appendix D.4 shows that whether survey responses are matched with previous-, same-, or next-day official reports makes little difference.

[13] The public dataset was first published on March 28, 2020, hence data for our first four survey dates, March 24–27, may include some ex-post updates applied by March 28.

deaths in the next 7 and 30 days are less understated, by 32, 6, and 14 percent respectively.

Demographics and general platform experience. We explore possible drivers of the Present and Future Cases gaps (see Appendix C.2). We find that they are generally stable not only across state and time, but also across demographic groups. We also find little effect of respondents' previous experience on the MTurk platform, using an additional, time-limited sample ($N = 255$) of Workers with less experience than the baseline sample (described in Appendix A.2).

Module and question order. We find that different randomized orders of the survey modules and questions (see Table 1 and surrounding text) have some meaningful quantitative, but no qualitative effect on the Present and Future Cases .gaps. Both gaps are consistently negative and, importantly, orders-of-magnitude smaller than the Risk-Cases gap (Appendix D.2).

Perceived cases vs. perceived confirmed cases. Perceived (actual) cases and perceived confirmed cases (hollow circles and diamonds in Figure 1) are close not only on average, but also at the individual level: they are identical in 37 percent of responses, and the distribution of differences when they are not is concentrated around zero (for example, in 76 percent of responses, one is at most twice the other). On average, perceived actual cases are 17 percent lower than perceived confirmed cases, suggesting an average belief in over-detection or over-reporting of cases (or both)—a belief that, intuitively, appears to have the wrong sign. But this small average difference, and its sign, are sensitive to randomly assigned order of the case-perceptions questions. In particular, the half of respondents who are asked to predict cases in a time-horizon order of [today]–[in a week]–[in a month] (see Table 1, modules B and C) perceive the above difference to be -9 percent on average, consistent with an average belief in little *under*-detection/reporting; the other half, who are asked in reverse time-horizon order ([in a month]–[in a week]–[today]), perceive the difference to be 37 percent on average. For more details and analysis of all order effects, see Appendices C.3 and D.2.

Importantly for our first and second main findings, perceptions of cases—whether confirmed or actual—are in the same order of magnitude as official case-count information, while risk perceptions are orders-of-magnitude larger than such perceptions.

2.2 Relation Between Case Perceptions and Risk Perceptions

Outliers. Our second main finding is a large Risk-Cases gap, between perceived average infection risk (risk perceptions, solid triangles in Figure 1) and its in-principle mathematical equivalent, percent of the population predicted to be newly infected (case perceptions, solid circles). The large gap is not driven by outliers. Its distribution is symmetric and bell-shaped, and 96 percent of the sample have a positive gap (Appendix C.4).

Demographics and general platform experience. As in Section 2.1 above, we explore (in Appendix C.2) possible drivers of the Risk-Cases gap and find that it varies little across state, time and demographic groups. Among less-experienced MTurk Workers, the gap is in fact 2.1 times larger, on average, than in our main sample (with a 95% confidence interval between 1.4 and 2.8).

Module and question order. We find (in Appendix D.2) that survey order has some effect on the Risk-Cases gap, but it is rather limited. Even when the question about infection cases in 30 days appears immediately before the question about perceived state-average infection risk in the next 30 days, the Risk-Cases gap only shrinks to 0.7 of the average gap (with a 95% confidence interval between 0.3 and 1.0).

2.3 Relations Between Perceptions and Self-Reported Behavior

Controlled regressions. Our third main finding is that behavior is much more strongly correlated with risk perceptions ($r = 0.20$ in Figure 1) than with case perceptions ($r = 0.02$).

Table 2 reports OLS regressions of self-reported health-protective behavior on: perceived state-average infection risk ("Risk perceptions" row); percent of the population predicted to be newly infected in the next 30 days ("Case perceptions"); their interactions with the survey position of the behavior module F, which appeared in the beginning ("Behavior first"), middle ("Behavior middle"), or end of the survey (omitted category); demographic controls;[14] and state and day fixed effects. The baseline dependent variable is defined as the sum of behaviors reported as adopted out of a list of nine private behaviors, i.e., excluding three public behaviors, which may be affected by state regulations.[15] Controlling for state and day fixed effects, an alternative dependent variable in column (6) includes all twelve behaviors, private and public.[16]

The regression results reinforce the main message of Figure 1's panel (c). Under all specifications, the perceived state-average infection risk is dramatically more strongly associated with behavior than the percent of population predicted to be newly infected, both economically (coefficient range = 0.19–0.25 vs. −0.00–0.04) and statistically, although the two elicit, in theory, a mathematically identical concept. A coefficient of, e.g., 0.22 on log perceived state-average infection risk (columns 1 and 3) means that an e-fold increase in perceived risk, or 0.48 standard deviations, is associated with an increase of 0.22 in the number of behaviors adopted, or 0.10 standard deviations. Adding order effects of the behavior module F, demographics and state and day fixed effects, and the three public behaviors (columns 4–6) makes

[14]These include: number of people in household; number of people above 18; gender (male/female/other); Hispanic origin (yes/no); race (White/Black/Asian/Native/other); year born; education (10 categories); marital status (6); employment status (6); economic attitudes (7-point scale from very liberal to very conservative); social attitudes (same 7-point scale); political self-identification (Republican/Democrat/Independent/other/none); combined household income (8 brackets for \$0–\$200,000; or above \$200,000); medical insurance coverage (bad/fair/good); have been infected with COVID-19 (yes/no/prefer not to answer); someone from immediate family has been infected (same options).

[15]Private behaviors: increasing hand-washing frequency relative to pre-pandemic habits by at least 5/10/15 times per day (three separate questions); cleaning or sanitizing incoming mail and deliveries; cleaning or sanitizing groceries; cleaning or sanitizing furniture or frequently touched items; avoiding touching own face; stopping breath when passing near others; coughing into elbow rather than palm. Public behaviors: avoiding contact with people from a high-risk group; avoiding meeting family and friends; avoiding public spaces, gatherings and crowds.

[16]All regressions report Driscoll and Kraay (1998) standard errors, which assumes a heteroskedastic error structure, possible cross-individual correlations, autocorrelation up to some time lag—chosen to be 4 days according to a formula provided by Hoechle (2007)—and using the standard Bartlett kernel.

Table 2: Perceptions and Behavior

Dependent variable: Self-reported protective behavior

	Only 9 private behaviors					All 12 behaviors
	(1)	(2)	(3)	(4)	(5)	(6)
Risk perceptions	0.22		0.22	0.25	0.19	0.24
	(0.02)		(0.02)	(0.02)	(0.02)	(0.03)
Case perceptions		0.02	−0.00	0.01	0.02	0.04
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Behavior first				−0.19	−0.17	−0.27
				(0.06)	(0.06)	(0.08)
Behavior first × Risk perceptions				−0.07	−0.06	−0.07
				(0.02)	(0.03)	(0.03)
Behavior first × Case perceptions				−0.02	−0.03	−0.01
				(0.04)	(0.03)	(0.04)
Behavior middle				−0.17	−0.18	−0.22
				(0.05)	(0.05)	(0.07)
Behavior middle × Risk perceptions				−0.01	−0.01	−0.02
				(0.04)	(0.04)	(0.05)
Behavior middle × Case perceptions				−0.02	−0.02	−0.03
				(0.03)	(0.03)	(0.03)
Constant	4.33	4.89	4.32	4.39		
	(0.06)	(0.05)	(0.07)	(0.08)		
Demographics	No	No	No	No	Yes	Yes
State fixed effects	No	No	No	No	Yes	Yes
Day fixed effects	No	No	No	No	Yes	Yes
N obs.	5398	5398	5398	5398	5398	5397
R^2	0.04	0.00	0.04	0.04	0.14	0.15

Notes: OLS regressions. Dependent variable: number of self-reported health-protective behaviors, out of nine private behaviors (columns 1–5) or twelve private and public behaviors (column 6). Independent variables: Risk perceptions: (log) perceived state-average infection risk in the next 30 days; Case perceptions: (log) percent of population predicted to be newly infected in the next 30 days; Behavior first/middle: survey position of the behavior module (F).
The non-binary interacted variables (Risk perceptions and Case perceptions) are centered around their means (to estimate the uninteracted order effects at the mean perception values). In parentheses: Driscoll-Kraay standard errors using Bartlett's kernel and a bandwidth of 4 days.

little difference. Comparing $R^2 = 0.04$ in column (1) with $R^2 = 0.14$ in column (5)—where all demographics, fixed effects, and order indicators are included—provides another indication that the perceived state-average infection risk has a *relatively* strong predictive power. *Absolutely*, however, these R^2 values show that the bulk of variation in behavior remains unexplained.

Disaggregated behaviors. Appendix C.5 shows that the above results do not depend on the specific way the twelve behaviors are aggregated: they generally hold separately for each

Figure 3: Non-parametric relations between perceptions and behavior

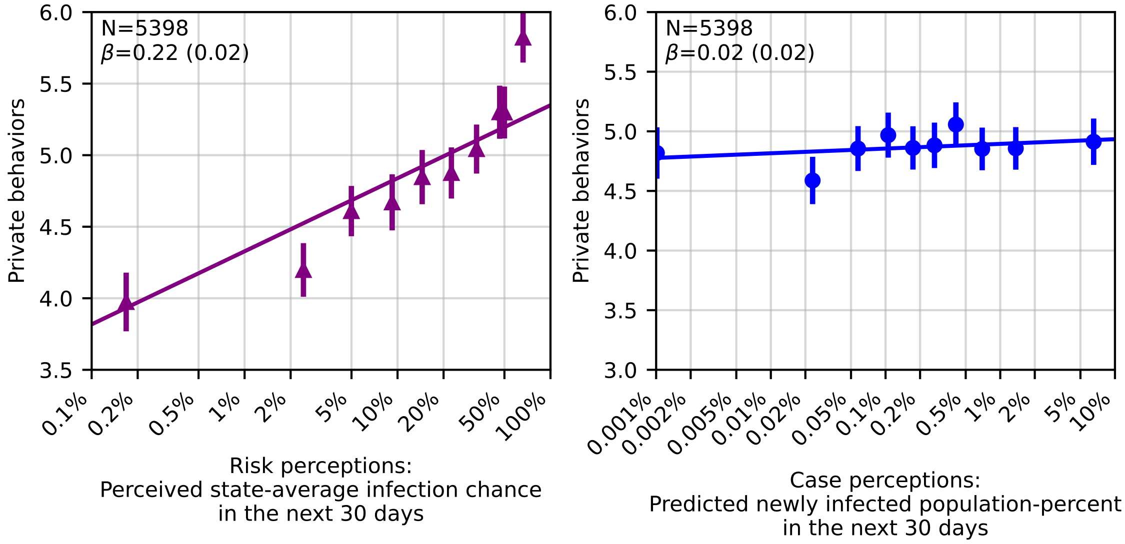

Notes: Triangles and circles: mean number of private behaviors by decile (their horizontal location is the exponentiated mean log percent perceptions within each decile). Error bars: bootstrapped 95% confidence intervals. Lines: OLS-regression estimates from columns (1) and (2) of Table 2. β's: regression coefficients (SEs).

behavior.

Non-parametric relation estimates. Figure 3 shows the regression lines from columns (1) and (2) of Table 2, as well as private-behavior averages by perception deciles. We find no evidence for non-monotonic relations, and the linear correlations and regressions above seem to summarize the (non-parametric) relations reasonably well. In the left panel, the bottom risk-perceptions decile (who perceive on average a 0.2 percent state-average infection risk) report on average 4.0 protective behaviors, while the top decile (66.0 percent) report 5.8—an increase of 0.8 standard deviations. In the right panel the relation is flat.

Behavior/perceptions order effects. While our correlational data do not allow us to identify a causal effect of risk perceptions on protective behavior, our findings can at least rule out the possibility that elicited beliefs are generated ad-hoc to merely match the protective behavior subjects have just reported (e.g., in order to appear consistent or avoid cognitive dissonance). Indeed, columns (4)–(6) in Table 2 show that when the behavior module F is

Figure 4: Correlations of Perceptions and Different Outcomes

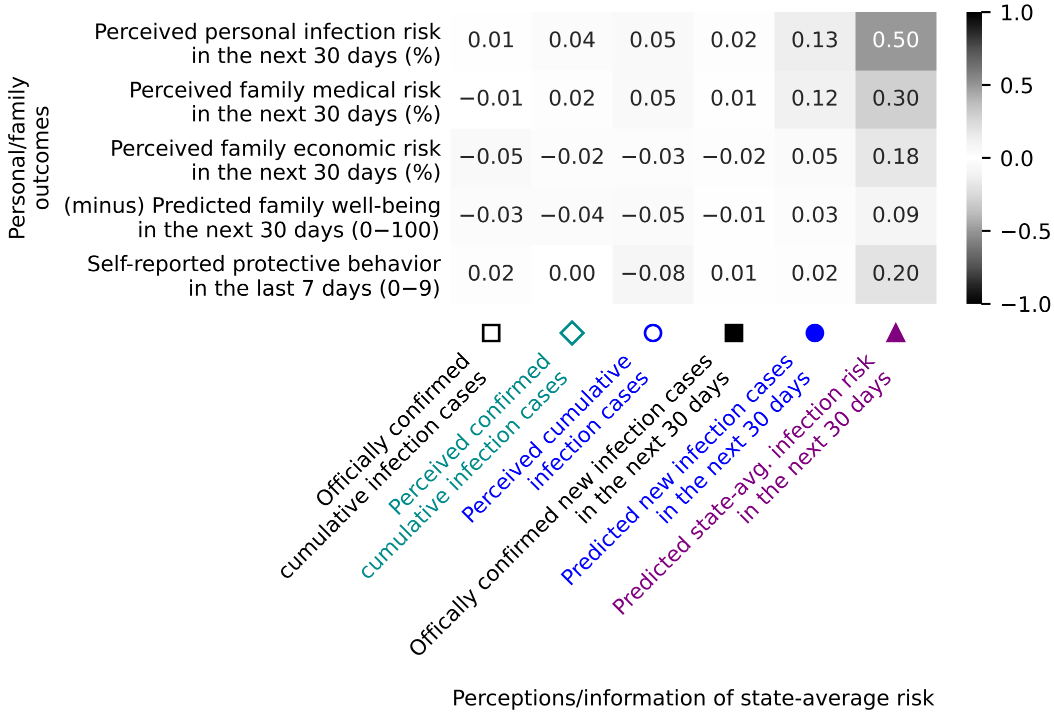

Notes: Correlations of officially confirmed cases, case perceptions and risk perceptions shown in Figure 1 with additional outcome variables, listed on the vertical axis. For full correlation tables see Appendix C.8.

presented first rather than last (the baseline), the coefficient on risk perceptions *decreases* (by 0.06–0.07 (SEs 0.02–0.03)). More generally, that these order effects are so much smaller than the baseline risk-perceptions coefficient (0.19–0.25 (all SEs 0.02)) is reassuring.

Correlations of perceptions with additional personal and family outcomes. Figure 4 extends the correlations bar from Figure 1's panel (c)—replicated as Figure 4's bottom row—to four additional outcome variables: three of personal and family risk perceptions and one of predicted family well-being. The original finding (bottom row), that perceived state-average infection risk is a dramatically stronger predictor of reported personal behavior than case perceptions and official reports, extends to these additional personal/family outcomes (top four rows).

Within-respondents correlations. Appendix C.6 further investigates the relations between perceptions and these four additional outcomes using only *within-respondent* variation.[17] It uses the subsample of respondents who completed the survey twice—once before and once on or after June 2 (N = 5,236 responses by 2,618 respondents). We find that controlling for individual fixed effects, Risk perceptions (as in Table 2) remain a stronger predictor than Case perceptions both when the dependent variable is perceived personal infection risk (the coefficients are 0.72 (SE 0.08) and 0.07 (0.04), respectively) and when it is family medical risk (0.34 (0.08) and 0.03 (0.04)). Risk and Case perceptions are both weak predictors when the dependent variable is family economic risk (0.16 (0.07) and −0.05 (0.08)) or (minus) predicted well-being (0.53 (0.39) vs. 0.19 (0.36)).

Explaining behavior-predictive perceptions with other variables. Can risk perceptions, which predict reported behavior as well as the additional personal and family outcomes in Figure 4, themselves be predicted from other variables? We find in Appendix C.7 that not only are they very far on average from case perceptions and official cases, risk perceptions are also hard to explain using other observables. Using all relevant available variables—officially confirmed cases (at present and newly added in the next 30 days), demographic characteristics, and state and day fixed effects—an $R^2 = 0.08$ suggests that little of the variance in risk perceptions is explained. At the same time, within the subsample of respondents who completed the survey twice, adding individual fixed effects increases the explained variance to $R^2 = 0.81$ (from a baseline $R^2 = 0.14$ in this subsample).[18] This finding suggests that the main determinant of risk perceptions in our data is a stable personal characteristic, consistent with, e.g., Giglio et al.'s (2021) findings in a stock-market context (see Manski 2018 for more findings on intra-personally stable characteristics of beliefs).

[17] We cannot investigate within-respondent variation in self-reported behavior because it was only collected after respondents were allowed to re-participate in the survey.

[18] Case perceptions are also poorly explained by these observables ($R^2 = 0.11$, 0.17 and 0.70, respectively, in the main sample, the responding-twice subsample, and that latter subsample including individual fixed effects); however, recall that case perceptions are on average dramatically closer to official cases at the state and day level. In comparison, perceptions of cumulative infections at present (hollow circles in Figure 1) are explained better ($R^2 = 0.22$, 0.33 and 0.75, respectively).

3 The Role of Elicitation Details

We study the role, in our three main findings, of several elicitation details, including response format and question language. We use our original survey and an additional survey conducted in early 2021.

Role of response format: non-causal evidence. Respondents in our survey report risk perceptions using percent chance (Figure 2b on page 12). However, they can choose to report case perceptions using either percent of population or absolute number of cases (Figure 2a). Respondents who choose to report absolute numbers—92 percent of all responses—may think about cases differently than those who choose to report percentages—the remaining 8 percent.[19] In particular, the latter actively use the same response format (percent) for reporting their case perceptions as all respondents use for reporting risk perceptions. Do they also report case perceptions that look more like risk perceptions?

Figure 5's two leftmost columns show that they generally do. The two columns compare the perceptions from Figure 1 and their correlations with behavior between absolute-number reporters and percent reporters. The middle panel shows that percent reporters indeed have a smaller Risk-Cases gap: their perceived state-average infection risk (triangles) is 14 times larger than predicted newly infected population percent (circles), compared with 92 times larger for absolute-number reporters. Not reported in the figure, percent reporters' case and risk perceptions are also more strongly correlated (0.25) than absolute-number reporters'

[19] Recall, as Figure 2a shows, our dual-format interface provides automatic, real-time, on-screen translation of one format into the other; we use rounding patterns to distinguish responses typed in using either format (see Appendix C.9 for details and robustness).

That so many of our respondents choose to use the absolute-number format is not surprising: in addition to appearing first on the screen (potentially making it a natural default), at the time of our study the absolute-number format was prominently used both by official sources and in the popular media to communicate information about COVID-19 cases. For example, as of August 19, 2020, absolute numbers were the default data (a) on a graph provided by Google following a Google search of the phrases "COVID 19 new cases," "COVID statistics," or similar phrases; (b) on data websites such as www.cdc.gov, www.worldometers.info, and coronavirus.jhu.edu/map.html; and (c) on the most popular news websites, including the following four out of the five most visited websites (according to ebizmba.com, ordered by popularity): Yahoo!, Google News, HuffPost, and New York Times (which reported both absolute numbers and proportions). (The fifth, ranked 4 on ebizmba.com, is CNN. Its default presentation is proportion out of 100,000 people.)

Figure 5: Main Results' Dependence on Elicitation Details

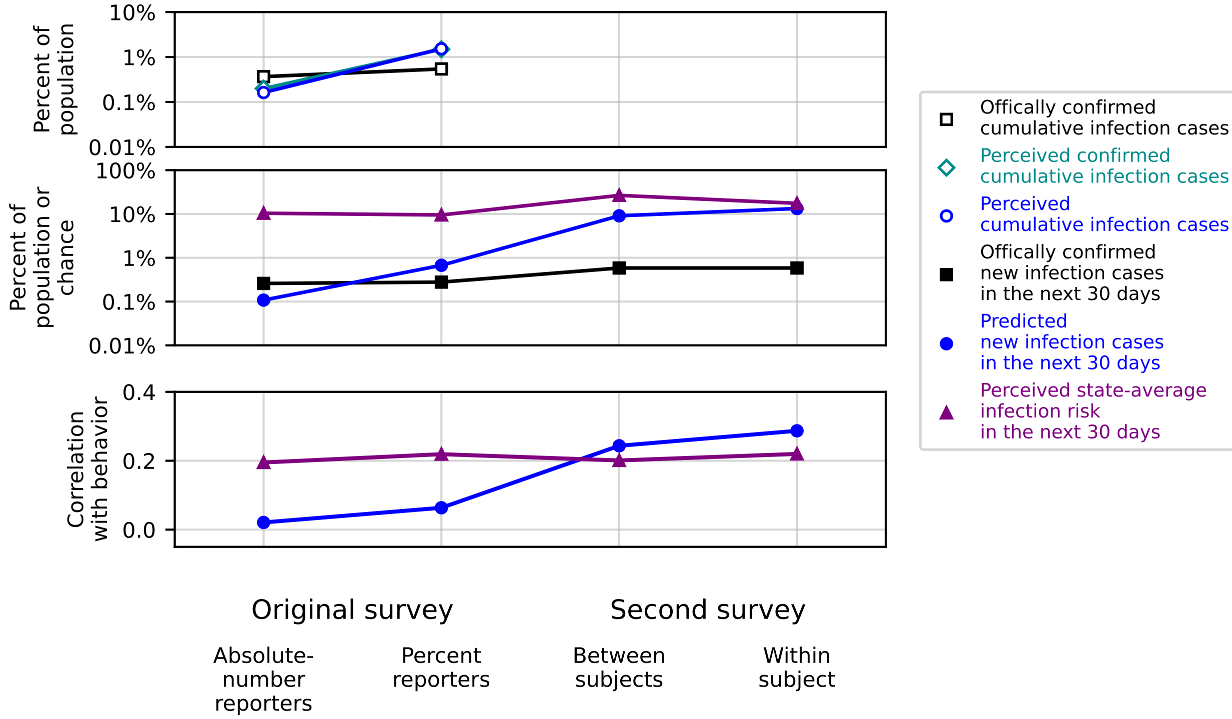

Notes: Two left columns: Figure 1's results, averaged over all states and days, split by respondents' self-selection to report case perceptions using absolute numbers ($N = 12,142$) vs. using percentages ($N = 1,014$) in the original survey. Two right columns: Figure 1's results from a second survey ($N = 1,530$), in which both case perceptions and risk perceptions are reported using percentages and are asked using similar wording. Top and middle panels: average officially confirmed cases, case perceptions and risk perceptions, listed in the legend. Bottom panel: correlation coefficients with self-reported (private) protective behavior. Error bars (top and middle panels; sufficiently narrow that they are hidden behind the markers) indicate bootstrapped 95% confidence intervals.

(0.12). At the same time, risk perceptions (triangles) are essentially identical across the two groups, and the smaller Risk-Cases gap is almost entirely explained by case perceptions (circles), which are higher among percent reporters than among absolute-number reporters. Relatedly, percent reporters overstate, rather than understate, present and future official counts (compare, across the two columns, hollow circles vs. squares in the top panel, and solid circles vs. squares in the middle panel), yielding positive Present Cases and Future Cases gaps (2.8- and 2.4-fold overstatement, respectively, compared with 56 and 58 percent understatement among absolute-number reporters). The bottom panel shows that the correlations between risk perceptions and behavior are 0.22 and 0.19 respectively; between case perceptions and behavior they are 0.06 and 0.02.

These results are qualitatively unchanged when controlling for demographics and fixed effects (see Appendix C.9), implying that the chosen response format is an important explanatory variable, only weakly explained by our other variables. We conclude that while this evidence is correlational—respondents self-select into the absolute-number- and percent-reporting groups—it points to a potential central role for response-format differences in the risk-cases inconsistency.

Joint role of response format and other elicitation details: evidence from a second survey. Due to self selection, causality in the above analysis could run in both directions. In addition, other elicitation differences in our survey between case and risk perceptions may have an even larger role in our main findings. These include, e.g., question wording and style, user interface, case perceptions being based on two questions rather than one, and the salient display of a state's population only in case-perceptions questions. To investigate the overall effect of these elicitation details jointly, we conducted another short survey. It includes only three questions, on three different pages, randomly ordered: case perceptions, risk perceptions, and self-reported behavior (see Appendix A.5 for screenshots).[20] Importantly, case perceptions are elicited more similarly to risk perceptions. Modifications include, e.g., replacing "How many people" with "What percent of the population"; using a simple, single-format (percent) textbox as in Figure 2b for responses; and not displaying total state populations.[21]

Figure 5's two rightmost columns show that in this second survey, case perceptions be-

[20] *Case perceptions question*: Give your best estimate: what percent (0-100) of the population in *[state]* will get infected with the coronavirus during the next month?

Risk perceptions question: Different people in *[state]* have different chances to get infected with the coronavirus during the next month. These chances depend on many things, such as personal circumstances, lifestyle, and behavior. Give your best estimate: what is the average chance (0-100 percent) for a person in *[state]* to get infected with the coronavirus during the next month?

Private behaviors: washing hands often, sanitizing groceries, avoiding touching face; *public behavior*: avoiding public spaces and crowds. (We chose this subset because it is diverse and has $R^2 = 0.96$ in a regression explaining the sum of Yes answers to all twelve behaviors in our main survey.)

[21] The questions were incorporated as a last module in a larger survey on a different topic (consumer expenditures), conducted between February 8, 2021 and March 10, 2021, on an online US sample that matches the adult population on several key demographics ($N = 1{,}530$). See Appendix A.5.

come qualitatively similar to risk perceptions, in both between-subjects ($N = 1,031$) and within-subject ($N = 1{,}530$) analyses.[22] In the middle panel, the Risk-Cases gap (solid triangles vs. circles) dramatically shrinks, from a 79-fold overestimation in the original survey's main sample to mere 2.9- and 1.3-fold overestimation, respectively, between and within subjects, while both case and risk perceptions dramatically overstate official numbers, with a Future Cases gap (circles vs. squares) growing to 16 and 23-fold *over*estimation. In the bottom panel, case perceptions become a bit *more* predictive of behavior than risk perceptions (correlation of 0.24 vs. 0.20 between- and 0.29 vs. 0.22 within-subjects), which remain as predictive of behavior as in the main sample (correlation of 0.20). (Not reported in the figure, case and risk perceptions themselves become highly correlated, at 0.76.)

In summary, our purposefully designed "case-less" case-perceptions question, or "cases-as-percent"-perceptions question, behaves qualitatively like a risk-perceptions question: it elicits gross overestimates of case numbers yet it is correlated with reported behavior. We draw two conclusions, both in line with the main message of this paper. First, highlighting the importance of the "Which Beliefs?" question in the paper's title, elicited beliefs crucially depend on elicitation details that standard economic theory is agnostic about. In our case, we find the differences between case and risk perceptions to be a function of question structure, wording and response format more than of the specific mathematical object the question appears to target. Second, consistent with the rest of the paper's title, the perceptions we elicit in our surveys are related to either information or behavior, but not both.

4 Discussion

We now relate our three main findings to existing evidence from psychology and economics, focusing on health contexts and in particular on COVID-19. We then discuss potential implications for public communication of health-risk information.

[22] For case and risk perceptions, respectively, the between-subjects analysis includes data from the first two pages only, for respondents who saw the survey orders [Case↔Behavior]–Risk and [Risk↔Behavior]–Case.

4.1 Existing Evidence of Inconsistent Beliefs

Our first finding that case perceptions are not far from official information benchmarks—i.e., limited Present and Future Cases gaps—and our second finding of an inconsistency between mathematically equivalent risk and case perceptions—i.e., a large Risk-Cases gap—are related to previously documented biases in people's elicited beliefs. One strand of literature finds that beliefs depend on whether they are elicited as probabilities (e.g., 10 percent or 1 in 10 probability) or as relative frequencies, i.e., as absolute counts in an imaginary sample (e.g., 100 out of 1,000 people). Gigerenzer and Hoffrage (1995) find that some well-documented probability-reasoning biases such as base-rate neglect or the conjunction fallacy significantly shrink in magnitude when both the questions and the responses are communicated using relative frequencies rather than probabilities. Slovic et al. (2000) find that psychiatrists' estimates of the risk that discharged patients will be violent are lower when they use relative frequencies rather than probabilities.

Another strand of literature finds that when asked to explicitly relate probabilities to their corresponding relative frequencies, a large proportion of people do not give the expected answer. For example, Galesic and Garcia-Retamero (2010) find that only 57.7% of a US-representative sample give the answer "10" to the question: "In the Bingo Lottery, the chance of winning a $10 prize is 1%. What is your best guess about how many people will win a $10 prize if 1,000 people each buy a single ticket for Bingo Lottery? ____ person(s) out of 1,000." See also Woloshin and Schwartz (2011).

Our evidence appears generally consistent with the findings of both strands of literature. We find that respondents (a) report more accurate perceptions and predictions using relative frequencies (out of their state's population); and (b) fail to relate population percentages to average probabilities, even when the questions are adjacent. Section 3 highlights that the main drivers of such results may be subtle "look and feel" differences between elicitation questions, e.g., in response format or wording, rather than the *conceptual* difference between probabilities and relative frequencies. Supportive of this view, Bordalo et al.'s (2020) respondents predict that around 5 percent of the population will get infected with COVID-19—getting closer to our respondents' high risk-perceptions levels—even though they are asked about infection *frequencies* (i.e., cases out of 1,000 people) and not about *probabilities*.[23]

4.2 Existing Evidence on the Relationship Between Beliefs and Protective Behavior

Our third main finding—that risk perceptions predict behavior better than case perceptions—contributes to a literature comparing belief-elicitation questions by their power to predict behavior, e.g., Windschitl and Wells (1996), Weinstein et al. (2007), and Dillard et al. (2012). While we only explore *numeric* response scales, that literature also explores response scales with *verbal* descriptions of uncertainty levels, e.g., asking "Without a flu shot, do you think you're likely to get the flu this year?" with six response options: "extremely likely," "very likely," "somewhat likely," "somewhat unlikely," etc. Two main findings of that literature are that eliciting beliefs using such verbal scales often predicts behavior better than using numeric scales, and that questions that involve a language of feelings predict better than questions with more objective language. Within the domain of numeric scales, we are, to the best of our knowledge, the first to compare the predictive power of behavior of relative-frequency perceptions (cases out of a given sample) with that of percent-chance perceptions. Because we compare two objective-language, numeric questions that can be directly interpreted as probabilities, our contribution is especially relevant for economic studies eliciting beliefs.

Another aspect of our third main finding is that the correlation between perceived risk and

[23] Their elicited frequencies are of future infection cases in the next 9 weeks among US subpopulations in May, 2020. Furthermore, their elicitation interface shows respondents, after they responded, their reported relative frequencies in terms of percentages and proportions out of 100,000 people, and allows them to revise their answers, somewhat similarly to our dual-format interface. This evidence may in addition suggest that our dual-format interface is also not the main driver of the relatively small Present and Future Cases gaps we observe.

protective behavior is positive. Theoretically, it could have either sign, change magnitude in different ranges, and even be nonmonotonic.[24] Empirically, that we find a positive correlation is consistent with previous findings in the literature on health-risk perceptions and risk-mitigating behavior, including findings in the COVID-19 context. Among studies similar to ours in their questions and analysis, Brewer et al. (2007) conduct a meta-analysis of twelve studies, mostly longitudinal, and report a pooled positive correlation of 0.26 between perceived flu infection risk (conditional on not being vaccinated) and getting vaccinated; and Allcott et al. (2020) find a positive standardized effect of 0.32 of perceived COVID-19 personal infection chances on self-reported social distancing, where beliefs are conditioned on a hypothetical scenario in which the person maintains a pre-COVID routine, controlling for individual characteristics and state fixed effects.[25] In contrast, Akesson et al. (2020) and Papageorge et al. (2020) find *negative* correlations between risk perceptions and protective behaviors.[26]

Other studies of COVID-19 are less similar to ours, but still find positive relations between risk perceptions and protective behaviors. These include, e.g., Fan et al. (2020), who study only differences between demographic groups, Wise et al. (2020), who only control for age in

[24] In a standard expected-utility model, holding behavior cost and personal characteristics constant, the sign of the protective-behavior response to increased risk depends on the returns to protective behavior. If washing hands, for example, reduces infection risk by a constant multiplying factor, then the higher the baseline risk, the higher the returns to washing hands, generating a positive sign. In contrast, if washing hands is less effective at high infection risk, the sign could be negative (or be first positive and then negative, for example).

[25] Notably, they also find a much weaker positive effect, of 0.07, of predicted future number of US COVID-19 cases on self-reported social distancing—strikingly similar to our relations between risk perceptions, case perceptions and behavior. Also like our findings, their case perceptions underpredict officially confirmed infections—though by a larger factor of 4.4. They use their case- and risk-perceptions measures to study partisan-differences, and do not investigate the inconsistencies between those perceptions.

[26] Studies that are most comparable to ours (a) elicit perceptions of risk that is exogenous to behavior and (b) attempt to identify within-individual correlations.

Criterion (a) is relevant due to a theoretical result, based on a standard expected-utility model, that any correlation sign between *endogenous* risk and behavior is consistent with our result of a positive correlation between *exogenous* risk and behavior (see Appendix E). That result renders comparisons with correlations of the former kind less informative. Criterion (b) is relevant because our analysis attempts to identify, to the extent possible given our data, within-individual correlations: we control for personal characteristics and external costs of protective behavior (by using either private behaviors or controlling for state and day fixed effects, e.g., in Table 2), and we report panel results (with outcomes different from behavior; see Section 2.3). Comparing our study to studies that do not include a rich set of controls, or that do not focus on private behaviors, is therefore less informative.

their analysis, and Bruine de Bruin and Bennett 2020 and Dryhurst et al. 2020, who elicit personal infection risk, which is endogenous to own behavior (in contrast to our exogenous state-average infection risk).

4.3 Implications for Public Communication of Health-Risk Information

The COVID-19 pandemic brought back important policy questions including how to effectively inform the public about risks and, separately, whether and how to communicate risk in order to induce behavioral change (especially when facing strong externalities). Our findings suggest that depending on policy goals, policymakers may want to explore alternatives to the case-count language that was so prominently used early in the pandemic.

The effective-communication question has become especially pronounced in light of large documented differences in protective behavior across demographic groups. In the US, for example, Allcott et al. (2020), Barrios and Hochberg (2020), Bruine de Bruin et al. 2020 and Fan et al. (2020) all find that people who consider themselves Democrats engage in more COVID-protective behavior than those who consider themselves Republicans. They also find that the behavior differences can be partially explained by differences in perceptions and media consumption. Consistent with these studies, we document in Appendix C.2 that (self-identified) Republicans' case perceptions are understated more than Democrats' (Future Cases gaps of 70 and 39 percent underestimation, respectively), and that Republicans perceive lower risks and engage less in protective behavior (4.7 private behaviors compared with Democrats' 5.2).

Such findings may lead policymakers to speculate, for example, that a national campaign providing the public with accurate facts—such as infection and death case counts—may help close such behavioral gaps. However, our second and third main findings, which continue to hold within each partisan group, may suggest otherwise: both Republicans and Democrats show huge Risk-Cases gaps (factors of 92 and 80); and their protective behavior is essentially uncorrelated with case perceptions (-0.04 and 0.00, respectively) while being moderately correlated with risk perceptions (even more so among Republicans: 0.25 and 0.15, respectively). Our findings therefore suggest that communicating accurate official counts more prominently may not be an effective way to close such cross-group gaps and, more generally, to affect behavior: in our data, both official count information and its perceptions are only weakly related to reported behavior.

At the same time, many studies do find *certain* risk-related information-provision policies to be rather effective. Shermohammed et al. (2021), for example, find that informing high-risk influenza patients about their high-risk status increases their likelihood to get vaccinated for flu by 5.7 percent. Outside the health-risk context, Abito and Salant (2019), for example, find that the demand for expensive extended warranties, largely related to exaggerated perceived failure-risk probabilities, decreases when these perceived probabilities are corrected. Such results may, again, depend on *how* information is communicated. Our finding that language matters is consistent with the possibility that, for example, communicating risk referring to *chances* or *population fractions* rather than *cases* may shift people's risk perceptions and behaviors—a possibility that should be further investigated. At the same time, our findings also suggest that to the extent that such communication is effective, it may in fact *reduce* protective behavior. As Figure 3 above shows, in our sample, respondents in the top risk-perceptions deciles, who engage the most in protective behaviors, grossly overpredict (30-day state-average) infection risk at 50 percent or above—arguably a public-panic level. To the extent that our correlations imply causation, correcting these respondents' risk perceptions may reduce their protective behavior and therefore worsen social health outcomes.

Finally, language and framing could also have other effects on behavior that should be considered and further investigated. Freeman et al. (2020) find that people rate the risk of dying of COVID-19 (conditional on infection) higher (on a scale from "very low risk" to "very high risk") when they are informed about the probability using spelled-out-fraction language that is similar to relative frequencies (e.g., a "120 out of 1,000" chance) than when using percent chances (e.g., a "12%" chance). In a different health domain, the Slovic et al. (2000) study mentioned above similarly finds that psychiatrists rate discharged patients' violence risk higher (on a "low"/"medium"/"high" scale) when they are informed about this risk using relative frequencies (e.g., "10 out of 100" patients) than when using probabilities expressed as percent chances (e.g., a "10%" chance). In our context, official communication of COVID-19 risk that uses case counts (i.e., frequencies) may therefore be a contributing *cause* of our finding of the public's exaggerated risk perceptions.

In summary, our evidence, while correlational, suggests that public-communication campaigns aimed at affecting protective behavior may benefit from exploring risk communication that emphasizes chances or population fractions, rather than case counts. Past evidence suggests in addition that public campagins aimed at correcting risk perceptions—irrespective of the potential protective-behavior consequences—may independently benefit from exploring alternatives to case counts. Yet, at the time of this writing—more than a year and a half into the pandemic—the case-count language that became the default way of public communication early in the pandemic appears to remain an overwhelmingly common standard.

5 Conclusion

In this study we elicit two types of forward-looking beliefs: about a population's future infection *case counts* and about the population's average future infection *risk*. While the two are mathematically equivalent, we find that they are reported at levels that are orders-of-magnitude apart and are only weakly correlated across respondents. Furthermore, we find this gap between beliefs to be closely related to an economically meaningful gap between two observable objects: information and protective behavior. Indeed, the correlation bar in Figure 1 shows that it is not only *beliefs* about COVID-19 infection case rates that are essentially unrelated to behavior, but also the *official* infection rate, i.e., the objective information. In our second survey, when replacing case language and numbers with caseless language and percentages, elicited beliefs move unrealistically away from information benchmarks but become correlated with behavior.

One implication of our findings is that merely communicating accurate factual information to people (e.g., about case counts) may be insufficient for affecting behavior even in life-and-death situations. Moreover, in our data, we do not find that such information is grossly misperceived or misremembered—indeed, our respondents report and predict case information with relatively limited (downward) bias. However, this information, while apparently successfully communicated and understood, does not seem to affect behavior, perhaps because it does not sufficiently shape behavior-relevant perceptions (e.g., risk perceptions). This distinction, between information that is merely understood (and remembered) and information that "sinks in" and affects behavior, should be further studied.[27]

Other implications concern researchers' ability to reliably elicit beliefs using standard survey questions. Of course, our findings may merely highlight the imperfection of existing belief-elicitation questions. Other questions that we have not tried may be related to both information and behavior (and, of course, behavior itself may not be perfectly captured by the *self-reported* behavior we elicit). But our findings may reflect deeper issues with elicited beliefs. That our modified case-perceptions question elicits beliefs that so easily move away from those elicited with our original case-perceptions question may support the idea that people report ad-hoc beliefs formed in response to the elicitation context and framing (e.g., Windschitl 2002; Benjamin 2019 reviews some other context dependencies). That our elicited risk perceptions have unrealistic levels and yet they are correlated with protective behavior may support the idea that elicited beliefs sometimes confound probabilities with preferences (see Manski 2018).

Finally, that the belief inconsistency we find is related to an apparent economically

[27] See Heffetz (2021) for a discussion of "sunk-in" beliefs in the context of expectations-based reference-dependent preferences. There, lab participants understand and remember objective probabilities; yet that demonstrably successful communication of information does not seem enough for those probabilities to become participants' reference point and affect behavior.

meaningful disconnect between information and behavior may have theoretical implications that go far beyond belief-elicitation measurement issues. Economic models routinely assume that information affects beliefs, and that beliefs affect behavior. Our findings highlight the possibility that information may affect *some* beliefs, while *some other* beliefs may affect behavior. Our findings may thus be consistent with work that questions the notion, standard in economics, that beliefs should be modeled as a *single* probability distribution over relevant outcomes. Dual process models, such as in Loewenstein et al. (2015), take a step in this direction by allowing beliefs to shape behavior in two different ways, associated with different processes. For example, choice may maximize a mix of "deliberate" utility, based on rational beliefs, and "affective" utility, based on probability-weighted beliefs. To explain our findings using such models, future research would need to investigate the relationships between different theoretically defined cognitive processes and different empirically elicited beliefs.

References

Abito, J. M. and Salant, Y. (2019). The effect of product misperception on economic outcomes: Evidence from the extended warranty market. *The Review of Economic Studies*, 86(6):2285–2318.

Akesson, J., Ashworth-Hayes, S., Hahn, R., Metcalfe, R. D., and Rasooly, I. (2020). Fatalism, beliefs, and behaviors during the COVID-19 pandemic. *NBER Working Paper*, w27245.

Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., and Yang, D. (2020). Polarization and public health: Partisan differences in social distancing during the Coronavirus pandemic. *Journal of Public Economics*, 191:104254.

Barrios, J. M. and Hochberg, Y. (2020). Risk perception through the lens of politics in the time of the COVID-19 pandemic. *NBER Working Paper*, w27008.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:69–186.

Bordalo, P., Coffman, K. B., Gennaioli, N., and Shleifer, A. (2020). Older people are less pessimistic about the health risks of COVID-19. *NBER Working Paper*, w27494.

Brewer, N. T., Chapman, G. B., Gibbons, F. X., Gerrard, M., McCaul, K. D., and Weinstein, N. D. (2007). Meta-analysis of the relationship between risk perception and health behavior: the example of vaccination. *Health Psychology*, 26(2):136–145.

Bruine de Bruin, W. and Bennett, D. (2020). Relationships between initial COVID-19 risk perceptions and protective health behaviors: A national survey. *American Journal of Preventive Medicine*, 59(2):157–167.

Bruine de Bruin, W., Saw, H.-W., and Goldman, D. P. (2020). Political polarization in us residents' COVID-19 risk perceptions, policy preferences, and protective behaviors. *Journal of Risk and Uncertainty*, 61(2):177–194.

Dillard, A. J., Ferrer, R. A., Ubel, P. A., and Fagerlin, A. (2012). Risk perception measures' associations with behavior intentions, affect, and cognition following colon cancer screening messages. *Health Psychology*, 31(1):106–113.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4):549–560.

Dryhurst, S., Schneider, C. R., Kerr, J., Freeman, A. L., Recchia, G., van der Bles, A. M., Spiegelhalter, D., and van der Linden, S. (2020). Risk perceptions of COVID-19 around the world. *Journal of Risk Research*, 23(7–8):994–1006.

Fan, Y., Orhun, A. Y., and Turjeman, D. (2020). Heterogeneous actions, beliefs, constraints and risk tolerance during the COVID-19 pandemic. *NBER Working Paper*, w27211.

Freeman, A. L., Kerr, J., Recchia, G., Schneider, C., Lawrence, A. C., Finikarides, L., Luoni, G., Dryhurst, S., and Spiegelhalter, D. J. (2020). Communicating personalised risks from COVID-19: Guidelines from an empirical study. *medRxiv*.

Galesic, M. and Garcia-Retamero, R. (2010). Statistical numeracy for health: a cross-cultural comparison with probabilistic national samples. *Archives of Internal Medicine*, 170(5):462–468.

Gigerenzer, G. and Hoffrage, U. (1995). How to improve bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102(4):684–704.

Giglio, S., Maggiori, M., Stroebel, J., and Utkus, S. (2021). Five facts about beliefs and portfolios. *American Economic Review*, 111(5):1481–1522.

Heffetz, O. (2021). Are reference points merely lagged beliefs over probabilities? *Journal of Economic Behavior & Organization*, 181:252–269.

Heffetz, O. and Rabin, M. (2013). Conclusions regarding cross-group differences in happiness depend on difficulty of reaching respondents. *American Economic Review*, 103(7):3001–21.

Hoechle, D. (2007). Robust standard errors for panel regressions with cross-sectional dependence. *The Stata Journal*, 7(3):281–312.

Kennedy, R., Clifford, S., Burleigh, T., Jewell, R., and Waggoner, P. (2018). The shape of and solutions to the MTurk quality crisis. *Available at SSRN*, `https://papers.ssrn.com/abstract=3272468`.

Loewenstein, G., O'Donoghue, T., and Bhatia, S. (2015). Modeling the interplay between affect and deliberation. *Decision*, 2(2):55–81.

Manski, C. F. (2018). Survey measurement of probabilistic macroeconomic expectations: Progress and promise. *NBER Macroeconomics Annual*, 32(1):411–471.

Papageorge, N. W., Zahn, M. V., Belot, M., van den Broek-Altenburg, E., Choi, S., Jamison, J. C., Tripodi, E., et al. (2020). Socio-demographic factors associated with self-protecting behavior during the COVID-19 pandemic. *Institute of Labor Economics (IZA)*.

Shermohammed, M., Goren, A., Lanyado, A., Yesharim, R., Wolk, D. M., Doyle, J., Meyer, M. N., and Chabris, C. F. (2021). Informing patients that they are at high risk for serious complications of viral infection increases vaccination rates. *medRxiv*.

Slovic, P., Monahan, J., and MacGregor, D. G. (2000). Violence risk assessment and risk communication: The effects of using actual cases, providing instruction, and employing probability versus frequency formats. *Law and Human Behavior*, 24(3):271–296.

Weinstein, N. D., Kwitel, A., McCaul, K. D., Magnan, R. E., Gerrard, M., and Gibbons, F. X. (2007). Risk perceptions: Assessment and relationship to influenza vaccination. *Health Psychology*, 26(2):146–151.

Windschitl, P. D. (2002). Judging the accuracy of a likelihood judgment: The case of smoking risk. *Journal of Behavioral Decision Making*, 15(1):19–35.

Windschitl, P. D. and Wells, G. L. (1996). Measuring psychological uncertainty: Verbal versus numeric methods. *Journal of Experimental Psychology: Applied*, 2(4):343–364.

Wise, T., Zbozinek, T. D., Michelini, G., Hagan, C. C., et al. (2020). Changes in risk perception and protective behavior during the first week of the covid-19 pandemic in the United States. *PsyArXiv*.

Woloshin, S. and Schwartz, L. M. (2011). Communicating data about the benefits and harms of treatment: a randomized trial. *Annals of Internal Medicine*, 155(2):87–96.