

Public reactions to communication of uncertainty: How long-term benefits can outweigh short-term costs

Running header: Public reactions to communicating uncertainty

Author affiliation notes:

CHRISTINE STEDTNITZ: Fellow, London School of Economics and Political Science, London, UK.

PAULA SZEWACH: Data Science Researcher, Barcelona Supercomputing Center, Barcelona, Spain.

ROBERT JOHNS* : Professor of Politics, Department of Government, University of Essex, Wivenhoe Park, Colchester, UK.

Corresponding author note:

Robert Johns, Department of Government, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK; email: rajohn@essex.ac.uk.

Supplementary Material:

Supplementary Material may be found in the online version of this article:
<https://doi.org...>

Acknowledgements:

The authors thank the editor and three anonymous reviewers for helpful comments.

Data availability statement:

REPLICATION DATA AND DOCUMENTATION are available at
https://osf.io/ahs9y/?view_only=111bb792abc24a26b69006a9aa0edc4b .

Abstract

Uncertainty is a fact of political life but not a fact of political communication. Elites are prone to make confident predictions and downplay uncertainty about future outcomes, presumably fearing that the acknowledgement of uncertainty would undermine public confidence in their predictions and the evidence they are based on. But this calculation might both exaggerate the costs and downplay the potential benefits of reporting uncertainty. On costs, the evidence from previous studies is mixed; on benefits, previous research has neglected the possibility that, by acknowledging that outcomes may be worse than expected, those communicating uncertainty will dampen public reactions to the bad news. Here, based on a two-stage online survey experiment (N=2165) from December 2020 about COVID-19 vaccines, we find results suggesting that governments are well advised to communicate uncertainty. The costs at Stage 1 were low: reporting a confidence interval around the safety and effectiveness of a hypothetical COVID-19 vaccine did not undermine belief in the statistics or intentions to take the vaccine. And there were indeed benefits at Stage 2: when outcomes turned out to be worse than expected but within that confidence interval, confidence in the vaccine was partly insulated from negative effects.

When governments introduce policies based on uncertain scientific evidence, how open should they be about that uncertainty? Five decades after Donald Campbell pointed out that the difficulties in communicating uncertainty meant that policy reforms “are typically advocated as though they were certain to be successful” (1969, 409), Jakob Jensen and colleagues were still lamenting that “public communication of science is generally devoid of caveats, limitations, or other forms of...uncertainty” (2017, 40). There is a stark contrast between the growing calls in the scientific community for “radical transparency” (Nature 2020, Manski 2019) and the general instinct of political and media elites to conceal rather than to communicate uncertainty (Peters and Dunwoody 2016).

That instinct is hard to justify on normative or democratic grounds but it is understandable given the assumptions that politicians are likely to make about public opinion. First, they may suppose that the language of uncertainty – probabilities, margins of error, and so on – adds too much complexity to what may already be hard to communicate to the general public. Second, they may be concerned that admissions of uncertainty will erode public trust in both the information and the policies based on it (Post and Maier, 2016). “Why should we follow rules based on predictions that might be wrong?” is the mass reasoning that elites have grounds to fear.

There are two lines of counter-argument, however. One is that the public does not react uniformly negatively to the reporting of uncertainty. Reviews of the literature (van der Bles et al. 2019; Gustafson and Rice 2020) point to much more mixed and contingent findings, indicating that the *costs* of reporting uncertainty are lower than supposed and may in some cases be absent altogether. The other counterargument emphasizes the potential down-the-line *benefits* of reporting uncertainty. The very definition of that uncertainty is that outcomes may end up being worse than first envisaged. The more clearly this uncertainty was communicated to the public in the past, the less jolted they are likely to be by the bad news in the present, and the more faith they are likely to retain in subsequent communications and predictions. That logic, yet to be tested in the literature summarised above, is central to this article.

Our empirical purpose is therefore to estimate both the costs and the potential benefits of reporting uncertainty – and thereby to test whether elite reluctance to do so is misplaced. We do so via a two-stage survey experiment on a sample of 2,165 British adults in December 2020, and based around that country’s emerging COVID-19 vaccination programme. Given the urgent calls for transparency in communication about these vaccines (e.g. Mahase, 2020), this

is an obvious and useful test of both the costs and the benefits of reporting uncertainty. As the first vaccines emerged, politicians across the world had to persuade people to take them based on limited information about the safety and effectiveness of those vaccines. They might reasonably have feared that acknowledging any sort of uncertainty, especially around potentially lethal side effects, would aggravate public suspicions that the vaccines were untested or risky – with potentially disastrous consequences for vaccine roll-outs. On the other hand, those communicators were shooting at a moving target. If their initial estimates proved over-optimistic and they had failed to acknowledge that possibility, this risked further scepticism or distrust down the road.

Our treatments manipulate the information available about the effectiveness or safety of a hypothetical vaccine. The dependent variables are respondents' evaluations of that effectiveness and safety and their willingness to get the vaccine. (While reporting uncertainty may also indirectly shape other important outcomes such as trust in the communicators or in science more broadly, our concern here is more directly with trust in the information communicated and willingness to act on it.) At Stage 1, we investigate the cost side, i.e., whether communicating uncertainty around the effectiveness and safety projections undermines vaccine intentions – and, if so, whether that effect could be mitigated by explaining why such uncertainty was inevitable. At Stage 2, we investigate the benefits side by providing respondents with an update based on new vaccine data. The updated data are worse than the initial point estimates but just within the confidence intervals around them, with the idea that the reporting of uncertainty had prepared respondents for the negative news.

We report two encouraging results. First, reporting uncertainty around estimates of the effectiveness and safety of a COVID-19 vaccine does not significantly dent confidence in those initial reports or vaccine intentions. Second, people provided with uncertainty information react less negatively – in terms of perceptions of the vaccine as well as intentions to have it – than people who only saw the over-optimistic point estimate. Both findings provide some instrumental back-up to the normative arguments for communicating uncertainty to the public.

Public reactions to the communication of uncertainty

We draw a clear distinction here between the immediate costs of reporting uncertainty and the down-the-line benefits of doing so. There is now plenty of empirical research into the potential costs of communicating uncertainty in terms of negative public reactions, and our task is to

highlight where our study fits into and adds to that literature. By contrast, our hypothesis about the potential benefits of *having reported* uncertainty when outcomes prove worse than expected has not been empirically tested, and so there we lay out the theoretical arguments underpinning that hypothesis.

The avoidable costs

For some time now, researchers have assessed the impact of communicating uncertainty on public attitudes and behaviour. The field has been strongly empirically driven, especially since the rush of studies triggered by the pandemic (Ratcliff et al., 2022), and it can be hard to pick out consistent results amid a flurry of mixed and null findings. However, two recent reviews (van der Bles et al., 2019; Gustafson and Rice, 2020) highlight three key variables that can account for much of this inconsistency. The first of these is the *type of uncertainty* or, related, the reason for it. Typologies and terminology vary across studies but the same broad distinctions recur and prove relevant for public reactions. First, reactions are most negative when uncertainty is due to disagreement among experts (Gustafson and Rice 2020; Paek and Hove 2020). By contrast, there is often no cost of reporting the other main types of uncertainty: aleatory or scientific, where some fundamental randomness in the world means that the future cannot be predicted with any certainty; deficient, where the future might be predicted but there is as yet too little data; and technical or epistemic, where some limits on existing data – statistical assumptions, methodological choices, sampling bias or variance, measurement error and so on – mean that predictions come with a degree of uncertainty (Spiegelhalter 2017; van der Bles et al. 2019). These types can overlap, as in the case of the early COVID vaccines, where a shortage of data meant heavy reliance on the limited testing done so far. What they have in common is that they are the result of limitations in the data rather than expert disagreement about it. For the purposes of our study, the key points are these: that uncertainty about vaccines was primarily a case of technical uncertainty; and this is the type to which the public is *least* likely to react negatively (Gustafson and Rice 2020).

The second key variable is the *method of conveying uncertainty*. Amid a range of possibilities here, the key distinction is between numerical ranges and verbal descriptions, with the latter carrying a heavier cost in terms of public reactions. Examining the effect of reporting uncertainty in news stories about topics like global warming or immigration, van der Bles et al. (2020) find that using words to describe uncertainty reduces trust in both the statistics and their source. In contrast, numbers, reporting a numerical range rather than a point estimate have

no such negative effect. From a review of uncertainty communication, Dhimi and Mandel conclude that: “The benefits of precise numeric expressions of uncertainty, coupled with receivers’ preference for numeric information when it really matters, suggests that senders ought to embrace numeric precision over vague words if they wish to communicate uncertainty clearly” (2022, 514).

It is probably no coincidence that the reporting of technical uncertainty and of numerical ranges both cost less in terms of public reactions. For one thing, there is a clear connection between the two: numerical ranges are the calculable expression of the uncertainty generated by the limitations in the data available. The other point is more psychological. Both imply a degree of control over uncertainty. The public need not know the term ‘confidence interval’ to have a sense that those communicating a specific range of possibilities are more confident about an uncertain future than those simply asserting its uncertainty. The qualitative analyses of reactions to uncertainty reported by Markon et al. (2013) and Maxim and Mansier (2013) both highlight the role of a sense of communicator control in reassuring the public.

The third key factor that explains the inconsistency of research findings is the different *outcome variables*. Researchers have investigated the effect of reporting uncertainty on four types of public reactions: trust in the source, trust in science more generally, belief in the specific statistics or claims being reported, and intention to act based on them. There is more work on the first two and here the evidence is broadly optimistic. Plenty of studies find that reporting of uncertainty has no effect on trust in the messenger (e.g. van der Bles et al. 2020; Wiedemann et al. 2021; Janssen et al. 2021); some others find that it actually boosts such trust (e.g. Jensen 2008; Johnson and Slovic 1995). These broader dimensions of trust are important, but they are in a sense less immediate than the latter two outcomes. In many contexts, of which COVID-19 vaccines are a vivid example, we are most directly interested in whether uncertainty undermines belief in the claims made and affects behaviour based on those claims. That is the focus of our study.

Here there has been less research (see Gustafson and Rice 2020, Table 2). In their review, van der Bles et al. (2019, 26) explicitly highlight behaviour and compliance as an area needing more systematic research. What research there is paints a mixed picture. Morton et al. (2011) find people less willing to take environmental action if the probability of extreme weather conditions due to climate change is presented as a range rather than as a point estimate, whereas Joslyn & LeClerc (2012) find that transparency about uncertainty in weather forecasts increases

appropriate precautionary action (while deterring unnecessary action). Then there are null findings: Gustafson and Rice (2019) find no effect of uncertainty frames on either belief in scientific claims or behavioural intentions in the areas of climate change, genetically-modified foods and machinery hazards. Daoust and Bastien (2021) find no effect of graphically depicted confidence intervals around COVID-19 death projections on either trust in the statistics or support for preventive public health measures.

To recap: there is clearly potential to reduce the costs of reporting uncertainty by emphasising its technical inevitability and expressing it as a numerical range, but we still have only limited evidence of the effectiveness of this approach when it comes to behavioural compliance. Our study thus makes its first contribution by estimating the costs of reporting uncertainty in the crucial case of public beliefs about and willingness to take a COVID-19 vaccine. It goes beyond two similar experimental studies in this area. Petersen et al. (2021) vary the level of transparency in descriptions of a hypothetical vaccine (which is the standard approach given the obvious ethical concerns with manipulating details about a real vaccine). The authors find that transparently communicating even negative features of the COVID-19 vaccines (e.g. that it was less effective, or safe than a flu vaccine) was barely more damaging to vaccine willingness than was vaguely reassuring information (e.g. stating it was ‘sufficiently effective’, or had ‘acceptable’ side effects). However, since their manipulation was of transparency rather than uncertainty, these relatively upbeat results may not generalise to our case. There is a more explicit focus on uncertainty in a similar study by Kelp et al. (2022). However, their more downbeat conclusion – that a high-uncertainty condition erodes vaccine acceptance, especially among the initially reluctant – could reflect confounds between uncertainty and vagueness or positivity. While their low-uncertainty condition reported specific and extremely high effectiveness rates (‘100% effective’ for severe cases), the high-uncertainty condition reported no numbers and generally resembled the vagueness condition that Petersen et al. (2021) also found to be unpopular (e.g. stating that ‘most’ vaccinated individuals did not get infected). What is needed – and we provide – is a study in which only uncertainty is manipulated: that is, the same specific point estimates are reported but, in the treatment group, accompanied by measures of uncertainty.

The potential benefits

The second and more innovative contribution of this study is its exploration of the potential down-the-line benefits of reporting uncertainty. The idea here is that reporting not only shapes

the public's immediate assessments of messages and their messengers; it can also shape the way in which they react to what happens next. Our basic argument is that having acknowledged uncertainty at T_1 can serve as an insurance policy against criticism at T_2 if initial estimates turn out to have been too optimistic. This is an important point: after all, scientific forecasts and the uncertainty around them are about a future that *will* emerge – and often in the public glare.

The driver of our argument is that, when people react to incoming information, their expectations matter. That premise can be derived from evolutionary psychology and the distinction between disposition and surveillance systems (e.g. Armony and Ledoux 1997; Ranganath and Rainer 2003; Reisenzein et al. 2019). If incoming reports are within the bounds of our expectations, we can continue to rely on our predispositions. If they are unexpected and potentially threatening, the surveillance system generates anxiety, which in turn prompts cognitive processing and potentially belief change (MacKuen et al. 2010). It is also central to many models of Bayesian reasoning, where unexpected or surprising updates – information that falls a long way from priors – generate particular attention (Baldi and Itti 2010) and have a disproportionate impact on belief change and behaviour (e.g. Lorini and Castelfranchi 2007; Choi and Hui 2014). If the goal of a communicator or policymaker is to deter overreaction to a negative update, then making that update less surprising is a theoretically proven route to doing so.

This argument has been applied to communication research, notably by Lamberson and Soroka (2018) who develop a model of 'attentiveness to outlying news'. Their contribution is noteworthy here for two reasons. First, they define 'outlyingness' in terms of 'deviation from expectations' and emphasise the extent of this deviation as the key driver of public opinion – that is, of change in beliefs about the state of the world. Second, the choice of the term 'outlying' implies that expectations form a range or interval rather than a point estimate. Incoming information is outlying if (and to the extent that) it falls outside that range of expectations. This is neatly consistent with our argument that, if the result reported in a negative update had previously been included within a confidence interval, it is no longer 'outlying' and so is less likely to shift beliefs and undermine behavioural intentions.

We emphasise those two outcomes (belief in the statistics, and intentions to act based on them) because, as noted in the previous section on the costs of reporting uncertainty, they are our central concern in this study. But it is worth noting that we might expect the same logic to operate with other outcomes such as trust in the messenger or in scientific claims more

generally. Communicators who expect a certain outcome but admit that things might end up worse should face a (slightly) more forgiving public if and when that happens. By contrast, communicators having to admit that they were wrong face a reputational hit (Freitag et al. 2023).

Because it is hard to test with the ‘one-shot’ studies that have dominated in research on communicating uncertainty, evidence on this potential ‘insurance policy’ benefit is very scarce. What we do know, courtesy for example of a one-shot study by Kreps and Kriner (2020) using COVID-19 death tolls, is that people react badly to a treatment reporting that a given outcome is worse than was previously estimated. What we do not yet know is whether, had that previous estimate come with a confidence interval that encompassed that outcome, people would be at least somewhat mollified. This is best tested by the two-stage design used in our study, in which uncertainty is reported at the first stage and then the eventual data are reported at the second.

The current study

The rollout of COVID vaccinations provides an excellent opportunity to study both the costs and benefits of reporting uncertainty. This is an unusually high-profile case of a government policy requiring mass compliance but based on less than perfect information (Larson 2020; Motta 2021; Dudley et al. 2021). There is an especially pressing need to understand whether, and if so to what extent, being transparent about the uncertainties around the new vaccines impacts take-up (Motta 2021; Petersen et al. 2021). Recent literature on COVID-19 vaccine hesitancy suggests that patients appreciate being informed about the risks involved (Schwartz 2020). The DELVE Initiative (2020) recommended that ‘clear, transparent communication’ be used to address ‘rational doubts and to enable informed decision-making’, and not to hide the potential limitations of vaccines (see also van Bavel et al. 2020).

In practice, as lamented at the outset, such transparency is not common. While information about uncertainty around the vital statistics of vaccines is provided by manufacturers and thus available to those who sought it, it is not routinely included in government communications or media reports. Of course, there are various reasons for reporting only point estimates without confidence intervals. For present purposes, the question is whether reporting such uncertainty would undermine vaccine willingness in the short term – or help to sustain it in the medium term if the point estimates prove too optimistic.

Following Motta's (2021) finding that safety and effectiveness are two principal drivers of vaccine willingness, we make those the basis for our manipulations. In contrast to previous work, we disentangle the two, exposing respondents to either safety or effectiveness statistics and thus providing two parallel tests of our hypotheses. On both safety and effectiveness tracks, the point estimates on these two criteria are held constant; what is manipulated is whether they are accompanied by confidence intervals.

Before setting out our hypotheses, we should note three points about our confidence intervals approach to reporting uncertainty. All three derive from its relative precision (cf. the vaguer verbal approach of Kelp et al. (2022), for example). First, it is a direct application of the theoretical point about 'outlying' news. The point about locating negative outcomes within confidence intervals is to render them no longer outlying. Second, the tendency in the public towards negativity bias and risk aversion (Soroka, Fournier, and Nir 2019) means that there is likely to be particular focus on the pessimistic end of the confidence interval. If that pessimistic end is given a numerical expression, particularly one that is not a huge distance from the point estimate, that gives worried citizens a firmer foothold and might make uncertainty easier to cope with – thus reducing the costs of reporting it. The third point has more to do with cognitive capacity. The precise meaning of confidence intervals is relatively complicated, needing at least a basic understanding of probability distributions and inferential statistics. This might lead some to doubt whether they are a helpful means of conveying uncertainty. But there is a much easier means of understanding them – as the range between the worst and best case scenarios – which, while not correct, serves as a reasonable heuristic. Hence Daoust and Bastien (2021) make a convincing theoretical case for them helping people to grasp uncertainty – and, as noted above, a convincing empirical case that they do nothing to undermine faith in the reliability of the statistics.

One thing is to understand the method of presenting uncertainty; another is to understand why uncertainty is there in the first place. The acceptance and processing of a confidence interval may be readier among those who know – or learn – why, especially in the early days of clinical trials based sometimes on restricted samples, the rates of effectiveness and the risk of side-effects can be estimated only imprecisely. To test this conjecture, another innovation in our experiment is that there are two uncertainty treatments: the basic version, simply reporting confidence intervals, and the expanded version in which there is also an explanation of the reasons why an exact estimate makes less sense than a range. We refer to this as a conjecture rather than a formal hypothesis given that there is little research on *how* to explain the reasons

for uncertainty, and our expanded treatments are effectively therefore pilot tests of potential wordings. Nonetheless, our general expectation here is that the explained version elicits a more positive response from respondents.

As already noted, existing research points to heartening null findings when it comes to the effect of reporting uncertainty intervals on trust in information and its source. Ours is a new and acid test given both the unusual public salience of COVID-19 and the extension from trust in sources to planned behaviour – here, vaccination intentions. Nonetheless, that previous work gives grounds for optimistic hypotheses about the costs of reporting trust.

H1 (No costs Hypothesis): Providing a confidence interval around estimates of vaccine safety/effectiveness does not significantly weaken:

- a. willingness to have the vaccine
- b. perceptions of the vaccine's safety/effectiveness

The more novel hypothesis here concerns the potential benefits of reporting uncertainty. These are likeliest in situations where estimates have to be adjusted downward. Our hypothesis is that, if that adjustment goes below the initial estimate but remains within its confidence interval, the audience reaction to the bad news is dampened. As before, we test the same hypothesis in parallel for updates on safety and effectiveness and test it for both vaccine intentions and perceptions of that safety or effectiveness.

H2 (Benefits Hypothesis): When an initial estimate of safety/effectiveness is corrected downward, providing a confidence interval around that estimate increases:

- a. willingness to have the vaccine
- b. perceptions of the vaccine's safety/effectiveness

The psychological mechanism driving H2 is the negativity bias referred to earlier. If respondents presented with a range of outcomes anticipate a result towards the more pessimistic end of that range, they will have less Bayesian-style updating to do when the disappointing new data come in. We can test for that mechanism by asking about perceptions of the updated data and the extent to which they fall below what respondents had been led to expect. If our logic holds, then:

H3 (Mechanism Hypothesis): Respondents who had seen confidence intervals around estimates of safety/effectiveness will be less likely to say that the update revealed the vaccine as less safe/effective than expected.

Data

We tested our hypotheses via a two-stage experiment embedded in an online survey. The survey was fielded on 22 December 2020, when medical trials were ongoing and prominent in the media, risks were salient, and there were big question marks around the effectiveness and safety of the vaccines coming up for government approval. Hence, while the questions were about a hypothetical vaccine, the context made this a highly realistic exercise.

The sample was of 2,165 UK residents recruited via the Prolific platform (Peer et al. 2017). This sample size gives the experiment considerable power to detect the costs and benefits of reporting uncertainty. Like other Prolific samples, this one is more diverse than many convenience samples but is still unrepresentative of the UK adult population in various (related) ways. The complete sample characteristics are reported in section A of of the Supplementary Material accompanying this article. In line with the advice of Miratrix et al. (2018, 275), we do not weight the data for the experimental analysis. Typically, the whole sample was available for analysis because we rarely included a “don’t know” or other non-response option, wanting instead to encourage respondent effort and to avoid any biases from differential non-response across experimental conditions. The only exception is the mechanism questions which asked respondents about change between the manufacturers’ initial estimate and subsequent reports of safety or effectiveness. (This is why the Ns in Table 3, which is based on those questions, are very slightly smaller than the full sample bases reported in Table 1.)

Design

The survey questionnaire (reported in full in Supplementary Material section D) opened with some general questions about COVID-19, including general vaccination intentions, trust in key actors, concern about the pandemic and personal experience of the virus. Then came the two stages of the experiment, separated by a battery of socio-demographic items. These act as a buffer to make the information update a little more realistic, and the priming from Stage 1 a little lighter, than if the two stages had been adjacent.

[TABLE 1 ABOUT HERE]

The experimental design is set out in Table 1. The crucial treatment is uncertainty and has three levels: a simple point estimate not conveying any uncertainty ('point estimate'), a point estimate along with a range in which the true estimate is likely to fall ('confidence interval'), and a point estimate with a range *and* an explanation why the point estimate could not be exact ('CI and explanation'). In contrast to earlier studies (Petersen et al., 2021; Kelp et al., 2022), we disentangle the effect of exposure to safety and effectiveness statistics by allocating respondents to one of two tracks, showing respondents either safety or effectiveness statistics. In other words, we cross-manipulate the uncertainty treatments with the two different domains of uncertainty. There was also a pure control condition – respondents asked about initial willingness to take the vaccine without receiving any data about safety or effectiveness – but, since our focus is on the effect of reporting uncertainty, our comparisons are focused on the three conditions of that manipulation.

The experiment has two stages. Respondents are first prompted to suppose they are offered a COVID-19 vaccine. Next, they are shown some information from the manufacturers about either how effective (in terms of percentage reduction in the transmission rate) or how safe (in terms of the probability of serious side effects) the vaccine is estimated to be. From this point, respondents are randomly assigned to one of the three uncertainty conditions in Table 1. As recommended by Dhimi and Mandel (2022), the uncertainty information is delivered numerically but not graphically: we report confidence intervals in the treatment texts rather than on a graph (cf. Daoust and Bastien, 2021). The additions in the explained uncertainty treatment differ a little across the two tracks, reflecting the slightly different reasons for technical uncertainty in the two cases. In the effectiveness treatment, the addition focuses on random sampling variance and on biases in the test sample; the safety treatment includes parallel wording on sampling bias but also an acknowledgment that side effects could emerge too slowly to be captured in the trials. Given the potential difficulties in understanding these treatments – particularly the more elaborate 'explained uncertainty' condition – we fielded a second survey aimed specifically at testing comprehension. The reassuring findings of that survey, showing widespread understanding even of the more involved treatments, are in Supplementary Material section C.

To test the No Costs Hypothesis (H1) at this first stage, we ask respondents how likely they are to get this vaccine and how safe or effective (depending on which track they were on) they feel it is. There are of course numerous drivers of vaccination intentions – many of them unrelated to the threat assessments that we manipulated here (Friedman 2019) – and many respondents will have entered the experiment with strong motivations for accepting or rejecting the information we provided (Taber and Lodge 2006). This is not disruptive to our purpose of gauging whether reactions vary across the randomly assigned treatment conditions. Nevertheless, that comparison is made cleaner by trying to iron out any pre-experimental differences in vaccine-relevant predispositions across those conditions, and so our regressions include controls for relevant dimensions of trust – in the NHS, in vaccine manufacturers, and in Prime Minister Boris Johnson – as well as age, gender, COVID history and self-reported vulnerability.

[TABLE 2 ABOUT HERE]

The second stage (see Table 2) is based on a putative update from the manufacturer on the vaccine's record after a few months of use. On both safety and effectiveness tracks, the updated results are presented as calculations rather than estimations and involve results that are at the negative limit of the previously reported range. There is no additional uncertainty manipulation at this stage; our interest is in the effect of the Stage 1 treatment on reactions to the Stage 2 information. To test the Benefits Hypothesis (H2) we therefore repeat the questions about the perceived safety or effectiveness of the vaccine, vaccination intention and confidence in that decision. (These questions were not asked of those in the 'No update' condition, which is shown in Table 2 for completeness but is not used in the analysis here.)

It is worth making a couple of remarks about this approach. While a between-subjects design is more common for one-shot information treatment experiments (and is therefore our approach for analysing Stage 1), using repeated measures within the same experiment is well suited to testing the effect of an information update (e.g. Tappin et al. 2020; Dixon et al. 2022; Zhou and Shen 2022; see Clifford et al. 2021 for a more general endorsement of this method). It is especially apt in a context where the information is explicitly an update on those things – the effectiveness and safety and thus the attractiveness of the vaccine – that are asked about a

second time around. The potential biases with swiftly repeated measures, namely memory effects and consistency biases, are not problematic because, with new information provided, it is clear to respondents both why the question is being repeated and why belief change would be as reasonable as belief consistency. In any event, even at Stage 2, our primary analysis remains a between-subjects comparison: that is, not the total amount of belief change, but the differences in belief change across the uncertainty conditions. If there were any residual biases from a repeated-measures approach, there is no reason why they would affect the uncertainty and no-uncertainty conditions differently, and so our core comparison is robust.

Results

Since the key outcome variables are on numerical (0-6) scales, it is reasonable to test our hypotheses using OLS regression. Full models are provided in section B of the Supplementary Material. Figure 1 reports our test of H1. For the safety and effectiveness tracks in turn, it graphs the effect of receiving either of the uncertainty treatments relative to the no uncertainty (point estimate) condition in Stage 1, when the vaccine is first introduced. In each graph there are two dependent variables corresponding to the two parts of H1: a) vaccination intentions and b) acceptance of the statistics – that is, perceived safety or effectiveness.

[FIGURE 1 ABOUT HERE]

First, then, does adding a confidence interval to the safety or effectiveness estimates drive down vaccination intentions? It does not. Neither the stand-alone confidence interval nor the fuller explained uncertainty version carried any cost in terms of vaccination intentions. If anything, the results fall on the benefits side. This was true for both tracks, confirming our first ‘no costs’ hypothesis (H1a). Turning to perceptions of vaccine safety and effectiveness (H1b), our results are again largely null but with one exception: in the ‘safety’ track, the explained confidence interval treatment had a small but just about significant negative impact on perceptions of safety (of 0.16 points along the seven-point scale). One reason why the explanation may have dented safety perceptions is that it explicitly mentions the possibility of side effects later on. However, since the confidence intervals for the two uncertainty treatments themselves overlap, we should hesitate before concluding that that explanation in itself has an effect. Hence, we accept our

second ‘no costs’ hypothesis (H1b) for the effectiveness track and only marginally and conditionally reject it for the safety track.

Overall, then, communicating uncertainty seems to carry minimal cost in the short run and, crucially, does nothing to dent vaccine willingness. We now turn to the Stage 2 update to examine what happens if, once data on a new vaccine emerge, the level of effectiveness drops to the bottom or the probability of serious side effects climbs to the top of those respective confidence intervals. Does acknowledging uncertainty offer a benefit by buffering against bad news?

The first thing to check is whether overall the update had the expected negative impact on vaccine intentions and perceptions of safety and effectiveness. Across the sample, paired-samples t-tests recorded statistically significant declines in all three dependent variables, although vaccination intentions fell by less (0.14 points) than the estimates of safety (0.24) or effectiveness (0.27) which were more directly informed by the update. None of these declines is particularly sharp, however, which may already be a hint that the uncertainty treatments prepared respondents for the new information and thereby helped to sustain willingness to accept the vaccine.

[FIGURE 2 ABOUT HERE]

A more direct test is provided by Figure 2. It does the same as Figure 1, reporting the effect of having seen an uncertainty treatment on vaccine intentions and perceived safety and effectiveness – but this time the Stage 2 responses on those outcome variables. On the effectiveness track, there is clear support for H2a: those who have received either uncertainty treatment are significantly more likely to want the vaccine after the Stage 2 update than those who had seen a point estimate prediction that proved over-optimistic. With that same outcome variable of vaccine intentions, the results are similar but not quite significant on the safety track. On balance, though, there is encouraging support for the first part of our benefits hypothesis, H2a. There is rather little support for H2b, by contrast. Even though the update more directly concerned the safety and effectiveness variables, there is little sign – except in the borderline case of the explained confidence interval on the effectiveness track – that those receiving an uncertainty treatment maintained more positive evaluations of the vaccine after the update.

H3 is intended to test the mechanism through which reporting uncertainty at Stage 1 delivers benefits at Stage 2. While those benefits have now been revealed as conditional, there is enough support for H2a to make H3 worth testing. Table 3 presents two simple crosstabulations, each time pairing the uncertainty treatment with a post-update question asking respondents whether the manufacturer's update had revealed the vaccine as more safe/effective or less safe/effective than expected. Two points are clear from the tables. The first is the support in each case for H3. Those seeing an uncertainty range that encompassed the updated figures were disproportionately likely to say that the vaccine was as safe or effective 'as expected' rather than 'a little less' safe or effective. (These differences are the main drivers of the significant chi-squared tests for each crosstabulation.) Second, there is an explanation here for the generally more muted or null results on the safety track: that is, not only is the treatment effect in the crosstabulation less strong but the safety update generally left respondents more unmoved overall, with quite a few even in the point-estimate condition (36 per cent) reading the update 'as expected'. Perhaps 6 in 10,000 seemed like such a small proportion that it is similar to the 2 in 10,000 originally estimated.

[TABLE 3 ABOUT HERE]

Conclusion

There is a strong ethical and democratic case for transparency in the public communication of science – including reporting the uncertainty surrounding scientific forecasts. Alongside those arguments, here we make an instrumental case for communicating uncertainty. It has two planks. The first is that reporting uncertainty carries little cost in terms of trust in data and willingness to act on it – in this case, to have a COVID-19 vaccine. The second is that it may offer benefits down the line. In this case, having reported uncertainty around initial estimates helps to sustain vaccination intentions when those estimates have to be corrected downward. And our data isolates the cognitive mechanism driving this: the pessimistic lower reaches of those uncertainty intervals become part of people's expectations such that, when the update comes in, it is less likely to be an unpleasant surprise.

This study thus contributes to a small but growing literature adding empirical backing to the normative case for political communicators to present proposals and make claims in a less bombastic fashion. They should know that acknowledging uncertainty need not erode trust and legitimacy but may in fact bolster them. From a policy perspective, two other findings are worthy of comment. One is the similarity of results across our basic and explained uncertainty treatments. Not only can people handle uncertainty but they grasp its implications without needing elaboration. This is crucial for communicators given the time and space needed for that elaboration. Simple confidence intervals – even expressed verbally numerically (and hence more economically than via a graph) – prove an intuitive means of conveying the likely range of outcomes. All of this is consistent with studies showing that clarity and brevity often prove sufficient. Heavier informational treatments often have minimal impact on vaccine intentions, even if they leave audiences feeling more informed (Kerr et al. 2021; Loomba et al. 2021).

The second is the fact that the benefits of reporting uncertainty are more limited in the safety context. Overall, the data suggest that respondents simply do less cognitive processing of the treatments in the safety track – perhaps they are feeling rather than analysing risk (Slovic 2010). The upbeat point is that, even in the safety context, there are no costs of reporting uncertainty: any heightened sense of anxiety does not create an inability to cope with uncertainty. This is important since talking about safety is not really avoidable for policymakers. If side effects emerge, no democratic government can contain that information.

There are three limitations of our study worth noting here, each relating to a different facet of external validity. One concerns the nature of our sample and specifically the high levels of vaccine willingness on show. This is partly contextual: the British public was relatively vaccine-willing anyway and sleeves in December 2020 were being enthusiastically rolled up as the roll-out got under way. But it may also reflect the particular willingness in a relatively educated and liberal sample. It is conceivable that a more vaccine-hesitant sample would have balked at the uncertainty that these respondents took in their stride. That said, given that vaccine intentions were much less politicised in the UK (as in much of Europe) than in the US, with all mainstream parties strongly promoting the vaccination campaign (Klymak and Vanders 2022), sampling bias seems unlikely to have made a big difference to the effects recorded here.

Second, we could only examine vaccination intentions, not actual vaccination decisions. While this is a common limitation given the expense and difficulty of panel studies, it is true that the

ideal design would track respondents across the real time during which new information – whether about the safety and effectiveness of a vaccine, or about any politically salient outcome – emerges. In a survey experimental context like this, indeed, vaccination intentions could be seen more as an attitude or feeling about the vaccine than as an intention highly proximate to behaviour. However, this does not undermine our findings because we were interested in exactly those feelings – and how they were affected by the communication of uncertainty. Many other factors could subsequently intervene in the real world, pushing people towards or away from a vaccine, but that is a point about attitude-behaviour disjuncture rather than a criticism of survey experiments *per se*.

Third, the COVID-19 pandemic represents (thankfully) an unusual case in terms of the nature and salience of the predictions being communicated. Vaccine effectiveness and safety matters a great deal to governments and citizens alike. The statistics being reported have direct implications for behaviour – unlike, say, economic forecasts which might influence consumer behaviour at the margins but are probably less immediately relevant. What is unclear pending further research is the implications of this for generalizing our findings. It could be that the heightened anxiety in this case triggers *more* cognitive processing of the uncertainty information (Mackuen et al., 2010). In that case, our advice would be less clear-cut in situations where the incentive for grasping and coping with uncertainty is weaker. A rival reading is that this case represents a stiff test of the case for communicating uncertainty. If citizens can respond sensibly to this information even in such a febrile context, then we will recommend reporting it as a general rule.

References

Armony, Jorge L., and Joseph E. LeDoux. 1997. "How the Brain Processes Emotional Information." In *Annals of the New York Academy of Sciences*. Vol. 821.

<https://doi.org/10.1111/j.1749-6632.1997.tb48285.x>.

Baldi, Pierre, and Laurent Itti. 2010. "Of Bits and Wows: A Bayesian Theory of Surprise with Applications to Attention." *Neural Networks* 23: 649–66.

<https://doi.org/10.1016/j.neunet.2009.12.007>.

Bavel, Jay J. Van, Katherine Baicker, Paulo S. Boggio, Valerio Capraro, Aleksandra Cichocka, Mina Cikara, Molly J. Crockett, et al. 2020. "Using Social and Behavioural Science to Support COVID-19 Pandemic Response." *Nature Human Behaviour* 4 (May): 460–71. <https://doi.org/10.1038/s41562-020-0884-z>.

Bles, Anne Marthe Van Der, Sander Van Der Linden, Alexandra L. J. Freeman, James Mitchell, Ana B. Galvao, Lisa Zaval, and David J. Spiegelhalter. 2019. "Communicating Uncertainty about Facts, Numbers and Science." *Royal Society Open Science* 6.

<https://doi.org/10.1098/rsos.181870>.

Bles, Anne Marthe van der, Sander van der Linden, Alexandra L. J. Freeman, and David J. Spiegelhalter. 2020. "The Effects of Communicating Uncertainty on Public Trust in Facts and Numbers." *Proceedings of the National Academy of Sciences of the United States of America* 117: 7672–83. <https://doi.org/10.1073/pnas.1913678117>.

Campbell, Donald T. 1969. “Reforms as Experiments.” *American Psychologist* 24: 409–29.

Choi, Darwin, and Sam K. Hui. 2014. “The Role of Surprise: Understanding Overreaction and Underreaction to Unanticipated Events Using in-Play Soccer Betting Market.” *Journal of Economic Behavior and Organization* 107. <https://doi.org/10.1016/j.jebo.2014.02.009>.

Clifford, Scott, Geoffrey Sheagley, and Spencer Piston. 2021. “Increasing Precision Without Altering Treatment Effects: Repeated Measures Designs in Survey Experiments.” *American Political Science Review* 115: 1048–65. <https://doi.org/10.1017/S0003055421000241>.

Daoust, Jean François, and Frédérick Bastien. 2021. “Should We or Should We Not Include Confidence Intervals in COVID-19 Death Forecasting? Evidence from a Survey Experiment.” *Political Studies Review* 19: 302–10. <https://doi.org/10.1177/1478929920985686>.

The DELVE Initiative. 2020. “SARS-CoV-2 Vaccine Development and Implementation; Scenarios, Options, Key Decisions.” Royal Society DELVE Initiative. <https://rs-delve.github.io/reports/2020/10/01/covid19-vaccination-report.html>.

Dhami, Mandeep K., and David R. Mandel. 2022. “Communicating Uncertainty Using Words and Numbers.” *Trends in Cognitive Sciences* 26. <https://doi.org/10.1016/j.tics.2022.03.002>.

Dixson, Henry G. W., Aimée F. Komugabe-Dixson, Fabien Medvecky, Jovana Balanovic, Helene Thygesen, and Edith A. MacDonald. 2022. “Trust in Science and Scientists: Effects

of Social Attitudes and Motivations on Views Regarding Climate Change, Vaccines and Gene Drive Technology.” *Journal of Trust Research* 12: 179–203.

<https://doi.org/10.1080/21515581.2022.2155658>.

Dudley, Matthew Z., Roger Bernier, Janesse Brewer, and Daniel A. Salmon. 2021. “Walking the Tightrope: Reevaluating Science Communication in the Era of COVID-19 Vaccines.”

Vaccine 39 (September): 5453–55. <https://doi.org/10.1016/j.vaccine.2021.08.037>.

Freitag, Joshua, Madeline Gochee, Mitchell Ransden, Brendan Nyhan, Kristy Roschke, and Dan Gillmor. 2023. “The Corrections Dilemma: Media Retractions Increase Belief Accuracy but Decrease Trust.” *Journal of Experimental Political Science*, 1–12.

<https://doi.org/10.1017/xps.2023.4>.

Friedman, Jeffrey A. 2019. “Priorities for Preventive Action: Explaining Americans’ Divergent Reactions to 100 Public Risks.” *American Journal of Political Science* 63: 181–96.

<https://doi.org/10.7910/DVN/ZSJA25>.

Gustafson, Abel, and Ronald E. Rice. 2019. “The Effects of Uncertainty Frames in Three Science Communication Topics.” *Science Communication* 41: 679–706.

<https://doi.org/10.1177/1075547019870811>.

———. 2020. “A Review of the Effects of Uncertainty in Public Science Communication.” *Public Understanding of Science* 29 (August): 614–33.

<https://doi.org/10.1177/0963662520942122>.

Janssen, Inse, Friederike Hendriks, and Regina Jucks. 2021. "Face Masks Might Protect You from COVID-19: The Communication of Scientific Uncertainty by Scientists Versus Politicians in the Context of Policy in the Making." *Journal of Language and Social Psychology* 40 (October): 602–26. <https://doi.org/10.1177/0261927X211044512>.

Jensen, Jakob D. 2008. "Scientific Uncertainty in News Coverage of Cancer Research: Effects of Hedging on Scientists' and Journalists' Credibility." *Human Communication Research* 34 (July): 347–69. <https://doi.org/10.1111/j.1468-2958.2008.00324.x>.

Jensen, Jakob D., Manusheela Pokharel, Courtney L. Scherr, Andy J. King, Natasha Brown, and Christina Jones. 2017. "Communicating Uncertain Science to the Public: How Amount and Source of Uncertainty Impact Fatalism, Backlash, and Overload." *Risk Analysis* 37 (January): 40–51. <https://doi.org/10.1111/risa.12600>.

Johnson, Branden B., and Paul Slovic. 1995. "Presenting Uncertainty in Health Risk Assessment: Initial Studies of Its Effects on Risk Perception and Trust." *Risk Analysis* 15: 485–94. <https://doi.org/10.1111/j.1539-6924.1995.tb00341.x>.

Joslyn, Susan L., and Jared E. LeClerc. 2012. "Uncertainty Forecasts Improve Weather-Related Decisions and Attenuate the Effects of Forecast Error." *Journal of Experimental Psychology: Applied* 18 (March): 126–40. <https://doi.org/10.1037/a0025185>.

Kachanoff, Frank J., Yochanan E. Bigman, Kyra Kapsaskis, and Kurt Gray. 2021. "Measuring Realistic and Symbolic Threats of COVID-19 and Their Unique Impacts on Well-Being and Adherence to Public Health Behaviors." *Social Psychological and Personality Science* 12 (July): 603–16. <https://doi.org/10.1177/1948550620931634>.

Kelp, Nicole C., Jessica K. Witt, and Gayathri Sivakumar. 2021. "To Vaccinate or Not? The Role Played by Uncertainty Communication on Public Understanding and Behavior Regarding COVID-19." *Science Communication* 44: 223–39.
<https://doi.org/10.1177/10755470211063628>.

Kerr, John R., Alexandra L. J. Freeman, Theresa M. Marteau, and Sander van der Linden. 2021. "Effect of Information about COVID-19 Vaccine Effectiveness and Side Effects on Behavioural Intentions: Two Online Experiments." *Vaccines* 9: 1–22.
<https://doi.org/10.3390/vaccines9040379>.

Kreps, S E, and D L Kriner. 2020. "Model Uncertainty, Political Contestation, and Public Trust in Science: Evidence from the COVID-19 Pandemic." *Science Advances* 6: 1–12.
<https://www.science.org>.

Lamberson, P. J., and Stuart Soroka. 2018. "A Model of Attentiveness to Outlying News." *Journal of Communication* 68: 942–64. <https://doi.org/10.1093/joc/jqy040>.

Larson, Heidi J. 2020. *Stuck. How Vaccine Rumors Start and Why They Don't Go Away*. Oxford University Press.

Loomba, Sahil, Alexandre de Figueiredo, Simon J. Piatek, Kristen de Graaf, and Heidi J. Larson. 2021. "Measuring the Impact of COVID-19 Vaccine Misinformation on Vaccination Intent in the UK and USA." *Nature Human Behaviour* 5 (March): 337–48.
<https://doi.org/10.1038/s41562-021-01056-1>.

Lorini, Emiliano, and Cristiano Castelfranchi. 2007. "The Cognitive Structure of Surprise: Looking for Basic Principles." *Topoi* 26: 133–49. <https://doi.org/10.1007/s11245-006-9000-x>.

Mackuen, Michael, Jennifer Wolak, Luke Keele, and Gero E. Marcus. 2010. "Civic Engagements: Resolute Partisanship or Reflective Deliberation." *American Journal of Political Science* 54: 440–58.

Mahase, Elisabeth. 2020. "Covid-19: Vaccine Trials Need More Transparency to Enable Scrutiny and Earn Public Trust, Say Experts." *The BMJ* 371 (October): 1–3. <https://doi.org/10.1136/bmj.m4042>.

Manski, Charles F. 2019. "Communicating Uncertainty in Policy Analysis." *Proceedings of the National Academy of Sciences of the United States of America* 116: 7634–41. <https://doi.org/10.1073/pnas.1722389115>.

Markon, Marie Pierre L., Joshua Crowe, and Louise Lemyre. 2013. "Examining Uncertainties in Government Risk Communication: Citizens' Expectations." *Health, Risk and Society* 15 (June): 313–32. <https://doi.org/10.1080/13698575.2013.796344>.

Maxim, Laura, Pascale Mansier, and Natalia Grabar. 2013. "Public Reception of Scientific Uncertainty in the Endocrine Disrupter Controversy: The Case of Male Fertility." *Journal of Risk Research* 16: 677–95. <https://doi.org/10.1080/13669877.2012.726245>.

Miratrix, Luke W., Jasjeet S. Sekhon, Alexander G. Theodoridis, and Luis F. Campos. 2018. "Worth Weighting? How to Think about and Use Weights in Survey Experiments." *Political Analysis* 26(3): 275–91.

Morton, Thomas A., Anna Rabinovich, Dan Marshall, and Pamela Bretschneider. 2011. "The Future That May (or May Not) Come: How Framing Changes Responses to Uncertainty in Climate Change Communications." *Global Environmental Change* 21 (February): 103–9. <https://doi.org/10.1016/j.gloenvcha.2010.09.013>.

Motta, Matthew. 2021. "Can a COVID-19 Vaccine Live up to Americans' Expectations? A Conjoint Analysis of How Vaccine Characteristics Influence Vaccination Intentions." *Social Science and Medicine* 272 (March): 1–8. <https://doi.org/10.1016/j.socscimed.2020.113638>.

Nature. 2020. "Vaccine Confidence Needs Radical Transparency." *Nature* 586: 8–8. <https://www.nature.com/articles/d41586-020-02738-y>.

Paek, Hye Jin, and Thomas Hove. 2020. "Communicating Uncertainties During the COVID-19 Outbreak." *Health Communication* 35 (December): 1729–31. <https://doi.org/10.1080/10410236.2020.1838092>.

Peer, Eyal, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. "Beyond the Turk: Alternative Platforms for Crowdsourcing Behavioral Research." *Journal of Experimental Social Psychology* 70: 153–63. <https://doi.org/10.1016/j.jesp.2017.01.006>.

Peters, Hans Peter, and Sharon Dunwoody. 2016. “Scientific Uncertainty in Media Content: Introduction to This Special Issue.” *Public Understanding of Science* 25 (November): 893–908. <https://doi.org/10.1177/0963662516670765>.

Petersen, Michael Bang, Alexander Bor, Frederik Jørgensen, and Marie Fly Lindholt. 2021. “Transparent Communication about Negative Features of COVID-19 Vaccines Decreases Acceptance but Increases Trust.” *PNAS* 118: 1–8. <https://doi.org/10.1073/pnas.2024597118/-/DCSupplemental>.

Post, Senja, and Michaela Maier. 2016. “Stakeholders’ Rationales for Representing Uncertainties of Biotechnological Research.” *Public Understanding of Science* 25: 944–60. <https://doi.org/10.1177/0963662516645039>.

Ranganath, Charan, and Gregor Rainer. 2003. “Cognitive Neuroscience: Neural Mechanisms for Detecting and Remembering Novel Events.” *Nature Reviews Neuroscience* 4: 193–202. <https://doi.org/10.1038/nrn1052>.

Ratcliff, Chelsea L., Rebekah Wicke, and Blue Harvill. 2022. “Communicating Uncertainty to the Public During the COVID-19 Pandemic: A Scoping Review of the Literature.” *Annals of the International Communication Association* 64: 260–89. <https://doi.org/10.1080/23808985.2022.2085136>.

Reisenzein, Rainer, Gernot Horstmann, and Achim Schützwohl. 2019. “The Cognitive-Evolutionary Model of Surprise: A Review of the Evidence.” *Topics in Cognitive Science* 11: 50–74. <https://doi.org/10.1111/tops.12292>.

Schwartz, Jason. 2020. "Evaluating and Deploying Covid-19 Vaccines. The Importance of Transparency, Scientific Integrity, and Public Trust." *New England Journal of Medicine* 383 (November): 1920–31. <https://doi.org/10.1056/nejmoa2022483>.

Slovic, Paul. 2010. *The Feeling of Risk: New Perspectives on Risk Perception*. Routledge.

Soroka, Stuart, Patrick Fournier, and Lilach Nir. 2019. "Cross-National Evidence of a Negativity Bias in Psychophysiological Reactions to News." *Proceedings of the National Academy of Sciences of the United States of America* 116 (September): 18888–92. <https://doi.org/10.1073/pnas.1908369116>.

Spiegelhalter, David. 2017. "Risk and Uncertainty Communication." *Annual Review of Statistics and Its Application* 4: 31–60. <https://doi.org/10.1146/annurev-statistics-010814-020148>.

Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50: 755–69. <https://doi.org/10.1111/j.1540-5907.2006.00214.x>.

Tappin, Ben M., Gordon Pennycook, and David G. Rand. 2020. "Bayesian or Biased? Analytic Thinking and Political Belief Updating." *Cognition* 204: 1–12. <https://doi.org/10.1016/j.cognition.2020.104375>.

Wiedemann, Peter, Franziska U. Boerner, and Frederik Freudenstein. 2021. “Effects of Communicating Uncertainty Descriptions in Hazard Identification, Risk Characterization, and Risk Protection.” *PLoS ONE* 16 (July): 1–18.

<https://doi.org/10.1371/journal.pone.0253762>.

Zhou, Yanmengqian, and Lijiang Shen. 2022. “Confirmation Bias and the Persistence of Misinformation on Climate Change.” *Communication Research* 49: 500–523.

<https://doi.org/10.1177/00936502211028049>.

Figure captions

Figure 1. Effect of uncertainty treatments (relative to point estimate condition) on Stage 1 vaccination intentions and perceived safety/effectiveness (n(safety track) = 1320; n(effectiveness track) = 1306)

Figure 2. Effects of uncertainty treatments (relative to point estimate condition) on Stage 2 vaccination intentions and perceived safety/effectiveness (n(safety track = 817; n(effectiveness track) = 810)

Tables

Table 1: Treatment groups at Stage 1

	Safety track	Effectiveness track
Point estimate	Suppose you were offered the COVID-19 vaccine. Our tests show that ... 2 in 10,000 will develop more serious side effects. The most common serious side effect is Bell's Palsy, a freezing of muscles on one side of the face. (N=273)	Suppose you were offered the COVID-19 vaccine. Our best estimate of the effectiveness of the vaccine is 90% effective. What that means is that our vaccine reduces transmissions of the virus by 90%. (N=271)
Confidence interval	To be more exact, we estimate that between 1 and 6 out of 10,000 people will experience serious side effects. (N=272)	To be more exact, our tests show that the true effectiveness level is very likely to fall between 83% and 97%. (N=269)
Confidence interval and explanation	We cannot be exact in these circumstances for two reasons. One reason is that we have only been able to follow our volunteers for a few weeks. It is possible that they will develop side effects later on. Another reason ... is that those tested were not a perfect reflection of the UK population. (N=272)	We cannot be exact in these circumstances for two reasons. One reason is that... any results based on a sample of people come with a margin of error.... Another reason ... is that those tested were not a perfect reflection of the UK population. (N=270)
Control	Suppose you were offered the COVID-19 vaccine. (N=538)	

Table 2: Treatment groups at Stage 2

	Safety track	Effectiveness track
No update	-	-
Negative update	Suppose that the vaccine has been in use for a few months and the manufacturers update their results. Based on the data we have now collected, we calculate that the vaccine causes serious side effects ... in 6 out of 10,000 cases.	Suppose that the vaccine has been in use for a few months and the manufacturers update their results. Based on the data we have now collected, we calculate that the vaccine is 83% effective.

Table 3. Assessments of manufacturers' updates of safety (upper panel) and effectiveness (lower panel), by uncertainty condition. Cells report row percentages.

<i>Safety Track</i>						
Treatment condition	A lot less %	A little less %	About as expected %	A little more %	A lot more %	<i>N</i>
Point estimate (%)	7	45	36	9	4	256
Confidence interval (%)	5	35	46	10	3	252
CI plus explanation (%)	6	32	48	11	3	256
<i>Effectiveness Track</i>						
Treatment condition	A lot less %	A little less %	About as expected %	A little more %	A lot more %	<i>N</i>
Point estimate (%)	4	80	11	5	1	256
Confidence interval (%)	2	59	30	9	0	259
CI plus explanation (%)	0	57	32	9	2	258