Spreading consensus: Correcting misperceptions about the views of the medical community has lasting impacts on Covid-19 vaccine take-up

Vojtěch Bartoš[1*], Michal Bauer[2,3], Jana Cahlíková[4], Julie Chytilová[2,3]

[1] Department of Economics, University of Munich, Geschwister-Scholl-Platz 1, D-80539 München, Germany

[2] CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Academy of Sciences of the Czech Republic, Prague, Politických vězňů 7, 111 21, Czech Republic

[3] Institute of Economic Studies, Faculty of Social Sciences, Charles University, Opletalova 26, 110 00, Prague, Czech Republic

[4] Department of Public Economics, Max Planck Institute for Tax Law and Public Finance, Marstallplatz 1, D - 80539 München, Germany

Correspondence should be addressed to: vojtech.bartos@econ.lmu.de.

[This version: October 1, 2021]

Abstract

Identifying sources of vaccine hesitancy is one of the central challenges in fighting the Covid-19 pandemic. In this study, we focus on the role of public misperceptions about doctors' views. Motivated by widespread concern that media reports create uncertainty in how people perceive expert opinions, even when broad consensus exists, we elicited trust in Covid-19 vaccines among 9,650 doctors in the Czech Republic. We found evidence of strong consensus: 90% of doctors trust the approved vaccines. Next, we conducted a nationally representative survey (N=2,101) and document systemic misperception of doctors' views: more than 90% of respondents underestimate doctors' trust; the most common belief is that only 50% of doctors trust the vaccines. Finally, we integrate randomized provision of information about the true views held by doctors into a longitudinal data collection and study its effects across eight waves of data collection covering five-months period during which the vaccine was gradually rolled out. We find that the treatment recalibrates beliefs, leads to a lasting and stable increase in vaccine demand, and that treated individuals are 5 percentage points more likely to be vaccinated five months after the intervention. This paper illuminates how the engagement of professional medical associations, with their unparalleled capacity to elicit individual views of doctors on a large scale, may help to create a cheap, scalable intervention that corrects misperceptions and has lasting impacts on behavior.

Introduction

People's reluctance to take up vaccines represents a significant challenge in containing the spread of deadly infectious diseases[1,2]. Covid-19 is a salient example of a disease with profound economic, social, and health impacts which can be controlled by large-scale vaccination if enough people choose to be vaccinated. Nevertheless, a large percentage of people are hesitant to get a vaccine, many countries are still relatively far from the threshold necessary to achieve herd immunity[3], which is typically estimated to be between 70-85% of the population[4,5]. Consequently, rigorous evidence on scalable approaches that can help to overcome people's hesitancy to take a Covid-19 vaccine is a global policy priority[6–8]. Existing research has made important progress in documenting the roles of providing financial incentives[9,10], reminders[7,8], information about efficacy of the vaccines[11,12], the role of misinformation[13] on the public's intentions to get vaccinated, and more recently, also on their actual decisions to take a vaccine[8] shortly after an intervention. However, little is known about whether cheap, scalable strategies with the potential to cause lasting increases in people's vaccination intentions exist, and whether they can also increase people's willingness to follow through on those intentions and actually get a vaccine. A focus on the persistence of intervention impacts is especially important for vaccines like those against Covid-19, which are distributed in phases to different demographic groups due to capacity constraints, and for which full vaccination requires multiple doses.

In many surveys across the globe, people report that the views of doctors are highly trusted[14]. This makes it crucial to understand how people perceive doctors' views about the Covid-19 vaccine. In this paper, we pursue the hypothesis that reluctance to adopt the vaccine originates, in part, in misperceptions about the distribution of aggregate views of the medical community: many people may fail to recognize that there is a broad consensus in favor of the vaccine among doctors. Further, we argue and show that professional associations can serve as aggregators of individual views in a medical community, by helping to implement surveys eliciting doctors' views on a large scale. Disseminating information of a broad consensus, when one exists, can lead to people updating their perceptions of doctors' views and, in turn, may induce lasting changes in vaccination demand and take up.

Our focus on public misperceptions of doctors' views is motivated by a widespread concern that media coverage can create uncertainty and polarization in how people perceive expert views, even when a broad consensus actually exists. In terms of traditional media, a desire to appear neutral often motivates journalists to provide a "balanced" view by giving roughly equal time to both sides of an argument[15,16], creating an impression of controversy and uncertainty[17]. Such "falsely-balanced" reporting has been shown to be a characteristic element of policy debates ranging from climate change[15,16] to health issues, including links between tobacco and cancer and potential side effects of vaccines[17,18]. In the context of the Covid-19 vaccines, casual observation suggests that media outlets often feature expert opinions highlighting the efficacy of approved Covid-19 vaccines together with skeptical experts who voice concerns about rapid vaccine development and untested side effects. The media usually do not specify which claims are supported by the wider medical community, leading the World Health Organization to warn media outlets against engaging in false-balance reporting[19]. Furthermore, polarization of beliefs can arise due to echo-chambers -- people choosing to be exposed to expert opinions or opinion programs that either fuel their pre-existing fears about the vaccine or those that strongly approve of it[20–22].

We study these issues in the Czech Republic, a suitable setting, given the observed level of vaccine hesitancy among a large share of its population, similar to the situation in many other countries. At the time of data collection, the vaccine acceptance rate in the Czech Republic was around 65%, compared to 55-90% in other countries globally. At the same time, the Czech

Republic ranks close to the median level of trust and satisfaction with medical doctors, based on a comparison of 29 countries[14]. Supplementary Information provides more background.

We start by documenting and quantifying public misperceptions about doctors' views on the Covid-19 vaccines. To do so, we partnered with the Czech Medical Chamber, to aggregate the views of a uniquely wide spectrum of the medical community. Shortly before the Covid-19 vaccine roll-out began in January 2021, we elicited responses via a short online survey from almost 10,000 doctors on their trust in the vaccines, their willingness to get vaccinated themselves, and to recommend the vaccine to their patients. We find strong evidence of consensus: 90% intend to get vaccinated themselves, 89% trust the approved vaccines, and 95% would recommend vaccination to their patients. The responses are broadly similar across gender, age, locality, and seniority. Next, in order to quantify the extent of public misperceptions of the views of the medical community, we conducted a survey with a nationally representative online sample of the adult population (N=2,101). The participants were asked to estimate what percentage of doctors trust the vaccines and want to get vaccinated themselves. We find evidence of systemic and widespread misperceptions: more than 90% of people underestimate trust and vaccination intentions of doctors, with most people expecting that only 50% of doctors trust the vaccine and intend to be vaccinated.

These findings set the stage for our main experiment, in which we test whether randomized provision of information about the actual views of doctors can recalibrate public beliefs and, more importantly, cause a lasting shift in intentions and ultimately increase vaccination take-up. The experimental design aims to make progress on two important empirical challenges that are common in experiments on determinants of demand for Covid-19 vaccines. First, since an intention–behavior gap has been documented in the context of flu vaccines and other health behaviors[23], measuring both vaccination intentions and vaccination take-up allows us to test whether treatment effects on vaccination intentions translate into behavioral changes of a similar magnitude. The initial set of studies on Covid-19 vaccination, typically implemented before the vaccines became available, only tested impacts on intentions[9,12,13], although a recent exception exists[8].

Second, most experiments designed to correct misperceptions about the views of others, and other information provision experiments in various domains, including migration, health, and political behavior, document treatment effects to be substantially smaller when measured with a delay[24,25]. In theory, the worry is that individual perceptions about doctors' views might shift significantly between the time when the treatment takes place and when people decide whether to actually get vaccinated, due to reasons such as regression of perceptions to the mean, biased recall, or motivated memory[26]. On the other hand, researchers have suggested that providing facts about a widely-shared consensus of trustworthy experts might be resilient to these forces[18], since the treatment may reduce incentives to seek new information and condenses complex information into a simple fact ("90% of doctors trust the approved vaccines"), which is easy to remember. Understanding whether providing information about medical consensus has temporary or lasting impacts on vaccination demand is informative for policy, in terms of whether a one-off information campaign is sufficient, or whether the timing of messages needs to be tailored for different groups of people who become eligible for a vaccine at different points in time, and also whether such an information campaign needs to be repeated in cases of multiple-dose vaccines.

To address these issues, our experiment is integrated into longitudinal data collection with low attrition rates. The treatment was implemented in March 2021. We use data from eight consecutive waves collected from March to August 2021, covering the early period when the vaccine was scarce, later when it became gradually available to more demographic groups, and at the end, when it was easily available to all adults. This is reflected in the vaccination rates, which increased in our sample from 9% in March to 20% in May and to nearly 70% in June, and then remained relatively stable. Such intensive longitudinal data-collection approach allows us to estimate whether disseminating information on the consensus view of the medical community has immediate impacts on beliefs and intentions to get the vaccination shortly after the intervention, and also whether the effects are lasting and translate into actually getting vaccinated, even though most of the participants became eligible for the vaccine only many weeks after the intervention.

Consensus of the medical community

We conducted a supplementary survey to gather doctors' views on the Covid-19 vaccines in February 2021. The survey was implemented in partnership with the Czech Medical Chamber (CMC), to maximize coverage of the medical community. The CMC contact list includes the whole population of doctors in the country, because membership is compulsory. All doctors who communicate with CMC electronically (70%) were asked to participate in a short survey, using the Qualtrics platform. 9,650 doctors (24% of those contacted) answered the survey. The doctors in our sample work in all regions of the country, are on average 52 years old, 64% are female, and 62% have more than 20 years of experience. The summary statistics are presented in Supplementary Table 1, which also documents that our sample of doctors is similar, in terms of age, gender, and location, to the overall population of medical doctors in the Czech Republic.

Figure 1 shows the distribution of doctors' responses. A clear picture arises, suggesting that a broad consensus on Covid-19 vaccines exists in the medical community: 89% trust the vaccine (9% do not know and 2% do not trust it), 90% intend to get vaccinated (6% do not know and 4% do not plan to get vaccinated) and 95% plan to recommend that their patients take a vaccine (5% do not). These responses are broadly similar across gender, age, years of medical practice, and size of the locality in which the doctors live: for all sub-groups we find the share of positive answers to all questions ranges between 85-100% (Supplementary Table 2).

Methods

Sample: Our main sample consists of participants of the longitudinal online data collection "Life during the pandemic", organized by the authors in cooperation with PAQ Research; the data was collected by the NMS survey agency. The panel started in spring 2020 to provide real-time data on developments in economic, health, and social conditions during the Covid-19 pandemic. We use data from eight consecutive waves of data collections conducted between March and August 2021.

The experimental manipulation was implemented on March 15, 2021, which we label here Wave0 and in which 2,101 individuals took part. The sample from Wave0 is our "base sample" (n = 2,101, 1,052 females / 1,049 males, mean age 52.9 (s.d. = 15.98), youngest 18, oldest 92). The base sample is broadly representative of the adult Czech population in terms of sex, age, education, region, municipality size, employment status before the Covid-19 pandemic, age x sex, and age x education. Prague and municipalities with more than 50,000 inhabitants are oversampled (boost 200%). Sample statistics are presented in Supplementary Table 3. The sample is also close to being representative of the adult Czech population in terms of attitudes to Covid-19 vaccines. The development of the proportion of people who got vaccinated in our sample very closely mimics the actual vaccination rate in the Czech Republic (Supplementary Figure 1), when we weight observations in our sample to be representative in terms of observable characteristics.

An important feature of the panel is that participants agreed to be interviewed regularly, and attrition between waves is relatively low. The response rate, as compared to the base sample from Wave0, was 92% in Wave1 (March), 92% in Wave2 (April), 90% in Wave3 (May), 89% in Wave4 (May), 85% in Wave5 (June), 77% in Wave6 (July) and 84% in Wave7 (August). In Supplementary Table 4, we show that attrition in all seven follow-up waves is uncorrelated with the treatment. Furthermore, 1,413 participants (67%) answered the survey in all eight waves of data collection: they form the "fixed sample". Consequently, in the analysis, we report the main results for (i) all participants from the base sample who responded in a given wave as well as for (ii) the "fixed sample", composed of individuals who participated in all eight waves, eliminating the potential role of differences in samples across waves.

Experimental manipulation. The participants were randomly allocated to either the CONSENSUS (n=1,050) or CONTROL (n=1,051) condition in Wave0. In the CONSENSUS condition, they were informed that the Czech Medical Chamber conducted a large survey of almost 10,000 doctors from all parts of the country to collect their views on Covid-19 vaccines. They were also informed that the views were similar for doctors of different genders, ages, and regions. Then, the participants viewed three figures displaying the distribution of responses of doctors regarding trust in the vaccines, willingness to get vaccinated themselves, and to recommend the vaccine to their patients. Each of the figures was complemented by a short written summary and was displayed on a separate screen. The exact wording and the figures are provided in the Supplementary Online Material. In the CONTROL condition, the participants did not receive any information about the survey of medical doctors.

Supplementary Table 3 provides evidence that the range of individual characteristics, including gender, age, household income, employment status, education, region and size of municipality, as well as vaccination intentions measured three weeks before the intervention are balanced across the experimental condition. The exceptions are "own vaccination" and "prior beliefs about opinions of the doctors" (measured in Wave0). Prior to the intervention, compared to participants in the CONTROL condition, the individuals in the CONSENSUS condition were less likely to be vaccinated themselves, and expected a smaller percentage of doctors to trust the vaccine and to intend to get vaccinated. Since these baseline differences could potentially contribute to underestimation of the treatment effects, we control for these variables in the main regressions. To test robustness, we also report the estimates without these controls.

Data. Supplementary Table 5 summarizes the timeline of the experiment. Ultimately, the outcome of interest is whether the participants in the CONSENSUS condition are more likely to actually get vaccinated in the long-term. Therefore, in all waves we asked whether respondents got vaccinated against Covid-19. The variable "Vaccinated" is equal to 1 if the respondent reported having obtained at least one dose of a vaccine against Covid-19. In the Supplementary Information, we also report the results for being partially vs. fully vaccinated, and the results are similar. The fact that the vaccination rate in our sample closely mimics the levels and dynamics of the overall adult vaccination rate in the whole country (Supplementary Figure 1) gives us confidence that the respondents did not misreport their vaccination take-up, perhaps due to desirability biases.

It is important to bear in mind that not everyone was able to get vaccinated from the very beginning of the data collection period. Different demographic groups became eligible to register for the vaccine at different points in time. The Supplementary Online Material provides more details. Thus, in early stages of the data collection the measure "Vaccinated" does not reveal whether most respondents were *willing* to get vaccinated. To gauge more short-term impacts on vaccination demand, we asked the respondents who had not been vaccinated whether they would get an approved Covid-19 vaccine when they became eligible. The variable

"Vaccination demand" is equal to 1 if the respondent reported being already vaccinated or registered for the vaccine, or reported being willing to get vaccinated. To summarize, in Wave0 and Wave1, we study the short-term effects of the CONSENSUS condition on "vaccination demand", because most respondents were not eligible for the vaccine during these waves. In later waves, we are primarily interested in the effects on the ultimate outcome of interest: getting vaccinated.

Prior the treatment in Wave0, we elicited prior beliefs about doctor's views in order to quantify misperceptions about doctors' opinions and to test whether the positive effect of the CONSENSUS condition on vaccination take-up is driven by those who *a priori* underestimate the prevalence of positive attitudes towards the vaccine among medical doctors. Specifically, the participants were asked to estimate (i) the percentage of doctors in the Czech Republic who trust the approved vaccines, and (ii) the percentage of doctors who are either vaccinated or intend to get vaccinated themselves.[1] Further, in Wave1, we elicited posterior beliefs to estimate whether people in the CONSENSUS condition indeed updated their beliefs about doctors' views based on the information provided.

In the analysis, we use detailed information about individual demographic and economic characteristics, and control for a pre-registered set of variables, including gender, age (6 categories), household size, number of children, region (14 regions), town size (7 categories), education (4 categories), economic status (7 categories), household income (11 categories) and prior vaccination intentions. In addition, we also control for respondent's own vaccination in Wave0 and prior beliefs about the views of doctors in Wave0, which turned out not to be perfectly balanced across treatments.

The research study was approved by the Commission for Ethics in Research of the Faculty of Social Sciences of Charles University. The experiment and analyses were pre-registered on the AEA RCT Registry (AEARCTR-0007396).

Results

Misperceptions about doctors' views. To quantify misperceptions about doctors' views on Covid-19 vaccines, we compare prior beliefs about doctors' views measured before the intervention with the actual views from the CMC survey. We find strong evidence of misperceptions. The average, median and modal guess is that 57%, 60% and 50% of doctors, respectively, want to be vaccinated (Figure 2, Panel A), while in reality 90% of doctors do. A vast majority of participants (90%) underestimate the percentage of doctors who want to be vaccinated.

The distribution of beliefs reveals that the large underestimation of doctor's vaccination intentions does *not* originate in two distinct groups of participants holding opposite views of medical consensus -- one thinking that most doctors want to get vaccinated and the other group thinking that most doctors do not want to get vaccinated. Instead, most people expect a large diversity of attitudes among individual doctors. 81% of subjects believe that the percentage of doctors who want to be vaccinated is between 20-80%.

Reassuringly, we arrive at similar conclusions based on our analysis of beliefs about medical doctors' trust in the vaccines (Figure 2, Panel B). The average, median and modal

---

[1] We did not elicit beliefs about the third type of information provided to respondents in the CONSENSUS condition (the willingness of doctors to recommend Covid-19 vaccines to patients), to economize on time, since we expected this type of belief to be highly correlated with the other two about doctors' views (indeed, the pairwise correlation coefficient between Wave0 beliefs about doctors' trust and vaccination intentions is $r(2{,}099)=0.60$, $p<0.01$).

guess about what percentage of doctors trust the vaccines are 61%, 62%, and 50%, respectively, while in practice 89% of doctors report trusting the vaccines. Again, most people (88%) underestimate doctor's trust and do not expect consensus among doctors.

We find several intuitive descriptive patterns that increase confidence in our measures of beliefs. First, beliefs about doctors' vaccination intentions and their trust in the vaccines are strongly positively correlated ($r(2{,}099)=0.60$, $p<0.001$). Second, beliefs about doctor's trust and vaccination intentions are highly predictive of respondents' own intentions and take up (Supplementary Table 5). In the next sub-section, we will explore whether this relationship is causal. Third, in Supplementary Figure 2, we show that misperceptions about doctor's views are unlikely to arise due to participants' inattention to the questions. The results are very similar when we exclude the 5% of subjects who did not pass all of the attention checks embedded in the survey, and when we exclude the 10% of participants with the shortest response time.

Finally, in Supplementary Table 6, we explore which personal characteristics predict beliefs about doctors' trust and vaccination intentions. Overall, we find that misperceptions are widespread across all demographic groups, based on age, gender, education, income, and geographical regions. Higher-income individuals are less likely to expect that doctors will hold negative views of the vaccine, and this relationship does not seem to be due to differences in education levels.

Effects of the intervention. We now estimate the effects of the intervention on (i) posterior beliefs about doctors' views, (ii) demand for getting vaccines shortly after the intervention, and (iii) the long-term dynamics of actual vaccine take-up.

Updating of beliefs. We find that the information provided shifts expectations about doctors' views (Figure 3 and Supplementary Table 7). Two weeks after the intervention (in Wave1), respondents in the CONSENSUS condition expected 72% of doctors to trust the vaccine, while participants in the CONTROL condition expected 67%, the difference being significant statistically (two-sided Wilcoxon rank-sum test, $z(1, N=1{,}940) = -5.92$, $p < 0.001$). We observe similar impacts on beliefs about doctors' intentions to get vaccinated: in the CONSENSUS condition, people expected 69% of doctors to get vaccinated, while in the CONTROL it was 63% (two-sided Wilcoxon rank-sum test, $z(1, N=1{,}940) = -7.48$, $p < 0.001$).

Effects on vaccination demand. In the CONTROL condition, at the outset of the experiment, 64% of participants indicated demand for the vaccine, by reporting that they either planned to get vaccinated or were already vaccinated. This proportion was gradually increasing during the next five months, reaching 77% in August. Figure 4 and Supplementary Tables 8 and 9 document that the CONSENSUS condition positively affects vaccination demand. We find that the information provided increases vaccination demand by 3 percentage points shortly after the treatment (Wave0, linear probability model, $t(2{,}047) = 2.53$, $p = 0.01$, CI= 0.01 - 0.06) as well as when participants are asked two weeks later (Wave1, linear probability model, $t(1{,}886) = 2.28$, $p = 0.02$, CI= 0.00 - 0.06). When restricting the sample to those who participated in all eight waves, we find the effects to be slightly larger: 6 p.p. in Wave0 (linear probability model, $t(1{,}360) = 3.70$, $p <0.01$, CI= 0.03 - 0.08) and 5.p.p. in Wave1 (linear probability model, $t(1{,}360) = 3.05$, $p<0.01$, CI= 0.02 - 0.08).

The effect on vaccination demand is not ephemeral, but lasting. For the fixed sample, we find the effects of the treatment to be around 4 percentage points and statistically significant at conventional levels in all eight waves of data collection. For the base sample, the observed coefficients also suggest remarkable stability of treatment effects over time, although the estimates vary somewhat more across waves and are not statistically significant in each of them.

Effects on vaccination take-up. Ultimately, we are interested whether the treatment increases not only demand but also actual vaccination take-up. We find systematic treatment effects on vaccine take-up (Figure 5 and Supplementary Tables 10 and 11). As expected, due to the gradual roll-out of the vaccine during the data collection period, the effect emerges gradually. The difference in the take-up rates between CONSENSUS and CONTROL conditions is negligible in the initial waves, but then steadily increases to 5 and 4 percentage points in Wave6 and Wave7, respectively (for the full sample, Wave6, linear probability model, $t(1,566) = 2.69$, $p<0.01$, CI= 0.01 - 0.08; Wave7 linear probability model, $t(1,716) = 2.55$, $p<0.01$, CI= 0.01 - 0.08) or 5 to 6 percentage points (for the fixed sample, Wave6, linear probability model, $t(1,360) = 2.81$, $p<0.01$, CI= 0.02 - 0.09; Wave7 linear probability model, $t(1,360) = 3.29$, $p<0.01$, CI= 0.02 - 0.10). During the last two waves, implemented in July and August when all adults were eligible for vaccination, the difference remains relatively stable.

These patterns show that the intervention has lasting impacts on vaccination demand and also suggest that the treatment increased the number of vaccinated individuals, rather than only accelerating vaccination, because we do not see the effect diminishing in the final waves. Further, given that Wave0 vaccination intentions and Wave7 vaccine take up are highly correlated ($r(1,623)=0.53$, $p<0.01$) and, more importantly, that the magnitude of the estimated treatment effects on vaccination intentions shortly after the intervention (4-5 percentage points) is very similar to the estimated magnitude of the effects on actual take-up several months after the intervention, the results suggest that vaccination intentions might be a relatively informative proxy of behavioral impacts, which may somewhat reduce concerns about measurements in studies that rely solely on measuring vaccination intentions.

Additional results. The results are robust (Supplementary Tables 7-11). First, exclusion of participants who arguably paid less attention has a negligible influence on the estimated coefficients. If anything, the results become stronger. Second, the effects on vaccination demand and vaccine take-up are robust to changes in the set of control variables. Third, the effects are statistically significant at conventional levels after adjustment for multiple hypotheses testing. Fourth, while in the main estimates we focus on the likelihood of respondent's getting at least one vaccine dose, qualitatively similar and statistically significant effects emerge when focusing on the likelihood of their completing the whole vaccination cycle of two doses required for most Covid-19 vaccines (Supplementary Figure 3).

Further, we asked respondents whether they would actively recommend vaccination to their friends and relatives. Although this question was not pre-specified as one of our main outcome variables, it may be informative about whether the intervention could potentially have positive spillovers on the decisions of those not directly targeted by the intervention. This does not seem to be the case. Approximately half of the respondents (55%) reported that they would recommend the vaccine to friends, but the treatment has little influence on this outcome (Supplementary Figure 4).

In secondary analyses, we examine how the treatment effects differ across different sub-samples of respondents (Supplementary Tables 7, 9 and 11). In line with the hypothesized mechanism, we find that the positive effects on belief updating, vaccination demand, and vaccine take-up are concentrated among those who underestimated doctors' trust and vaccination intentions, while no systematic effects are observed among over-estimators. In addition, the effects are driven by those who reported that they did not intend to get vaccinated prior to Wave0, in line with the interpretation that the intervention changed the views of individuals who were initially skeptical about the vaccine, rather than helping those already intending to get vaccinated to follow through.

A common challenge in information provision experiments is to rule out the role of priming or the experimenter demand effect in explaining treatment effects[25]. These confounds are unlikely to play a role in our estimates. We observe the treatment effects to be lasting, while these potential confounds should, in principle, only affect responses immediately after reading about the survey of doctors (in Wave0).

Discussion

Our results shed light on the role that misperceptions of the distribution of expert views play in vaccine hesitancy, and also show how this barrier can be lifted by providing accurate information. We provide evidence that (i) the vast majority of Czech medical doctors trust the approved Covid-19 vaccines and intend to get vaccinated themselves, (ii) the vast majority of respondents in a nationally-representative survey substantially underestimate the percentage of doctors with positive views of the vaccine, and (iii) correcting these misperceptions has lasting positive impacts on vaccine take-up, with the effects being driven by persons who did not intend to get vaccinated prior to the intervention.. While existing experiments have made progress in identifying low-cost strategies to increase vaccination intentions[7,11–13] and take-up[8] measured shortly after the intervention, this paper integrates the experiment in longitudinal online data collection and contributes by identifying a low-cost, scalable treatment that has lasting effects on behavior.

Scientists, and the medical community as a whole, have invested enormous efforts to develop and deliver Covid-19 vaccines. However, much less collective effort has been directed at informing the public of the high levels of trust in the vaccine among the broad medical community. Here, we show that professional medical associations might serve as aggregators of individual doctors' views, by facilitating opinion polls among doctors. Resulting data can then be used in campaigns to tackle vaccine hesitancy, and also as an input for media reports. Although we cannot empirically pin down the source of misperceptions observed in our study, we suspect that they originate, at least in part, in a journalistic norm in which balance is often considered a mark of objective and impartial reporting, as well as a way that helps to attract attention of news consumers[27]. Our results strengthen the case for supplementing contrasting views on controversial issues with information about how prevalent such views are[28].

A natural open question is how broadly applicable these findings are beyond the context studied. In theory, this type of intervention should have larger effects (i) the greater is trust in medical doctors in a given country and (ii) the greater is the prevalence of misperceptions about the views of doctors towards a vaccine. We studied this intervention in a country with an approximately median level of trust in doctors[14], which provides some confidence that our findings from the Czech Republic may extend to other settings. At the same time, because this is the first paper that provides direct evidence of the prevalence and size of misperceptions about doctors' views on Covid-19 vaccines, we can only speculate how widespread such misperceptions are in other settings. Given that the likely sources of the misperceptions – false-balance reporting and echo chambers - are not specific to the Czech Republic, and given that misperceptions about scientific consensus have been documented in other countries in other domains, including health and climate change[24,29], we suspect that this bias in beliefs about Covid-19 vaccines is relatively widespread. We hope to see more research on this front.

References

1. Greenwood, B. The contribution of vaccination to global health: Past, present and future. *Philos. Trans. R. Soc. B Biol. Sci.* **369**, 1–9 (2014).
2. Banerjee, A. *et al. Selecting the Most Effective Nudge: Evidence from a Large-Scale*

Experiment on Immunization. *NBER working paper 28726* (2021).
3. Lazarus, J. V. *et al.* A global survey of potential acceptance of a COVID-19 vaccine. *Nat. Med.* **27**, 225–228 (2021).
4. D'Souza, G. & Dowdy, D. What is Herd Immunity and How Can We Achieve It With COVID-19? *John Hopkins Bloomberg School of Public Health* (2021). Available at: https://publichealth.jhu.edu/2021/what-is-herd-immunity-and-how-can-we-achieve-it-with-covid-19. (Accessed: 6th September 2021)
5. McDermott, A. Herd immunity is an important—and often misunderstood—public health phenomenon. *Proc. Natl. Acad. Sci. U. S. A.* **118**, e2107692118 (2021).
6. Jarrett, C. *et al.* Strategies for addressing vaccine hesitancy - A systematic review. *Vaccine* **33**, 4180–4190 (2015).
7. Milkman, K. L. *et al.* A megastudy of text-based nudges encouraging patients to get vaccinated at an upcoming doctor's appointment. *Proc. Natl. Acad. Sci. U. S. A.* **118**, 10–12 (2021).
8. Dai, H. *et al.* Behavioral Nudges Increase COVID-19 Vaccinations. *Nature* **597**, 404–409 (2021).
9. Serra-Garcia, M. & Szech, N. Choice Architecture and Incentives Increase COVID-19 Vaccine Intentions and Test Demand. *SSRN Electron. J.* (2021). doi:10.2139/ssrn.3818182
10. Klüver, H., Hartmann, F., Humphreys, M., Geissler, F. & Giesecke, J. Incentives can spur COVID-19 vaccination uptake. *Proc. Natl. Acad. Sci.* **118**, e2109543118 (2021).
11. Alsan, M. & Eichmeyer, S. *Experimental Evidence on the Effectiveness of Non-Experts for Improving Vaccine Demand*. *NBER working paper 28593* (2021). doi:10.1086/279125
12. Ashworth, M., Thunström, L., Cherry, T. L., Newbold, S. C. & Finnoff, D. C. Emphasize personal health benefits to boost COVID-19 vaccination rates. *Proc. Natl. Acad. Sci.* **118**, e2108225118 (2021).
13. Loomba, S., de Figueiredo, A., Piatek, S. J., de Graaf, K. & Larson, H. J. Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nat. Hum. Behav.* **5**, 337–348 (2021).
14. Blendon, R. J., Benson, J. M. & Hero, J. O. Public Trust in Physicians — U.S. Medicine in International Perspective. *N. Engl. J. Med.* **371**, 1570–1572 (2014).
15. Shapiro, J. M. Special interests and the media: Theory and an application to climate change. *J. Public Econ.* **144**, 91–108 (2016).
16. Boykoff, M. T. & Boykoff, J. M. Balance as bias: Global warming and the US prestige press. *Glob. Environ. Chang.* **14**, 125–136 (2004).
17. Dixon, G. & Clarke, C. The effect of falsely balanced reporting of the autism-vaccine controversy on vaccine safety perceptions and behavioral intentions. *Health Educ. Res.* **28**, 352–359 (2013).
18. Van Der Linden, S. L., Clarke, C. E. & Maibach, E. W. Highlighting consensus among medical scientists increases public support for vaccines: Evidence from a randomized experiment Health behavior, health promotion and society. *BMC Public Health* **15**, 1–5 (2015).
19. Wolrd Health Organization. *COVID-19 An informative guide. Advice for journalists*. *World Health Organization* (2021).
20. Haidt, J. Why social media is terrible for multiethnic democracies. *Vox* (2016).
21. Jennings, W. *et al.* Lack of trust, conspiracy beliefs, and social media use predict COVID-19 vaccine hesitancy. *Vaccines* **9**, 1–14 (2021).
22. Bursztyn, L., Rao, A., Roth, C. & Yanagizawa-Drott, D. *Opinions as Facts*. *Unpublisehd Working Paper* (2021).

23. Bronchetti, E. T., Huffman, D. B. & Magenheim, E. Attention, intentions, and follow-through in preventive health behavior: Field experimental evidence on flu vaccination. *J. Econ. Behav. Organ.* **116**, 270–291 (2015).
24. Bursztyn, L. & Young, D. Misperceptions about Others. *Annu. Rev. Econom.* **forth.**, (2021).
25. Haaland, I., Roth, C. & Wohlfart, J. Designing Information Provision Experiments. *J. Econ. Lit.* (2021). doi:10.2139/ssrn.3638879
26. Bordalo, P., Gennaioli, N. & Shleifer, A. Memory, attention, and choice. *Q. J. Econ.* **April**, 1399–1442 (2020).
27. Dunwoody, S. Scientists, journalists, and the meaning of uncertainty. in *Communicating Uncertainty: Media Coverage of New and Controversial Science.* (eds. Friedman, S. M., Dunwoody, S. & Rogers, C. L.) 59–79 (Lawrence Erlbaum, 1999).
28. Schmid, P., Schwarzer, M. & Betsch, C. Weight-of-evidence strategies to mitigate the influence of messages of science Denialism in public discussions. *J. Cogn.* **3**, 1–17 (2020).
29. Van Linden, S. L. D., Leiserowitz, A. A., Feinberg, G. D. & Maibach, E. W. The scientific consensus on climate change as a gateway belief: Experimental evidence. *PLoS One* **10**, 2–9 (2015).

Data availability. The dataset generated and analyzed for the main experiment, together with replication files, is available in the Harvard Dataverse repository (https://doi.org/10.7910/DVN/RH0T6R). The availability of the dataset from the Supplementary survey with medical doctors is subject to approval of the Czech Medical Chamber.

Acknowledgments. We thank the Czech Medical Chamber, and esp. Milan Kubek, Marek Matoušek, Radek Mounajjed, Jakub Studený and Daniel Valášek for implementing the Supplementary survey of medical doctors and providing further statistics, and PAQ Research (esp. Daniel Prokop) and NMS Market Research for the data collection used in our main experiment. VB thanks the German Research Foundation (CRC TRR 190 and 444754857). MB and JCh thank the Czech Science Foundation (20-11091S). The funders had no role in the design of the study, data collection and analysis, decision to publish, or preparation of the manuscript. Darya Korlyakova provided excellent research assistance.

Authors contributions. All authors contributed equally.

Competing interests. We declare no competing interests.

Additional information. Supplementary Information is available for this paper. It contains supplementary figures and tables and provides more background information about the Covid-19 pandemic and vaccination in the Czech Republic.

The research study was approved by the Commission for Ethics in Research of the Faculty of Social Sciences of Charles University. Participation was voluntary and all respondents provided their consent to participate in the survey.

Figures

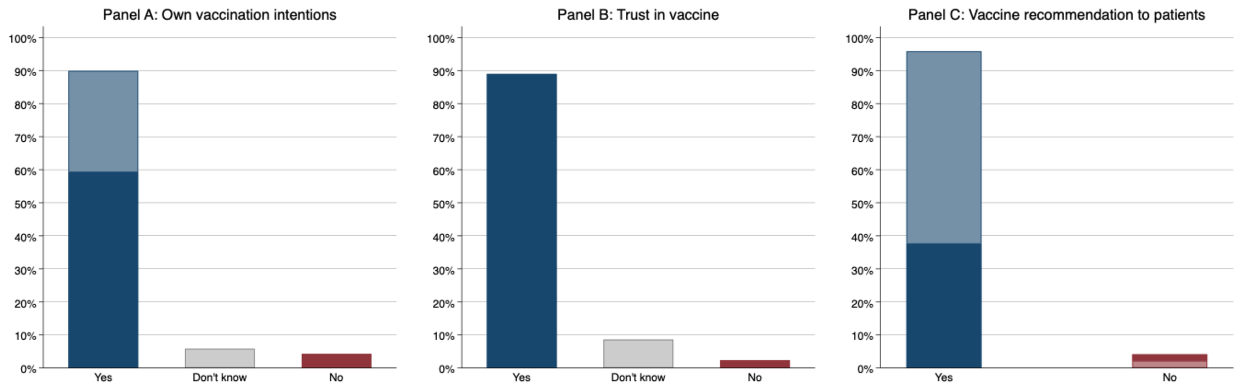

Figure 1. Doctors' views on Covid-19 vaccines (Supplementary Study, N=9,650). In **a**, we report the distribution of responses to the question "Will you personally be interested in getting vaccinated, voluntarily and free of charge, with an approved vaccine against Covid-19?". Among participants who answered yes, the dark (light) blue refers to those who reported already being vaccinated (plan to get vaccinated). In **b**, we report responses to the question "Do you trust Covid-19 vaccines that have been approved by the European Medicines Agency (EMA) approval process?". In **c**, we report responses to the question "Will you recommend Covid-19 vaccination to your healthy patients to whom you would recommend other commonly-used vaccines?" Among participants who answered yes, the dark (light) blue refers to those who will recommend the vaccines even without being asked (only when asked). In Supplementary Table 2, we show that the distribution of views is similar across various demographic groups and level of seniority.

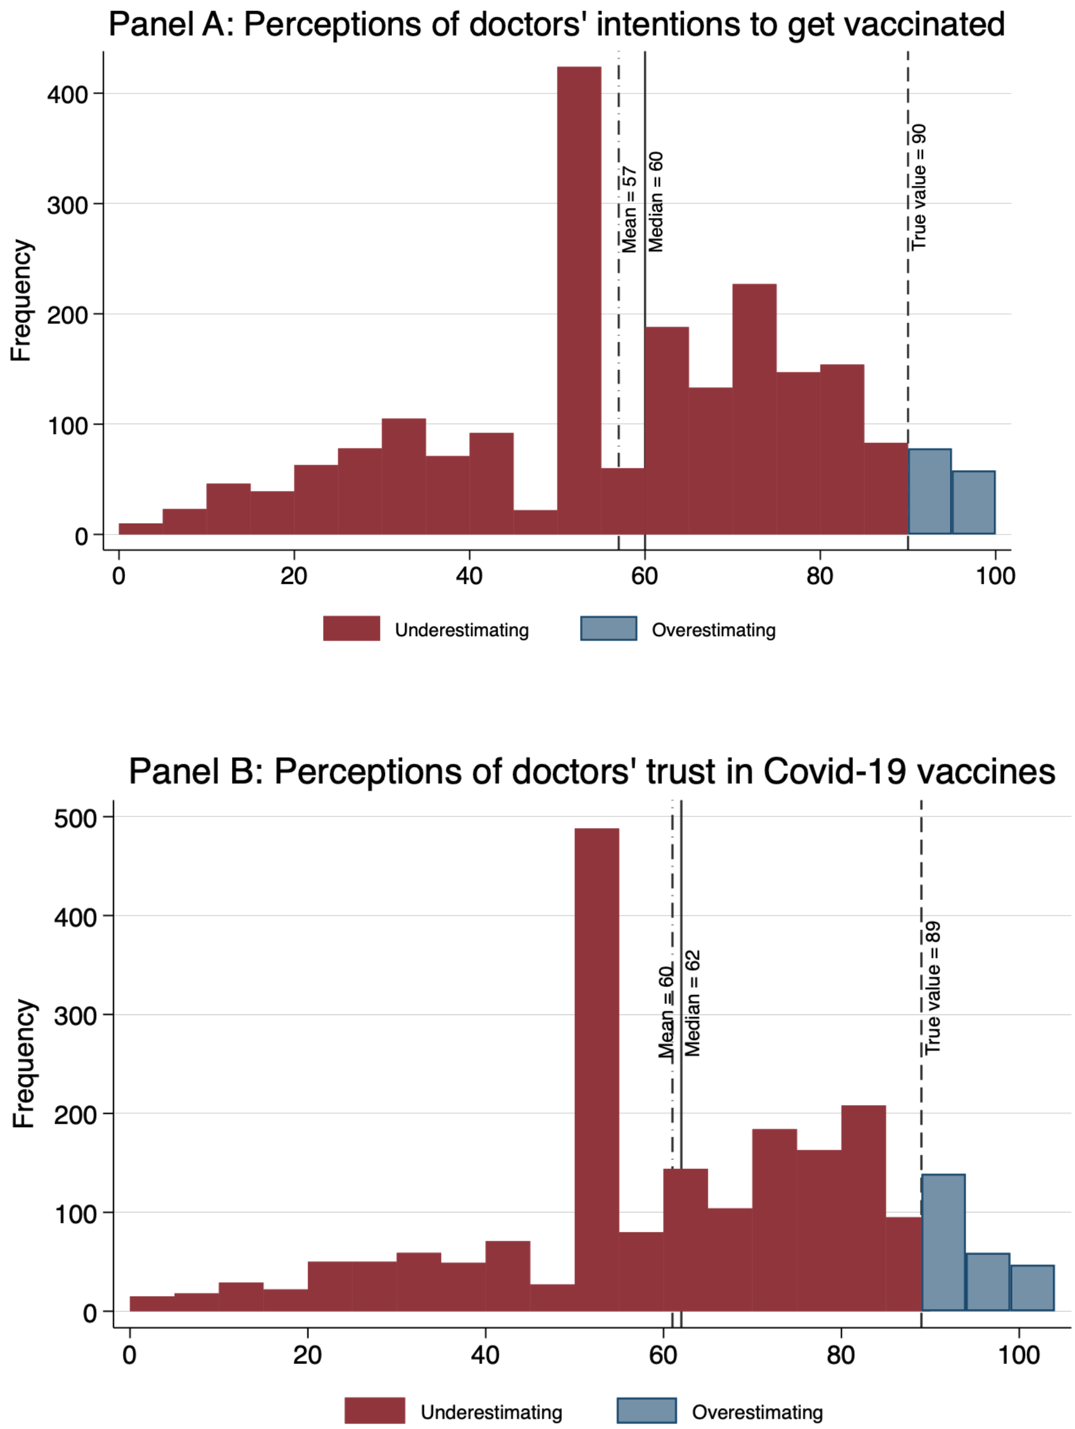

Figure 2. Perceptions of doctors' views on Covid-19 vaccines (Main Experiment, N=2,101). In **a**, we report the distributions of respondents' prior beliefs about what percentage of doctors would like to get vaccinated. In **b**, we report the distributions of respondents' beliefs about what percentage of doctors trust approved Covid-19 vaccines. The dashed line shows the true value, based on responses of doctors in the Supplementary study. The red (blue) color illuminates the percentage of those who underestimate (overestimate) doctors' own vaccination intentions (panel a) and trust in the Covid-19 vaccines (panel b).

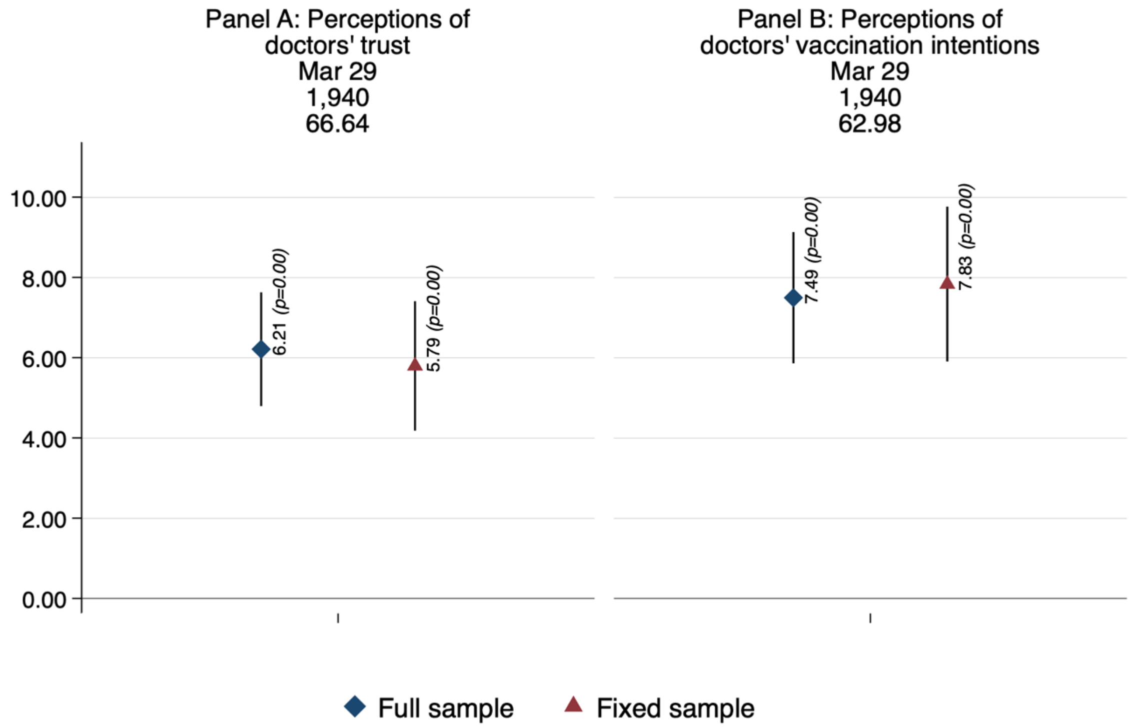

Figure 3. Effects of the CONSENSUS condition on posterior beliefs about doctors' views (Main Experiment). This figure plots estimated treatment effects on beliefs about what percentage of medical doctors plan to get vaccinated (Panel A) and on beliefs about what percentage of doctors trust approved Covid-19 vaccines (Panel B), using linear probability regression. The posterior beliefs were measured only in Wave1 (March 29). The whiskers denote the 95%-confidence interval based on Huber-White robust standard errors. We control for the following set of control variables: gender, age category (6 categories), household size, number of children, region (14 regions), town size (7 categories), education (4 categories), economic status (7 categories), household income (11 categories), and baseline vaccination intentions, take-up, and beliefs about the views of doctors. We report estimates for (i) all observations (diamond) and (ii) for a sub-sample of participants who took part in all eight waves (triangle). In the upper part of the Figure, we report the timing, the total number of observations, and control mean for each wave. Supplementary Table 7 shows the regression results in detail.

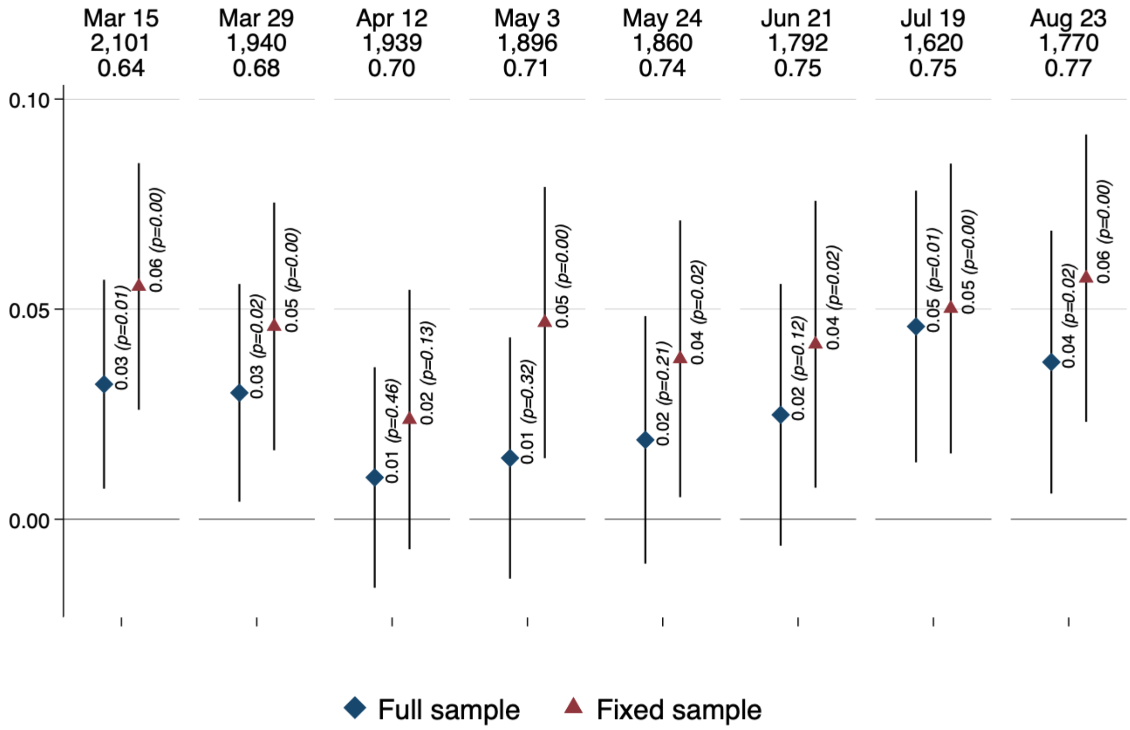

Figure 4. Effects of the CONSENSUS condition on vaccination demand (Main Experiment). This figure plots estimated treatment effects by survey wave on the likelihood of intending to or of being vaccinated, using linear probability regression. The whiskers denote the 95%-confidence interval based on Huber-White robust standard errors. We control for the same set of control variables as in Fig. 3. In the upper part of the Figure, we report the timing, the total number of observations, and control mean for each wave. For each wave, we report estimates for (i) all observations (diamond) and (ii) for a sub-sample of participants who took part in all eight waves (triangle). Supplementary Table 8 shows the regression results in detail.

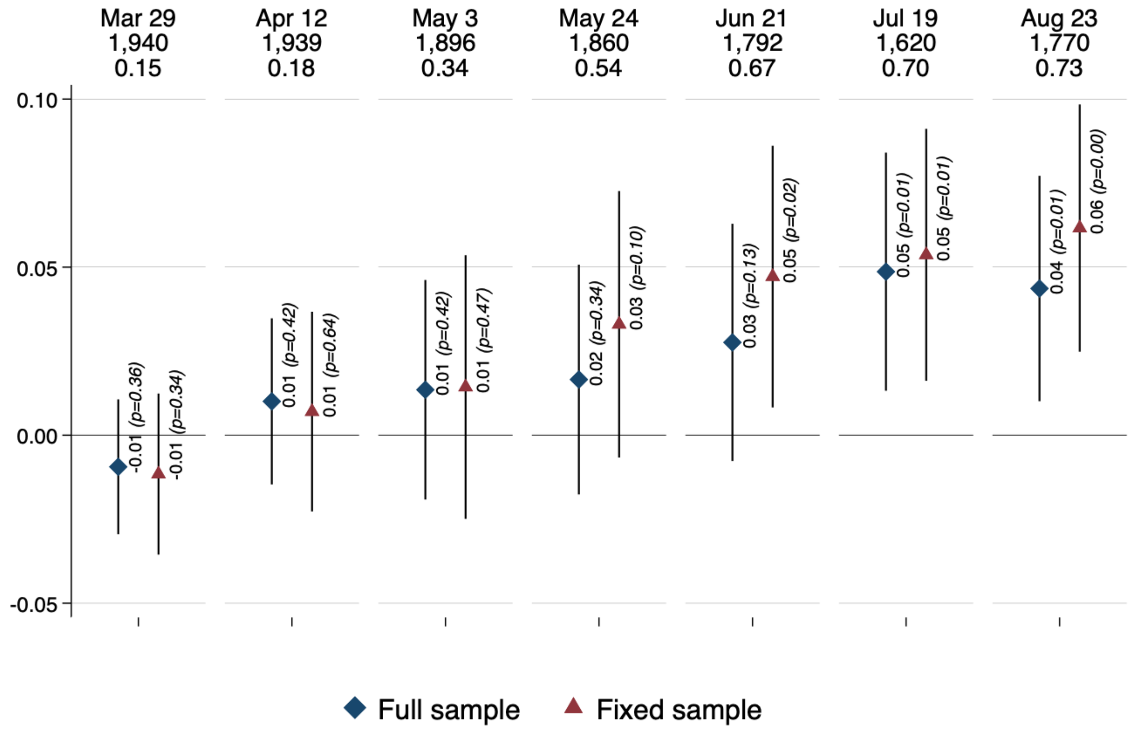

Figure 5. Effects of the CONSENSUS condition on vaccine take-up (Main Experiment). This figure plots estimated treatment effects by survey wave on the likelihood of vaccine take-up, using linear probability regression. The whiskers denote the 95%-confidence interval based on Huber-White robust standard errors. We control for the same set of control variables as in Fig. 3. In the upper part of the Figure, we report the timing, the total number of observations, and control mean for each wave. For each wave, we report estimates for (i) all observations (diamond) and (ii) for a sub-sample of participants who took part in all eight waves (triangle). Supplementary Table 10 shows the regression results in detail.