



HAL
open science

Information delivered by a chatbot has a positive impact on COVID-19 vaccines attitudes and intentions.

Sacha Altay, Anne-Sophie Hacquin, Coralie Chevallier, Hugo Mercier

► To cite this version:

Sacha Altay, Anne-Sophie Hacquin, Coralie Chevallier, Hugo Mercier. Information delivered by a chatbot has a positive impact on COVID-19 vaccines attitudes and intentions.. *Journal of Experimental Psychology: Applied*, 2021, 10.1037/xap0000400 . ijn_03451205

HAL Id: ijn_03451205

https://hal.science/ijn_03451205v1

Submitted on 26 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Information Delivered by a Chatbot Has a Positive Impact on COVID-19 Vaccines Attitudes and Intentions

Sacha Altay¹, Anne-Sophie Hacquin¹, Coralie Chevallier² †
& Hugo Mercier¹ †

† equal contribution

¹ Institut Jean Nicod, Département d'études cognitives, Ecole normale supérieure, Université PSL, EHES, CNRS, 75005 Paris, France

² LNC², Département d'études cognitives, Ecole normale supérieure, Université PSL, INSERM, 75005 Paris, France.

ABSTRACT

The COVID-19 vaccines will not end the pandemic if they stay in freezers. In many countries, such as France, COVID-19 vaccines hesitancy is high. It is crucial that governments make it as easy as possible for people who want to be vaccinated to do so, but also that they devise communication strategies to address the concerns of vaccine hesitant individuals. We introduce and test on 701 French participants a novel messaging strategy: a chatbot that answers people's questions about COVID-19 vaccines. We find that interacting with this chatbot for a few minutes significantly increases people's intentions to get vaccinated ($\beta = 0.12$) and has a positive impact on their attitudes towards COVID-19 vaccination ($\beta = 0.23$). Our results suggest that a properly scripted and regularly updated chatbot could offer a powerful resource to help fight hesitancy towards COVID-19 vaccines.

Data, scripts, ESM, pre-registration, and materials: <https://osf.io/8q3b2/>

Keywords: COVID-19; Vaccination; Chatbot; COVID-19 vaccines; Vaccine refusal; Attitude change.

Public Significance Statement: Interacting a few minutes with a chatbot answering the most common questions about COVID-19 vaccines increased people's intention to get vaccinated and had a positive impact on their attitudes towards the vaccines. Chatbots could be a powerful resource to fight COVID-19 vaccines hesitancy.

INTRODUCTION

Most countries face the issue of vaccine hesitancy, with sizeable fractions, or sometimes the majority, of the public opposing some vaccines (de Figueiredo et al., 2020). The problem is particularly acute in the case of COVID-19 vaccination: first, a high uptake of COVID-19 vaccines is necessary to reach and sustain herd immunity; second, and to the best of our knowledge, no country is currently planning on making COVID-19 vaccination mandatory, making public approval essential. Unfortunately, hesitancy towards COVID-19 vaccines is high in many countries (for an international meta-analysis see: Robinson et al., 2020; for France see: Hacquin et al., 2020; Ward et al., 2020). It is therefore crucial that health authorities make it as easy as possible for people who want to be vaccinated to do so, but also that they devise communication strategies to reassure vaccine hesitant individuals. After having briefly reviewed related work, we introduce a novel messaging strategy: the use of a chatbot that answers people's questions about COVID-19 vaccines.

Using communication to increase vaccine uptake has proven difficult. Systematic reviews suggest that communication to the public often has a modest effect or no effect at all on attitudes towards vaccination, vaccination intentions, or vaccine uptake (Brewer et al., 2017; Community Preventive Services Task Force, 2015; Dubé et al., 2015; Kaufman et al., 2018; Sadaf et al., 2013). Several studies even reported backfire effects, with participants who were

initially the most opposed to vaccination becoming even more hesitant after the intervention (Betsch & Sachse, 2013; Nyhan et al., 2014; Nyhan & Reifler, 2015; although backfire effects remain exceptional as we will see below).

Most messaging efforts related to COVID-19 have borne on behavior such as handwashing, social distancing, and mask wearing. The effects of these information campaigns have been mixed, with studies revealing fleeting and hard to replicate effects (Barari et al., 2020; Bilancini et al., 2020; Capraro & Barcelo, 2020; Favero & Pedersen, 2020; Hacquin, Mercier, et al., 2020; Jordan et al., 2020). Likewise, studies that have attempted to boost COVID-19 vaccination intentions have had little success. One study found that messages emphasizing the risks of the virus, or the safety of vaccination, had no effect on vaccination intentions (Duquette, 2020). Another study found that a message providing people with information about the coverage needed to reach herd immunity decreased the time they wanted to wait before being vaccinated, but the effect was small, and did not replicate in another condition that contained the same message in addition to another message (Trueblood et al., 2020).

These results show that, as in many other domains (Mercier, 2020), changing people's minds at scale is a difficult endeavor. A major obstacle for communication campaigns is their inability to address most counter-arguments. When people encounter a message that aims at changing their minds, they typically generate counter-arguments (e.g. Greenwald, 1968). If they do not have an interlocutor who can address these counter-arguments (e.g. if they read a leaflet), they are less likely to change their minds. This likely explains why small-group discussion, in which counter-arguments can be addressed in the back and forth of discussion, is vastly more effective at changing people's minds than the simple presentation of arguments (even for logical arguments, see Trouche et al., 2014; Claidière et al., 2017; more generally,

on the effectiveness of small-group discussion to change minds, see Laughlin, 2011; Mercier, 2016; Mercier & Sperber, 2017). In line with this, direct communication with trustworthy professionals appears to be an efficient lever to increase vaccination acceptance. In an intervention involving vaccination experts engaging in a Q&A with an audience about the H1N1 vaccine, researchers found that, after having discussed with the experts on vaccination, participants were more willing to vaccinate (Chanel et al., 2011). More broadly, discussion with politicians (Minozzi et al., 2015), canvassers (Broockman & Kalla, 2016), or scientists (Altay & Lakhli, 2020; Goldberg et al., 2019) can lead to significant and durable changes of mind (Broockman & Kalla, 2016), which tend to be larger than those observed with standard messaging techniques (Chanel et al., 2011; Minozzi et al., 2015). The interactivity that group discussions and Q&A sessions offer is known to improve learning and comprehension, as well as motivation to learn (Freeman et al., 2014; Johnson et al., 2000; King, 1990; Prince, 2004; Shi et al., 2020).

The interactivity that small-group discussion provides is, however, difficult to scale up. A potential solution is to gather the most common counter-arguments and to offer rebuttals to each of them. A list of counter-arguments, which can be phrased explicitly as counter-arguments or as questions, can then be provided to people, along with the rebuttals. Since not every rebuttal is relevant to everyone, chatbots can work as an interesting alternative to long-texts presenting every possible argument. When interacting with a chatbot, people select the questions (or counter-arguments) that are most relevant to them and read the corresponding answers, which can then raise further questions and answers. Tentative evidence suggest that chatbots and automated computer-based conversational agents can be useful to change people's mind (Andrews et al., 2008; Rosenfeld & Kraus, 2016), and that chatbots asking users what they are concerned about increased chatbots' efficacy by providing users with

more relevant counterarguments (Chalaguine et al., 2019). In the lines below we will detail the experimental protocol of the first study to systematically test the effectiveness of chatbots in a large sample (Altay et al., 2020). In one condition, participants were provided with the most common counter-arguments against Genetically Modified Organisms (GMOs) along with their rebuttals, presented by a chatbot. In two control conditions, participants were either presented with a standard pro-GMOs message citing the scientific consensus on their safety, or with a brief description of GMOs. Participants' attitudes towards GMOs were measured before and after the treatment. When participants had access to the chatbot, their attitudes towards GMOs became significantly more positive than in the control conditions, with a large effect size ($\beta = 0.66$). Finally, in a last condition, participants were presented with a non-interactive version of the chatbot. The formatting and the interface were the same as the chatbot, but participants scrolled through the arguments instead of clicking on them—which makes it easy to find relevant information. This condition had an even larger effect ($\beta = 0.85$), probably because it had led participants to spend more time on the task.

Here, we test a chatbot on COVID-19 vaccination hesitancy by addressing the most common questions about COVID-19 vaccines. We identified the most common questions about COVID-19 vaccines by relying on a survey conducted on a representative sample of the French population documenting the reasons why people were willing, or not, to take a COVID-19 vaccine (Hacquín, Altay, et al., 2020). We also relied on press articles refuting common myths about the COVID-19 vaccines, and resources from health institutions. Answers to these common questions were drafted based on a wide variety of publicly available information and checked by several experts on vaccination. Overall, the questions and answers formed a long text of 9021 words.

Participants were randomly assigned to a Chatbot condition, in which they had the opportunity to interact with the chatbot for as long as they wanted, or to a Control Condition, in which they read a brief text (93 words) describing the way vaccines work. Note that our design is not meant to compare the efficacy of an interactive Chatbot compared to a non-interactive Chatbot or a long text (see Altay et al. 2020 for such design). Instead, the present design is primarily meant to test the efficacy of a Chatbot to inform people about COVID-19 vaccines. The Control Condition allows us to control for potential demand biases. Between one and two weeks after the experiment, we surveyed the participants again to measure whether the effect of the chatbot would last in time. We will refer to the first experiment as Wave 1 and the follow-up as Wave 2. All our hypotheses, sample size, and analysis plan were preregistered (<https://osf.io/8q3b2/>).

METHOD

Pre-registered hypotheses

Our first two hypotheses were that participants' attitudes towards the COVID-19 vaccines (H₁) and their intention to get vaccinated (H₂) would shift more positively compared to participants in the Control condition. If these shifts occurred in response to the information provided in the chatbot, rather than as a result of a task demand, we expected that attitude shifts (H₃) and intention shifts (H₄) would be modulated by the time participants spent on interacting with the chatbot. The longer participants interact with the chatbot, the more information in favor of vaccination they will be exposed to, which should lead to larger attitude and intention change in favor of vaccination. Several studies have found backfire effects among participants most opposed to vaccination. However, most of the empirical literature fails to identify backfire effects (see, e.g., Guess & Coppock, 2018; Swire-Thompson et al., 2020; Wood & Porter, 2019). We therefore hypothesized that our main

effects (H_1 and H_2) would be observed in all participants, including in the tercile most opposed to vaccination (H_5).

Participants

Based on an *a priori* power analysis (two tailed, power = 95%, $\alpha = 5\%$, $d = 0.2$; see the pre-registration on OSF) we recruited 701 French participants between the 23rd and the 28th of December 2020 on the crowdsourcing platform Crowdpanel. Participants were paid 2€ to spend 15 minutes on the survey. We excluded 42 participants who said that they had not been able to access the chatbot, and 16 participants who had spent less than 20 seconds on the Chatbot (a pre-registered exclusion criterion), leaving 643 participants (291 women, $M_{\text{age}} = 38.58$, $SD_{\text{age}} = 12.40$). A week later, between the 5th and the 12th of January 2021, participants who had taken part in the first wave were contacted to answer more questions, and 614 answered (attrition rate = 12.5%). This time participants were paid 0.27€ to spend two minutes on the survey.

Experimental procedure

Participants in both conditions provided informed consent form and then answered a baseline questionnaire. Participants were then randomized to the Control or Chatbot condition. Finally, participants in both conditions completed an endline questionnaire.

Materials

Baseline questionnaire. Participants first answered five questions to measure their attitudes towards the COVID-19 vaccines using a seven-point Likert scale (“In total disagreement”, “Disagree”, “Somewhat disagree”, “Neither strongly agree nor strongly disagree”, “Somewhat agree”, “Agree”, “Totally agree”): “I think vaccines against COVID-19 are

safe”, “I think vaccines against COVID-19 are effective”, “I think we know enough about the COVID-19 vaccines.”, “I think we can trust the people who produce the COVID-19 vaccines.”, “I think it is important to get vaccinated against COVID-19”. These five variables are treated as a single composite variable, “the COVID-19 vaccines attitude” variable, in all analyses. This composite measure of COVID-19 vaccines attitude had a good internal consistency ($\alpha_{\text{wave 1}} = 0.89$; $\alpha_{\text{wave 2}} = 0.92$). Next, participants’ intention to take a COVID-19 vaccine was queried with the following question: “Do you personally wish to be vaccinated against COVID-19?”, on a three points-Likert scale (“Yes, as soon as the vaccine is available for me”, “Yes, but I will wait some time before getting vaccinated”, “No, I will not get vaccinated”). Participants were then asked the extent to which they trusted two types of sources regarding vaccination: “How much do you trust medical and health advice from medical workers, such as doctors and nurses...?”, “To what extent do you trust the medical advice of alternative medicine (homeopathy, naturopathy, energetic medicine, etc.)?”, on a five points-Likert scale (“No trust at all”, “Somewhat not trusted”, “Neutral”, “Somewhat trusted”, “Totally trusted”). The two trust questions will be combined in a single composite variable (trust in medicine minus trust in alternative medicine), that we refer to as the “trust in medicine” variable. Finally, participants were asked the following question to measure their information seeking behavior: “How often do you look for information on COVID-19 or the COVID-19 vaccine?” on a five points-Likert scale (“Never”, “Less than once a week”, “Several times a week”, “Daily”, “Several times a day”).

Treatment phase. Participants were randomly assigned to the Control Condition or to the Chatbot Condition by a pseudo-randomizer on the survey platform “Qualtrics” (i.e. a randomizer that ensures an equal number of participants is attributed to each condition).

Participants were told that they were paid to spend approximately ten minutes to interact with

the chatbot, but that they were free to spend as much time as they wanted. Time spent interacting with the chatbot was measured by Qualtrics.

The description of the COVID-19 vaccines used in the Control Condition was taken from the French government website and read as follows:

When we get sick, our immune system defends itself by making antibodies. They are designed to neutralize and help eliminate the virus that causes the disease.

Vaccination is based on the following process: it introduces into our body an inactivated virus, part of the virus or even a messenger RNA. Our immune system produces antibodies in response to this injection. Thus, the vaccine allows our immune system to specifically recognize the infectious agent if it enters our body. It will then be detected, neutralized and eliminated before it can make us sick.

To develop the responses to the most common questions about COVID-19 vaccines presented in the chatbot, we relied on a wide variety of publicly available information (primary scientific literature, governmental websites, etc.). The text was checked by several experts on vaccination, and was 9021 words long.

The questions and responses were used to build the chatbot. Participants were exposed to the most common questions that we gathered about the COVID-19 vaccines, as well as the responses to these questions. Participants had to select (by clicking on them) the questions about the COVID-19 vaccines they wanted to ask, and they were provided with the responses to their questions.

The chatbot was organized as follows. Participants were first asked whether they had any questions about the COVID-19 vaccines, and were given a choice of six questions to select from: "Are COVID-19 vaccines safe?," "Are COVID-19 vaccines effective?," "Do we know enough about the COVID-19 vaccines?," "Can we trust the people who produce it?," and "Do I need to be vaccinated?". Participants were able to select, at any stage, an option "Why should I trust you?," that informed them of who we are, who funded us, and what our goals are (all the materials are available on the Open Science Framework (OSF) at <https://osf.io/8q3b2/>).

Every time participants selected a question, the chatbot offered an answer. Participants could choose between several sub-questions that the initial answer might not have addressed. In total the chatbot offered 51 questions and answers about the COVID-19 vaccines. The chatbot did not allow participants to write open-ended questions, participants only had the option of choosing among our fixed set of questions, which were each coupled with a predefined answer. The responses were displayed in separate discussion bubbles (see Figure 1).

The screenshot shows a chatbot interface. At the top left, a small robot icon is next to the text: "Bonjour, je suis un petit robot conversationnel." Below this, a question bubble on the right asks: "Avez-vous des questions sur le vaccin contre la Covid-19 ?". A timestamp "04/01/2021 10:20" is visible. A large black bubble on the right contains the question: "Est-ce qu'on a assez de recul sur les vaccins contre la Covid-19 ?". Another timestamp "04/01/2021 10:20" is on the right. Below this, a robot icon is next to a response: "Par rapport aux vaccins précédents, la mise en circulation de certains vaccins contre la Covid-19 est très rapide. Nous devons cette rapidité à la mobilisation de centaines d'équipes de recherche et de volontaires du monde entier." This is followed by another response: "Cependant, tous les vaccins, y compris les vaccins contre la Covid-19, passent par les mêmes procédures avant d'être distribués." Then: "Les essais ayant commencé il y a plusieurs mois, nous avons désormais un bon recul. En effet, les effets indésirables des vaccins surviennent quasiment toujours au cours du premier mois suivant l'administration du vaccin." At the bottom, three blue bubbles offer choices: "Je veux en savoir plus sur le recul que l'on a sur les vaccins contre la COVID-19.", "Est-ce que les vaccins sont développés de manière trop précipitée ?", and "Connaît-on assez cette maladie pour élaborer un vaccin sûr ?". At the very bottom, a grey bar contains the text: "Veuillez choisir parmi les choix ci-dessus".

Figure 1. The beginning of a conversation with the chatbot. The left-justified dialogue bubbles correspond to chatbot's responses. The right-justified black bubble corresponds to the first question asked to the chatbot. The left-justified blue bubbles at the bottom of the screenshot are questions the participant can choose from at this stage of the interaction. Translation from top to bottom: 1- Hello, I'm a little conversational robot. 2- Do you have questions about COVID-19 vaccines? 3- Do we know enough about the COVID-19 vaccines? 4- Compared to previous vaccines, the release of some Covid-19 vaccines is very rapid. We owe this speed to the mobilization of hundreds of research teams and volunteers from all over

the world. 5- However, all vaccines, including COVID-19 vaccines, go through the same procedures before being distributed. 6- Since the trials started several months ago, we now have a lot of information about the COVID-19 vaccines. Indeed, adverse vaccine reactions almost always occur within the first month after vaccine administration.

Responses contained hyperlinks to scientific articles, reports from scientific agencies, media articles, and Wikipedia. At any time, users had the option of coming back to the first four basic questions of the main menu. In addition to the interactive part of the chatbot, participants could display all the questions and answers on the page at once, and scroll through them instead of clicking on them.

Endline questionnaire. Once participants had read the text in the Control condition, or once they had finished interacting with the chatbot, they answered the same questions as those presented in the baseline questionnaire regarding their COVID-19 vaccines attitudes and vaccination intentions. Participants then answered the following question: “Imagine you are talking to someone telling you that the COVID-19 vaccines are not safe and effective, and that we cannot trust it. What would you tell them?” (free text entry)¹. Finally, participants provided basic demographic information (age, gender, education, trust in government). Trust in the government was measured by the following question: "In general, are you satisfied with the Government's handling of the Coronavirus crisis?", on a 4-item Likert scale ranging from “Not at all satisfied” to “Very satisfied”. Interpersonal trust was measured by the following question: “Generally speaking, would you say that most people can be trusted or that you can’t be too careful in dealing with people?”. In addition, participants in the Chatbot

¹ We initially planned to analyze participants’ responses to this question with the following research question: “RQ6 will investigate the arguments in favor of the Covid-19 vaccine given by participants in the Chatbot Condition and in the Control Condition. This investigation will be exploratory.” However, we have not found a good way of rigorously analyzing these responses yet.

Condition were asked whether they had been able to access the chatbot, whether the Chatbot was intuitive, pleasant, frustrating, whether the information provided in the chatbot were too simple or too complicated, and whether they had unanswered questions that they wish the chatbot had addressed (free text entry).

Wave two questionnaire. Between one and two weeks after the experiment, participants were contacted again, and asked to the same questions as in Wave 1 to measure their attitudes towards COVID-19 vaccines and their intention to take a COVID-19 vaccine. In addition, participants were asked how many people they had tried to convince of their opinion on COVID-19 vaccines and whether they had used the information presented during Wave 1 for that purpose. Finally, participants were asked whether they had interacted a chatbot during Wave 1. Those who answered “No” were presented with a link to the chatbot; participants who answered “Yes” were also presented with the link and asked whether they intended to share it.

All the materials (including the full text of the chatbot) can be found on OSF at <https://osf.io/8q3b2/>. The Chatbot was displayed on a custom-made website created by [La Fabrique à Chatbots](#).

Methods for statistical analyses

All analyses were done with R (v.3.6.1; Team, 2017), using R Studio (v.1.1.419; Team, 2015). All statistical tests are two-sided. We refer to “statistically significant” as the p-value being lower than an alpha of 0.05. We controlled for multiple comparisons applying the Benjamini-Hochberg method.

All the statistical analyses reported below are regressions. When comparing conditions, we controlled for participants' initial attitudes by adding them as a predictor in the model. Attitude change corresponds to participants' attitudes after the treatment minus participants' initial attitudes (a positive score corresponds to more positive attitudes after the treatment). Intention change corresponds to participants' intentions after the treatment minus participants' initial intentions (a positive score corresponds to more positive intentions after the treatment). Attitudes and intentions before the treatment, together with time spent on the chatbot, were mean centered in order to facilitate the interpretation of the intercept. More details about the statistical analyses are available on OSF.

Results

Descriptive results

	Pre-treatment <i>(Wave 1)</i>	Post-treatment <i>(Wave 1)</i>	Post-treatment <i>(Wave 2)*</i>
Control Condition	3.81 (1.41)	3.93 (1.45)	4.15 (1.43)
Chatbot Condition	3.82 (1.28)	4.26 (1.35)	4.27 (1.38)

Table 1. Mean and standard deviation (in parentheses) of participants attitudes towards COVID-19 vaccines in the Control Condition and in the Chatbot Condition, pre- and post-treatment, on a scale of 1 (negative attitudes) to 7 (positive attitudes). * = Due to a technical issue, participants in Wave 2 were matched based on their answers to the question “did you interact with a chatbot during the first survey?” Results of Wave 2 are thus less reliable than those of Wave 1.

		No, I will not get vaccinated	Yes, but I will wait some time before getting vaccinated	Yes, as soon as the vaccine is available for me
Control Condition	<i>Pre-treatment (Wave 1)</i>	110 (36%)	133 (44%)	62 (20%)
	<i>Post-treatment (Wave 1)</i>	107 (35%)	134 (44%)	64 (21%)
	<i>Post-treatment (Wave 2)*</i>	93 (30%)	135 (43%)	87 (28%)
Chatbot Condition	<i>Pre-treatment (Wave 1)</i>	123 (36%)	156 (46%)	59 (17%)
	<i>Post-treatment (Wave 1)</i>	99 (29%)	168 (50%)	71 (21%)
	<i>Post-treatment (Wave 2)*</i>	85 (29%)	131 (44%)	82 (28%)

Table 2. Number and percentage of participants declaring that they do not intend to get vaccinated, will wait some time before getting vaccinated, or who will get vaccinated as soon as a vaccine is available for them, in the Control Condition and in the Chatbot Condition, pre-treatment and post-treatment. * = Due to a technical issue, participants in Wave 2 were matched based on their answers to the question “did you interact with a chatbot during the first survey?” Results of Wave 2 are thus less reliable than those of Wave 1.

Before interacting with the chatbot, 145 out of 338 participants had positive attitudes toward the COVID-19 vaccine, after interacting with the chatbot they were 199, which corresponds to a 37% increase. Before interacting with the chatbot, 123 out of 338 participants said they did not want to take the COVID-19 vaccine, after interacting with the chatbot they were 99, which corresponds to a 20% decrease.




Figure 2. Density plots representing the distributions of participants' attitudes towards COVID-19 vaccines in Wave 1 before treatment (left panel) and after treatment (right panel), in the Chatbot Condition (blue) and the Control Condition (beige).

Confirmatory Analyses

Participants held more positive attitudes towards the COVID-19 vaccines after the experimental task in the Chatbot Condition than in the Control Condition ($\beta = 0.23$, [0.17, 0.29], $t(640) = 7.59$, $p < .001$). This relation held among the third of the participants initially holding the most negative attitudes towards the COVID-19 vaccines ($\beta = 0.37$, [0.20, 0.53], $t(207) = 4.31$, $p < .001$).

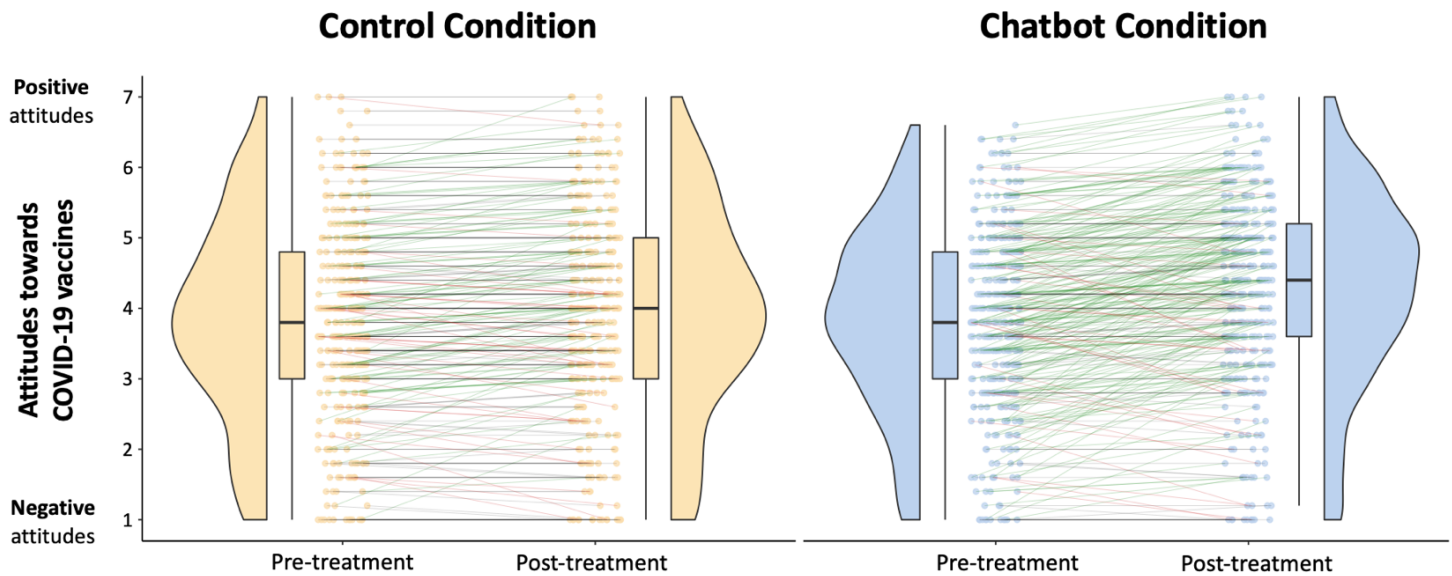


Figure 3. Evolution of participants' attitudes toward the COVID-19 vaccines in the Chatbot Condition and the Control Condition (Wave 1). Grey lines represent participants whose attitude toward the COVID-19 vaccines was similar after the treatment and before (i.e. a change of at most $\frac{1}{5}$ of a point on the COVID-19 vaccines attitude scale). Among the other participants, green (resp. red) lines represent participants whose attitude toward the COVID-19 vaccines was more positive (resp. negative) after the treatment than before.

Participants were more likely to report being willing to take the COVID-19 vaccines after the experimental task in the Chatbot Condition than in the Control Condition ($\beta = 0.12$, [0.07, 0.18], $t(640) = 4.37$, $p < .001$). This relation held among the third of the participants initially least willing to take the COVID-19 vaccines ($\beta = 0.50$, [0.25, 0.76], $t(231) = 3.96$, $p < .001$).

In the Chatbot Condition, time spent on the task was associated with more positive attitudes towards the COVID-19 vaccines after the experimental task ($\beta = 0.21$, [0.10, 0.31], $t(336) = 3.90$, $p < .001$).

In the Chatbot Condition, time spent on the task did not lead to a significantly greater willingness to take the COVID-19 vaccines after the experimental task ($\beta = 0.09$, $[-0.02, 0.19]$, $t(336) = 1.59$, $p = .13$).

Exploratory questions

We now turn to a series of pre-registered exploratory questions. First, we looked at what predicted holding positive attitudes towards the COVID-19 vaccines at baseline. We found that men ($\beta = 0.11$, $[0.04, 0.17]$, $p = .002$), older participants ($\beta = 0.08$, $[0.02, 0.14]$, $p = .023$), more educated participants ($\beta = 0.08$, $[0.01, 0.14]$, $p = .028$), participants with higher interpersonal trust ($\beta = 0.12$, $[0.06, 0.18]$, $p < .001$), participants who were more satisfied with the way the government handled the COVID-19 crisis ($\beta = 0.35$, $[0.28, 0.41]$, $p < .001$), participants trusting medical experts more than pseudo-medicine ($\beta = 0.32$, $[0.25, 0.38]$, $p < .001$), and participants higher in information seeking ($\beta = 0.10$, $[0.04, 0.16]$, $p = .004$) initially held more positive attitudes towards the COVID-19 vaccines.

Second, we look at what predicted intentions to take COVID-19 vaccines at baseline. We found that older participants ($\beta = 0.15$, $[0.08, 0.22]$, $p < .001$), participants with higher interpersonal trust ($\beta = 0.11$, $[0.04, 0.18]$, $p = .003$), participants who were more satisfied with the way the government handled the COVID-19 crisis ($\beta = 0.25$, $[0.19, 0.32]$, $p < .001$), participants trusting medical experts more than pseudo-medicine ($\beta = 0.30$, $[0.23, 0.37]$, $p < .001$), and participants higher in information seeking ($\beta = 0.14$, $[0.07, 0.21]$, $p < .001$) were initially more willing to take the COVID-19 vaccines. More educated participants ($\beta = 0.07$, $[0.00, 0.13]$, $p = .067$) and men ($\beta = 0.07$, $[0.00, 0.13]$, $p = .068$) were slightly, but not significantly more likely to be initially more willing to take the COVID-19 vaccines.

Third, we looked at what predicted positive attitudes change toward the COVID-19 vaccines after having interacted with the chatbot. We found that participants initially holding more negative attitudes toward the COVID-19 vaccines ($\beta = 0.28$, [0.14, 0.42], $p < .001$) and participants who were more satisfied with the way the government handled the COVID-19 crisis ($\beta = 0.21$, [0.10, 0.33], $p = .001$) displayed more positive attitudes change toward the COVID-19 vaccines after having interacted with the chatbot. Other variables were not significant (Gender: $\beta = 0.01$, [-0.12, 0.10], $p = .91$, Age: $\beta = 0.01$, [-0.12, 0.10], $p = .90$; Education: $\beta = 0.10$, [-0.01, 0.21], $p = .11$; Interpersonal trust: $\beta = 0.07$, [-0.04, 0.18], $p = .260$; Trust in medical experts: $\beta = 0.09$, [-0.02, 0.21], $p = .13$; Information seeking: $\beta = 0.03$, [-0.08, 0.14], $p = .59$).

After having interacted with the chatbot, and compared to the Control Condition, participants held more positive attitudes towards the COVID-19 vaccines on all five dimensions tested: safety ($\beta = 0.25$, [0.17, 0.32], $t(640) = 6.58$, $p < .001$), effectiveness ($\beta = 0.15$, [0.08, 0.22], $t(640) = 3.98$, $p < .001$), sufficient knowledge about the COVID-19 vaccines ($\beta = 0.30$, [0.20, 0.41], $t(640) = 5.59$, $p < .001$), trust in the people who produce the vaccines ($\beta = 0.21$, [0.14, 0.29], $t(640) = 5.44$, $p < .001$), and importance of vaccination ($\beta = 0.10$, [0.03, 0.17], $t(640) = 2.72$, $p = .010$).

Next, we examined the relationship between participants' initial attitudes, and attitude change. Specifically, we tested the interaction between participants' initial attitudes and the experimental condition on attitude change. We found that, compared to the Control Condition, participants initially holding more negative attitudes displayed slightly, but not significantly, more attitude change in favor of the COVID-19 vaccines ($\beta = 0.14$, [0.00, 0.29], $t(639) = 1.90$, $p = .073$).

On average, participants deemed the chatbot to be very intuitive (*Median* = 4, *M* = 3.53, *SD* = 0.98), their interaction with the chatbot to be quite pleasant (*Median* = 3, *M* = 3.26, *SD* = 0.99), and not very frustrating (*Median* = 4, *M* = 4.22, *SD* = 0.86). They also found the information presented in the chatbot to be neither too complex nor too simple (*Median* = 3, *M* = 2.99, *SD* = 0.55).

Exploratory analyses of the second wave

Due to a technical problem, we were not able to match participants between the first and the second wave. To infer the condition participants had been randomized to in Wave 1, we relied on their answers to the question “did you interact with a chatbot during the first survey?” We did not exclude any participants. 298 participants declared having interacted with the chatbot and 315 participants declared not having interacted with the chatbot. As a result of these limitations, we treat these results as exploratory, and urge caution in their interpretation.

Participants in the Chatbot Condition had more positive attitudes towards COVID-19 vaccines in Wave 2 (*M* = 4.27, *SD* = 1.38) than at baseline in Wave 1 (*M* = 3.82, *SD* = 1.28; *d* = 0.34, [0.18, 0.49], $t(608.77) = 4.23$, $p < .001$). However, this was also true for participants in the Control Condition (pre-treatment attitudes: *M* = 3.81, *SD* = 1.41; wave two attitudes: *M* = 4.15, *SD* = 1.43; *d* = 0.24, [0.08, 0.39], $t(617.79) = 2.93$, $p = .003$). In Wave 2, there was no significant difference between participants' attitudes in the Chatbot and in the Control conditions (*d* = 0.09, [-0.07, 0.25], $t(610.75) = 1.10$, $p = .27$); a pattern that is similar for vaccination intentions. For participants in the Chatbot Condition, intentions remained higher at Wave 2 than at baseline in Wave 1 (pre-treatment intentions: *M* = 1.81, *SD* = 0.71; wave two intentions: *M* = 1.99, *SD* = 0.75; *d* = 0.25, [0.09, 0.40], $t(614.01) = 3.08$, $p < .002$), but

intentions also increased in the Control Condition (pre-treatment intentions: $M = 1.84$, $SD = 0.74$; wave two intentions: $M = 1.98$, $SD = 0.76$; $d = 0.19$, $[0.03, 0.34]$, $t(617.99) = 2.31$, $p = .021$), leading to an absence of difference during Wave 2 ($d = 0.01$, $[-0.17, 0.15]$, $t(609.7) = 0.15$, $p = .88$).

In the Chatbot Condition, 45% of participants reported having tried to convince other people (typically, between two and five) of their position on the COVID-19 vaccines, and these participants were more likely to have positive attitudes and intentions towards the COVID-19 vaccines (attitudes: $\beta = 0.28$, $[0.21, 0.36]$, $p < .001$; intentions: $\beta = 0.33$, $[0.25, 0.40]$, $p < .001$). 72% of these participants reported having used information from the Chatbot in their attempts to convince others. 38% of the participants reported being willing to share the chatbot in at least one way (social networks, 11%; entourage, 37%; other means, 9%), and these participants had more positive attitudes and intentions towards the COVID-19 vaccines (attitudes: $\beta = 0.26$, $[0.15, 0.37]$, $p < .001$; intentions: $\beta = 0.25$, $[0.14, 0.36]$, $p < .001$).

Discussion

Using a simple chatbot, we gave participants access to a relatively exhaustive list of questions and answers about the COVID-19 vaccines. We compared participants who had interacted with the chatbot to a control group who only read a brief text about how vaccines work in general. Participants' attitudes towards the COVID-19 vaccines, and their intention to get vaccinated were measured before and after treatment. In contrast with the Control Condition, participants in the Chatbot Condition developed more positive attitudes towards the COVID-19 vaccines (on all five dimensions evaluated), and they declared being more willing to take the vaccine. The effects were $\beta = 0.23$ and $\beta = 0.12$ respectively.

The amount of change in attitudes was related to time spent interacting with the chatbot, which suggests that participants did change their minds thanks to the information provided by the chatbot. Importantly, we did not observe any backfire effect. On the contrary, and in line with previous findings (e.g. Altay et al., 2020; Altay & Lakhlifi, 2020; Bode & Vraga, 2015; Vraga et al., 2020; Vraga & Bode, 2017), the participants whose initial attitudes were the most negative shifted the most towards positive attitudes (for the most negative third, average attitude change = 0.54 on a scale of 1 to 7, and 0.39 for the other two thirds).

Unfortunately, our Wave 2 results are compatible with two interpretations. The first is that the gains in attitudes and intentions after interacting with the chatbot persisted, but that participants in the control condition were also exposed to pro-vaccination information, because of an intense media coverage of the vaccination campaign in France. This may have led them to catch up with the participants who had already acquired that information through the chatbot. The second interpretation is that participants in the chatbot condition quickly reverted to their original attitudes and intentions, and that those were then buoyed by the media coverage, along with those of the participants in the control condition. Parsimony favors the first explanation, but the evidence remains inconclusive.

The second wave survey showed that nearly half of the participants (45%) who recalled having seen the chatbot in Wave 1 had tried to convince others to share their views on vaccination, and 72% of them reported to have used information provided by the chatbot during these conviction attempts. Moreover, 38% of participants—which were more likely to be pro-vaccination—declared wanting to share the chatbot in one way or another. These results suggest that the chatbot could play a useful role beyond providing information to those directly exposed to it, as people use and share the chatbot to others (as in two-step and

multistep flow models of communication, see, e.g., Ahn et al., 2014; Katz & Lazarsfeld, 1955). More generally, the role of interpersonal communication in the dissemination of reliable information is being increasingly recognized in the field of risk communication (see, e.g., Altay & Lakhlifi, 2020; Altay & Mercier, 2020; Goldberg et al., 2019; Sloane & Wiles, 2019).

An exploratory analysis of users' behavior on the chatbot in ESM suggests that they were more interested in learning about the safety of COVID-19 vaccines than about their efficacy and that the non-interactive chatbot option was used as a complement to the interactive chatbot.

Consistent with previous findings on COVID-19 vaccines hesitancy in France (Hacquin, Altay, et al., 2020; Ward et al., 2020), we found that being a woman, being young, being less educated, and being unsatisfied with the way the government handled the COVID-19 crisis, were associated with more negative attitudes towards the COVID-19 vaccines. Overall, and by far, the best predictors of COVID-19 vaccines hesitancy (in intentions and attitudes) were being dissatisfied with the way the government handled the COVID-19 crisis ($\beta = 0.25$ & $\beta = 0.35$) and low trust in medical experts compared to alternative medicine ($\beta = 0.30$ & $\beta = 0.32$).

Interacting with the COVID-19 chatbot led to less attitude change than interacting with the GMOs chatbot in Altay and colleagues (2020) ($\beta = 0.23$ compared to $\beta = 0.66$). Two main reasons likely explain this difference. First, the arguments (in terms of number of scientific publications, etc.) in favor of the safety of GMOs were stronger than the arguments in favor of the COVID-19 vaccines, especially at the time when the study was conducted (i.e.,

December 2020). Second, everything else being equal, chatbots should be most effective at changing people's mind when they are the least informed. As a result, the more people know about a given topic, the harder it should be to change their mind. In this regard, COVID-19 vaccines were a more challenging test for the chatbot than GMOs. COVID-19 vaccines were in the media spotlight when we conducted the study. This was not the case for GMOs. People likely had stronger priors and opinions about COVID-19 vaccines than on GMOs (for instance in the U.K. a large share of people declare having no opinion on GM food, Burke, 2004).

The effect observed in the present study, even if of a small size, could have important practical consequences at a population level. For instance, if the chatbot had been deployed on the COVID mobile application developed by the French government "TousAntiCovid," and that it had been used by its 20 million users, it could have swayed 1.4 million vaccine hesitant individuals towards vaccination. This calculation doesn't take into account the indirect effects of the chatbot, by which participants discuss with their peers the information presented by the chatbot, and which could amplify its effects (especially in light of the finding that one third of participants at Wave 2 had used information gleaned on the chatbot in discussions).

More broadly, chatbots could be particularly useful to fill the gap between public opinion and scientists when laypeople are uninformed (see the Deficit Model of Communication, Sturgis & Allum, 2004). However, chatbots are less likely to be effective if the gap stems from politically motivated science denialism (e.g. Kahan et al., 2011, 2012). The use of chatbots to facilitate scientific communication (Altay et al., 2020) has been theorized to be effective on the basis of the interactive theory of reasoning (Mercier & Sperber, 2017). Even if the results proved inconclusive in terms of testing specific predictions from this theory, it is still

noteworthy that the theory could be used as a heuristic to develop effective means of communication.

Limitations

The present study has several limitations. First, its scope, as we did not investigate the mechanisms that led to the positive attitude change in the Chatbot Condition. Previous work suggests that the interactivity of the Chatbot is not central (Altay et al. 2020), but the dialogic format—which makes it easy to find relevant information—could be. In sum, this paper offers evidence that a chatbot can be used to inform people about the COVID-19 vaccines, but not why it is the case (for an investigation of these mechanisms see, Altay et al. 2020). Future work should try to disentangle the effect of interactivity from the effect of the dialogic format (for instance by having a text organized in a non-dialogic format, an interactive chatbot, and a non-interactive chatbot). Moreover, interactivity could have difficult-to-measure benefits, such as increasing people's motivation to read and engage with the arguments.

A second limitation of the present study is the unknown about its impact in the wild. Outside of experimental settings, we don't know how willing people would be to interact with the chatbot. This metric is key to measure the chatbot's conversion rate and have a good estimate chatbot's potential impact if it were widely deployed. Other ways of communicating information, e.g., short videos in a TikTok format, could be as efficient, if not more efficient, at capturing people's attention and ultimately conveying information to the general public.

A third limitation concerns the declarative nature of our dependent variables. Vaccine attitudes and declared intentions to get vaccinated are only indirect and imperfect measures of behaviors. We know that attitudes don't always translate into behaviors (e.g., Mainieri et al.,

1997). The existence of this gap between attitudes and behaviors suggests that even the most efficient communication campaigns won't be enough on their own: they are necessary, but not sufficient. This is why, in addition to effective communication campaigns, governments should do their best to facilitate vaccination, for instance by making it free and easy to access (Chevallier et al., 2021).

The fourth limitation regards its reception among diverse segments of the population. In contrast with a representative sample of the French population, our sample is younger (below 35: 46% [26%], between 35 and 65: 51% [51%], over 65: 3% [23%]), more educated (more than a high school diploma: 66% [53%], high school diploma: 23% [17%], less than a high school diploma: 10% [30%]), and more masculine (54% men [48%]). It is safe to assume that the chatbot can be used by a young and educated population. However, before deploying the chatbot at large scale in the general population, its efficacy should be tested on people with less than a high school diploma and, importantly, on people over 65 whose digital skills tend to be lower.

Conclusion

Messages that aim to change people's attitudes towards vaccines, or to increase their intention to take vaccines, often fall on deaf ears. One reason why people might be so reluctant to change their minds is that health messages tend to be brief, failing to anticipate most of the concerns people might have. To address this issue, we presented participants with a chatbot that answers the most common questions about the COVID-19 vaccines, as well as questions these answers might raise in turn.

Compared to a control group that had only been exposed to a brief text explaining the general concept of vaccination, participants given the opportunity to interact with the chatbot developed more positive attitudes towards COVID-19 vaccines, and higher intentions to vaccinate. Participants spent a significant amount of time interacting with the chatbot (between 5 to 12 minutes for half of the participants), and the more time they spent, the more they changed their minds. The effects were substantial, with a 37% increase in participants holding positive attitudes, and a 20% decrease in participants saying they would not get vaccinated. Moreover, we did not find evidence for backfire effects. In fact, the participants who held the most negative views changed their opinions the most. Finally, although exploratory, results from a second wave taking place between one and two weeks after the initial experiment suggest that the changes in attitudes and intentions might persist beyond the initial exposure. The second wave also shows that the chatbot can be leveraged by people to convince others, either as they rely on the chatbot's information, or as they share it with others.

Our results suggest that a properly scripted and regularly updated chatbot could offer a powerful resource to help fight hesitancy towards COVID-19 vaccines. Besides its direct effect on vaccine hesitant individuals, the chatbot could prove invaluable to pro-vaccination individuals, including professionals looking for information to use in interpersonal communication with vaccine hesitant individuals.

Data availability

The data associated with this research, together with the materials, are available at the following address: <https://osf.io/8q3b2/>.

Code availability

The R scripts associated with this research are available at the following address:

<https://osf.io/8q3b2/>. For data visualization, especially Figure 3, see: van Langen, 2020.

Ethics information

The present research received approval from an ethics committee (CER-Paris Descartes; N° 2019-03-MERCIER). Participants had to give their informed consent to participate in the study.

Acknowledgements

We are grateful to Tom Stafford, Charlotte Brand and Pierre Verger for having reviewed the manuscript (version 3 on PsyArXiv) for *Rapid Review: COVID-19*. Their reviews, and our response, can be found here:

<https://rapidreviewscovid19.mitpress.mit.edu/pub/akskfghv/release/1>.

We would like to warmly thank the two medical experts who carefully checked the chatbot's information and made numerous corrections: Odile Launay and Jean-Daniel Lelièvre. We are grateful to Camille Lakhlifi, Camille Rozier, Mariam Chammat, Delphine Grison, Rita Abdel Sater, and Léonard Guillou for their proofreading and suggestions on the chatbot's information. We also thank Vincent Laine from La Fabrique à Chatbot for his invaluable help, providing us with a tailor-made chatbot in no time, and for being very patient with our never-ending requests.

The main funding for this work was an ANR grant COVID-19_2020_BEHAVIRAL. We also received funding from the two following grants: ANR-17-EURE-0017 to FrontCog, and ANR-10-IDEX-0001-02 to PSL.

References

Ahn, T., Huckfeldt, R., & Ryan, J. B. (2014). *Experts, activists, and democratic politics: Are electorates self-educating?* Cambridge University Press.

Altay, S., & Lakhlifi, C. (2020). Are science festivals a good place to discuss heated topics? *Journal of Science Communication*, 19(1), A07. <https://doi.org/10.22323/2.19010207>

Altay, S., & Mercier, H. (2020). Framing messages for vaccination supporters. *Journal of Experimental Psychology: Applied*, 26(4), 567–578. <https://doi.org/10.1037/xap0000271>

Altay, S., Schwartz, M., Hacquin, A., Allard, A., Blancke, S., & Mercier, H. (2020). *Scaling up Interactive Argumentation by Providing Counterarguments with a Chatbot*. <https://doi.org/10.6084/m9.figshare.13122527.v1>

Andrews, P., Manandhar, S., & De Boni, M. (2008). Argumentative human computer dialogue for automated persuasion. *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*, 138–147.

Barari, S., Caria, S., Davola, A., Falco, P., Fetzner, T., Fiorin, S., Hensel, L., Ivchenko, A., Jachimowicz, J., & King, G. (2020). Evaluating COVID-19 public health messaging in Italy: Self-reported compliance and growing mental health concerns. *MedRxiv*.

Betsch, C., & Sachse, K. (2013). Debunking vaccination myths: Strong risk negations can increase perceived vaccination risks. *Health Psychology*, 32(2), 146.

Bilancini, E., Boncinelli, L., Capraro, V., Celadin, T., & Di Paolo, R. (2020). The effect of norm-based messages on reading and understanding COVID-19 pandemic response governmental rules. *ArXiv Preprint ArXiv:2005.03998*.

- Bode, L., & Vraga, E. K. (2015). In related news, that was wrong: The correction of misinformation through related stories functionality in social media. *Journal of Communication, 65*(4), 619–638.
- Brewer, N. T., Chapman, G. B., Rothman, A. J., Leask, J., & Kempe, A. (2017). Increasing vaccination: Putting psychological science into action. *Psychological Science in the Public Interest, 18*(3), 149–207.
- Broockman, D., & Kalla, J. (2016). Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science, 352*(6282), 220–224.
- Burke, D. (2004). GM food and crops: What went wrong in the UK? *EMBO Reports, 5*(5), 432–436.
- Capraro, V., & Barcelo, H. (2020). Priming reasoning increases intentions to wear a face covering to slow down COVID-19 transmission. *ArXiv Preprint ArXiv:2006.11273*.
- Chalaguine, L. A., Hunter, A., Hamilton, F. L., & Potts, H. W. (2019). Impact of Argument Type and Concerns in Argumentation with a Chatbot. *ArXiv Preprint ArXiv:1905.00646*.
- Chanel, O., Luchini, S., Massoni, S., & Vergnaud, J.-C. (2011). Impact of information on intentions to vaccinate in a potential epidemic: Swine-origin Influenza A (H1N1). *Social Science & Medicine, 72*(2), 142–148.
- Chevallier, C., Hacquin, A.-S., & Mercier, H. (2021). COVID-19 vaccine hesitancy: Shortening the last mile. *Trends in Cognitive Sciences*.
- Claidière, N., Trouche, E., & Mercier, H. (2017). Argumentation and the diffusion of counter-intuitive beliefs. *Journal of Experimental Psychology: General*.
- Community Preventive Services Task Force. (2015). *Increasing Appropriate Vaccination: Provider Education When Used Alone*.
<https://www.thecommunityguide.org/sites/default/files/assets/Vaccination-Provider-Education-Alone.pdf>

de Figueiredo, A., Simas, C., Karafillakis, E., Paterson, P., & Larson, H. J. (2020). Mapping global trends in vaccine confidence and investigating barriers to vaccine uptake: A large-scale retrospective temporal modelling study. *The Lancet*, 396(10255), 898–908.

Dubé, E., Gagnon, D., MacDonald, N. E., & SAGE Working Group on Vaccine Hesitancy. (2015). Strategies intended to address vaccine hesitancy: Review of published reviews. *Vaccine*, 33(34), 4191–4203. <https://doi.org/10.1016/j.vaccine.2015.04.041>

Duquette, N. (2020). “Heard” immunity: Messages emphasizing the safety of others increase intended uptake of a COVID-19 vaccine in some groups¹. *Covid Economics*, 39.

Favero, N., & Pedersen, M. J. (2020). How to encourage “Togetherness by Keeping Apart” amid COVID-19? The ineffectiveness of prosocial and empathy appeals. *Journal of Behavioral Public Administration*, 3(2).

Freeman, S., Eddy, S. L., McDonough, M., Smith, M. K., Okoroafor, N., Jordt, H., & Wenderoth, M. P. (2014). Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences*, 111(23), 8410–8415.

Goldberg, M. H., van der Linden, S., Maibach, E., & Leiserowitz, A. (2019). Discussing global warming leads to greater acceptance of climate science. *Proceedings of the National Academy of Sciences*, 116(30), 14804–14805.

Greenwald, A. G. (1968). Cognitive learning, cognitive response to persuasion, and attitude change’. In A. G. Greenwald, T. C. Brock, & T. M. Ostrom (Eds.), *Psychological Foundations of Attitudes* (pp. 147–170). Academic Press.

Guess, A., & Coppock, A. (2018). Does counter-attitudinal information cause backlash? Results from three large survey experiments. *British Journal of Political Science*, 1–19.

Hacquin, A.-S., Altay, S., de Araujo, E., Chevallier, C., & Mercier, H. (2020). *Sharp rise in vaccine hesitancy in a large and representative sample of the French population: Reasons for*

vaccine hesitancy. <https://doi.org/10.31234/osf.io/r8h6z>

Hacquin, A.-S., Mercier, H., & Chevallier, C. (2020). *Improving preventive health behaviors in the COVID-19 crisis: A messaging intervention in a large nationally representative sample.*

Johnson, D. W., Johnson, R. T., & Stanne, M. B. (2000). *Cooperative learning methods: A meta-analysis.*

Jordan, J., Yoeli, E., & Rand, D. (2020). *Don't get it or don't spread it? Comparing self-interested versus prosocially framed COVID-19 prevention messaging.*

Kahan, D., Jenkins-Smith, H., & Braman, D. (2011). Cultural cognition of scientific consensus. *Journal of Risk Research, 14*(2), 147–174.

Kahan, D., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., & Mandel, G. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change, 2*(10), 732–735.

Katz, E., & Lazarsfeld, P. F. (1955). *Personal influence: The part played by people in the flow of mass communications.* Free Press.

Kaufman, J., Ryan, R., Walsh, L., Horey, D., Leask, J., Robinson, P., & Hill, S. (2018). Face-to-face interventions for informing or educating parents about early childhood vaccination. *Cochrane Database of Systematic Reviews, 5.*

King, A. (1990). Enhancing peer interaction and learning in the classroom through reciprocal questioning. *American Educational Research Journal, 27*(4), 664–687.

Laughlin, P. R. (2011). *Group problem solving.* Princeton University Press.

Mainieri, T., Barnett, E. G., Valdero, T. R., Unipan, J. B., & Oskamp, S. (1997). Green buying: The influence of environmental concern on consumer behavior. *The Journal of Social Psychology, 137*(2), 189–204.

Mercier, H. (2016). The argumentative theory: Predictions and empirical evidence. *Trends in*

Cognitive Sciences, 20(9), 689–700.

Mercier, H. (2020). *Not Born Yesterday: The Science of Who We Trust and What We Believe*. Princeton University Press.

Mercier, H., & Sperber, D. (2017). *The enigma of reason*. Harvard University Press.

Minozzi, W., Neblo, M. A., Esterling, K. M., & Lazer, D. M. (2015). Field experiment evidence of substantive, attributional, and behavioral persuasion by members of Congress in online town halls. *Proceedings of the National Academy of Sciences*, 112(13), 3937–3942.

Nyhan, B., & Reifler, J. (2015). Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information. *Vaccine*, 33(3), 459–464.

Nyhan, B., Reifler, J., Richey, S., & Freed, G. L. (2014). Effective messages in vaccine promotion: A randomized trial. *Pediatrics*, 133(4), e835–e842.

Prince, M. (2004). Does active learning work? A review of the research. *Journal of Engineering Education*, 93(3), 223–231.

Robinson, E., Jones, A., & Daly, M. (2020). International estimates of intended uptake and refusal of COVID-19 vaccines: A rapid systematic review and meta-analysis of large nationally representative samples. *MedRxiv*.

Rosenfeld, A., & Kraus, S. (2016). Strategical argumentative agent for human persuasion. *Proceedings of the Twenty-Second European Conference on Artificial Intelligence*, 320–328.

Sadaf, A., Richards, J. L., Glanz, J., Salmon, D. A., & Omer, S. B. (2013). A systematic review of interventions for reducing parental vaccine refusal and vaccine hesitancy. *Vaccine*, 31(40), 4293–4304.

Shi, Y., Yang, H., MacLeod, J., Zhang, J., & Yang, H. H. (2020). College students' cognitive learning outcomes in technology-enabled active learning environments: A meta-analysis of the empirical literature. *Journal of Educational Computing Research*, 58(4), 791–817.

Sloane, J. D., & Wiles, J. R. (2019). Communicating the consensus on climate change to

college biology majors: The importance of preaching to the choir. *Ecology and Evolution*.

Sturgis, P., & Allum, N. (2004). Science in society: Re-evaluating the deficit model of public attitudes. *Public Understanding of Science*, 13(1), 55–74.

Swire-Thompson, B., DeGutis, J., & Lazer, D. (2020). *Searching for the backfire effect: Measurement and design considerations*.

Team, R. C. (2017). R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria*. URL <https://www.R-project.org>.

Team, Rs. (2015). RStudio: Integrated development for R. RStudio. Inc., Boston, MA, 700.

Trouche, E., Sander, E., & Mercier, H. (2014). Arguments, more than confidence, explain the good performance of reasoning groups. *Journal of Experimental Psychology: General*, 143(5), 1958–1971.

Trueblood, J. S., Sussman, A. B., & O’Leary, D. (2020). The Role of General Risk Preferences in Messaging About COVID-19 Vaccine Take-Up. Available at SSRN 3649654.

van Langen, J. (2020). *Open-visualizations in R and Python*. <https://github.com/jorvlan/open-visualizations>

Vraga, E. K., & Bode, L. (2017). Using expert sources to correct health misinformation in social media. *Science Communication*, 39(5), 621–645.

Vraga, E. K., Bode, L., & Tully, M. (2020). Creating news literacy messages to enhance expert corrections of misinformation on Twitter. *Communication Research*, 0093650219898094.

Ward, J. K., Alleaume, C., & Peretti-Watel, P. (2020). *The French public’s attitudes to a future COVID-19 vaccine: The politicization of a public health issue*.

Wood, T., & Porter, E. (2019). The elusive backfire effect: Mass attitudes’ steadfast factual adherence. *Political Behavior*, 41(1), 135–163.

