



HAL
open science

The causal effect of a health treatment on beliefs, stated preferences and memories

Alberto Prati, Charlotte Saucet

► **To cite this version:**

Alberto Prati, Charlotte Saucet. The causal effect of a health treatment on beliefs, stated preferences and memories. *Journal of Health Economics*, 2024, 94, pp.102864. 10.1016/j.jhealeco.2024.102864 . hal-04512861

HAL Id: hal-04512861

<https://hal.science/hal-04512861v1>

Submitted on 27 Mar 2024

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.

The causal effect of a health treatment on beliefs, stated preferences and memories

Alberto Prati^{a,*}, Charlotte Saucet^{b,*}

^a*University College London, University of Oxford, and London School of Economics, United Kingdom.*

^b*University of Paris 1 Panthéon-Sorbonne, France.*

Abstract

The paper estimates the causal effect of a health treatment on patients' beliefs, preferences and memories about the treatment. It exploits a natural experiment which occurred in the United Kingdom during the COVID-19 vaccination campaign. UK residents could choose to opt into the vaccination program, but not which vaccine they received. The assignment to a vaccine offered little objective information for learning about its qualities, but triggered strong psychological demand for reassuring beliefs. We surveyed a sample of UK residents about their beliefs on the different COVID-19 vaccines *before* and *after* receiving their jab. *Before* vaccination, individuals exhibit similar prior beliefs and stated preferences about the different vaccines. *After* vaccination, however, they update their beliefs overly optimistically about the safety and effectiveness of the vaccine they received, state that they would have chosen it if they could, and have distorted memories about their past beliefs. These results cannot be explained by conventional experience effects. At the aggregated level, they show that random assignment to a health treatment predicts a polarization of opinions about its quality. At the individual level, these findings provide evidence in line with the predictions of motivated beliefs and over-inference from weak signals in a real-world health setting.

Keywords: Natural experiment | Behavioral health economics | COVID-19 | Motivated beliefs | Motivated memory | Over-inference

JEL: I12, I18, D91

*E-mails : a.prati@ucl.ac.uk; charlotte.saucet@univ-paris1. Addresses: University College London, CNET Office 321, 16 Taviton St, London WC1H 0BW, United Kingdom; Maison des Sciences Économiques - Université Paris 1 Panthéon-Sorbonne, Bureau 410, 112 Bd de l'Hôpital, 75013 Paris. We would like to thank the editor Owen O'Donnell and two anonymous referees for their constructive comments, as well as Roland Bénabou, Jeanne Hagenbach, Nicolas Jacquemet, Shengwu Li, Armando N. Meier, Isaac Parkes, Alice Soldà, Rémi Suchon and Michael Thaler. We are grateful to the seminar participants at the Behavior Seminar (Paris School of Economics), Dep. of Economics External Seminar (University of Malaga), Dep. of Economics Postdoctoral Seminar (University of Oxford), "Motivated Reading of Information" Workshop (Sciences Po Paris), BEER Seminar (GATE Lyon) and the Affective Brain Lab seminar (UCL). This research has been funded by the *Fondation Nationale des Sciences Politiques* and the European Research Council, grant 850996, MOREV (Motivated Reading of Evidence) project.

1. Introduction

Beliefs are powerful predictors of decisions. This is true also in the health sector, where patients' beliefs about a health treatment predict their subsequent decisions to accept or decline that treatment. Evidence abounds across a wide range of treatments, from Type II diabetes to HIV, from depression to asthma (see Horne et al., 2013, for a meta-analytic review). Overall, beliefs play a bigger role in the adherence to a treatment than demographic and clinical factors (Horne and Weinman, 1999), and even unintentional non-adherence (due to forgetting or carelessness) can be predicted by patients' beliefs about how safe and effective their treatment is (Gadkari and McHorney, 2012).

Little is known, however, about the converse: how does a health treatment influence subsequent beliefs? Investigating this question is not trivial. When confronted with important health decisions, people typically have agency and choose what they believe to be the best option for them. Undergoing a specific treatment is thus an endogenous choice, with self-selection problems that prevent causal identification. To make progress, we exploit a natural experiment which occurred in the UK during the COVID-19 pandemic, where individuals were *exogenously* assigned either the Pfizer or the Moderna vaccine. We use this random variation to study the causal effect of getting a vaccine on patients' beliefs, preferences and memories about that vaccine.

We find that, after vaccination, individuals exhibit optimistic beliefs about the characteristics of their own vaccine, compared to those who received another vaccine. When asked to make hypothetical choices, many change them in favor of their own vaccine.¹ When asked about what they used to think before receiving the vaccine, they tend to distort their memories in favor of the vaccine they were assigned to. These effects can be causally attributed to the injection of a specific vaccine. At the aggregated level, this asymmetric updating (each group updates beliefs about its own vaccine in an overly optimistic way) generates a dynamic of belief divergence and thereby shows that beliefs polarization can emerge as a consequence of exogenous assignment of individuals to different treatments.

This study was made possible by a particular institutional setting in the UK during the COVID-19 vaccination campaign. UK residents could choose to opt into the vaccination program, but not which vaccine they received. Once age and health restrictions were taken into account, the assignment to a specific vaccine was exogenous. Although the majority of people declared that they would choose a type of vaccine (Pfizer) if they had the choice, about 20% of them were assigned to another one (Moderna). This event was likely to trigger strong psychological needs, that pointed in different directions according to the received vaccine. We interviewed 640 UK residents about their beliefs on the different COVID-19 vaccines *before* and *after* receiving their jab, and study how people who received Moderna (treatment group) update their beliefs compared to those who received Pfizer (control group). We focus on the safety and effectiveness of the Moderna vaccine as well its desirability perceived ex-ante and ex-post.

We document three main results. First, we observe that individuals are overly optimistic in updating their beliefs about the vaccine they received. *Before* vaccination, individuals hold similar prior beliefs about the different vaccines. *After* vaccination, however, those who received Moderna believe it to be safer and more effective than they used to think, and more so than those who received Pfizer. Such asymmetric updating is observed also when controlling for personal experience and socio-demographic characteristics. It is not explained by some peculiarities of the Moderna vaccine, as beliefs about Pfizer's safety display a similar (reversed) pattern,

¹The questions were "If you had the choice, which vaccine would you choose to receive?" (Wave 1) and "If you have had the choice, which vaccine would you have chosen to receive?" (Wave 2). Although economists typically favor data derived from incentivized choices, our analysis relies on stated preferences elicited from hypothetical scenarios. This represents a limitation of our approach, as stated and actual choices might differ, but it was the only implementable option in this context. See Benjamin et al. (2012) for a honest discussion of the limits and potentials of hypothetical choices in surveys.

and no difference emerges for beliefs about vaccines that were not assigned in our sample, i.e. AstraZeneca and Janssen. Second, we show that this observed asymmetric updating is driven by individuals who were particularly skeptical about a vaccine but then were assigned to it, and by individuals who were particularly keen on a vaccine but then did not receive it. This is more than an instance of simple regression-to-the-mean: even among people exhibiting low (resp. high) priors, we observe a significant asymmetry between those who received different vaccines. Third, we find that people adapt not only their beliefs, but also their hypothetical choices and the memory of their past beliefs. Before vaccination, only 7% of people declare that they would choose Moderna. Later on, about 25% of those who received Moderna state that they would have chosen this vaccine if they have had the choice. Moreover, after vaccination, those who received Moderna recall it to be safer and more effective than those who received Pfizer.

These results meet the predictions of Caplin and Leahy (2019)’s model, according to which individuals’ desires causally impact their beliefs, and thus a divergence in desires will result in a divergence in beliefs. In sum, people seem to believe what they would like to be true. This phenomenon, known as *motivated beliefs* (Kunda, 1990; Bénabou, 2015; Zimmermann, 2020; Möbius et al., 2022) or *wishful thinking* (Mayraz, 2011; Engelmann et al., 2019; Caplin and Leahy, 2019; Mayraz, 2019) appears whenever individuals update their beliefs in a fashion that is overly favorable for themselves given the available information.² Identifying motivated beliefs in the field –and particularly in the health domain– is far from simple. First, because the desire for good health is innate and unchanging, it is almost impossible to observe a change in desires, and thus a change in beliefs. As a consequence, one can often only observe the resulting beliefs.³ Second, because individuals decide according to their preferences, they have a priori no interest in changing their beliefs or memories. In our survey, *before* vaccination, individuals’ desire for good health was likely not geared toward a specific vaccine. *After* vaccination however, it became vaccine-dependent. This exogenous event motivates beliefs in different directions, and offers a rare opportunity to test the predictions of motivated beliefs in the real world. We focus our analysis on Moderna because, among those who received it, only a minority (7%) would have chosen it ex-ante (we report the same analyses for Pfizer in the Supplementary Material, “SM” hereafter). To the best of our knowledge, this setting provides the closest replication of a treatment-control test of motivated beliefs in a real-world health setting.

Yet, there is actually another potential explanation for our findings: individuals may update their beliefs based on personal experience with their vaccine that provides little information in most cases. The ambiguity in the interpretation of the underlying mechanism comes in part from the fact that beliefs and motivations are fundamentally unobservable, and in part from the main limitation of this study, i.e., that individuals in the treatment and control group arguably did not receive *exactly* the same information. In the weeks after being vaccinated, participants might have actively looked for different information based on the vaccine that they received. Furthermore, they had some personal experience about the safety and effectiveness of one vaccine only. We show that individuals do not display enhanced knowledge about the benefits of their own vaccine, nor do they reveal a preference for reading about those benefits in an information-selection task. We also show that mere Bayesian updating could not explain the size of the update. Since the mean prior subjective belief of experiencing severe side effects was 0.00001%, experiencing no severe side effects offers little objective information for Bayesian updating. Nevertheless, we cannot rule out the possibility of some forms of purely

²That is, they are deluded in the perception of the situation and therefore evaluate their outcome as more favorable than an impartial observer would do. This process is akin to “desirability bias” (Krizan and Windschitl, 2009), “desirability effect” (Bar-Hillel and Budescu, 1995), the “good news-bad news effect” (Eil and Rao, 2011) and “rationalisation” (Elster, 1983). It encompasses “unrealistic optimism” (Weinstein, 1980) and “optimism bias” (Sharot et al., 2011), which refer to the wishful prediction of future events.

³See, for example, Weinstein et al. (2005); Jansen et al. (2011); Hanoch et al. (2019); Brnstrm and Brandberg (2010).

cognitive (i.e., stripped of motivational factors) non-Bayesian update from their personal experience. This interpretation is motivated by recent evidence that people *over*-infer from weak signals (Augenblick et al., 2021). Over-inference might be enhanced by other contextual factors, such as the first-hand nature of information (Conlon et al., 2022), and the salience of the vaccination experience itself (Gennaioli and Shleifer, 2010; Bordalo et al., 2018, 2020).

This study belongs to the emerging field of behavioral health economics (Cox et al., 2016; Galizzi and Wiesen, 2018). The results reported herein contribute to the identification of behavioural factors that impact the demand-side of health treatments. They do so in two ways.

First, they clarify the dynamic of beliefs formation about health treatments in general, and COVID-19 vaccines in particular. Two important differences are worth mentioning. With respect to medical choices, we differ from the literature that investigated how beliefs impact the willingness to get tested, cured, or vaccinated (Oster et al., 2013; Ganguly and Tasoff, 2017). Our study explores the *reversed* causality. We look at how the effect of getting vaccinated in turn impacts beliefs about vaccines. With respect to the burgeoning literature on beliefs about COVID-19 vaccines, we differ from the studies that identified the determinants of vaccination hesitancy (Jamieson et al., 2021; Kaplan and Milstein, 2021; Mahmud et al., 2021) and that compared various types of interventions to reduce it (Ashworth et al., 2021; Campos-Mercade et al., 2021; Schneider et al., 2023). Conversely, we look at the dynamic of beliefs as a *consequence* of COVID-19 vaccination. This was made possible by considering differences in individuals' beliefs about each specific vaccine, rather than studying beliefs toward COVID-19 vaccinations in general (all vaccines confounded).

Second, our results provide evidence in line with the predictions of motivated beliefs in a *real-world health setting*. When the answer to a medical question has different comforting properties to different people, beliefs are updated in the direction of the most reassuring option. Motivated beliefs have been documented as a potential cause of belief divergence either in artificial situations (Schwardmann et al., 2022) or for individuals with different ex-ante beliefs, e.g. republican and democrats (Kahan, 2012; Levin et al., 2023). A few natural experiments documented that different beliefs tend to *converge* after a public event that affects everyone, e.g. an election (Beasley and Joslyn, 2001; Mullainathan and Washington, 2009). In contrast, we study a rare setting where people with *similar* beliefs are privately and randomly exposed to two *different* treatments, whose desirability depends on *public* (scientific) information. In this context, predictions of confirmation bias (Lord et al., 1979) and attribution theory (Bradley, 1978) do not apply, since individuals share similar priors and know that the vaccination outcome is outside of their control.⁴

There are very few studies that can identify a psychological mechanism by finding the data to be inconsistent with all other potential explanations, and ours is no exception. As already mentioned, we explore the role of motivated beliefs since the vaccination plausibly triggered strong psychological demand for reassuring beliefs, but our results are also consistent with the hypothesis that individuals simply learnt too much from their positive vaccination experience.⁵ Further research will be needed to pin down the psychological micro-foundation of the pattern that we document. Yet, the resulting beliefs are clear. We observe that when different people are assigned to different health treatments, beliefs tend to diverge. This divergence can be predicted, with corresponding implications for public policies, such as informing the debate on whether people should be able to choose the type of their vaccine during an epidemic.

⁴These results were not trivial ex-ante, as motivated beliefs may seem unwarranted when individuals cannot choose their outcome and therefore bear no responsibility for what happens to them. Instead, theories of cognitive dissonance (Festinger, 1962; Suzuki, 2019) and of rationalization of past decisions (Eyster et al., 2021) relate to actual choices. For instance, Suzuki (2019, p.25) argues that “when there is no choice, the decision-maker has no room to experience post-decision dissonance as her choice cannot be ‘wrong’”.

⁵99% of the sample reported that they experienced no severe side effects.

We organize the paper as follows. In Section 2 we describe the natural experiment and the fundamental features of our longitudinal survey; we also spell out the identifying assumptions and test them. In Section 3, we describe the main results about beliefs updates, memories and states preferences. Section 4 discusses the potential mechanisms underlying the results and Section 5 concludes.

2. The natural experiment

2.1. Timeline

We surveyed a panel of individuals on two occasions.

In June-July 2021 (Wave 1), via Prolific (Palan and Schitter, 2018), we interviewed 1,285 UK residents who had not been vaccinated against COVID-19. In the UK, the vaccination campaign had different rules based on age. We interviewed only people between 18 and 29 years old because they became eligible for vaccination in June 2021 (thus, mitigating self-selection problems), and because they were eligible for two vaccines only, Pfizer-BioNtech or Moderna (thus, reducing the dimensionality problem to one pairwise comparison). Participants had to answer a survey made of questions on their beliefs and preferences toward the different vaccines (see 2.4 *Measurements*).

In November-December 2021 (Wave 2), by which time individuals in sample age group had been given the possibility of receiving COVID-19 vaccination, we asked the same participants to complete a second survey that consisted of the same questions and additional questions about their memory (see 2.4 *Measurements*). Thus, if we compare beliefs elicited in Wave 1 and in Wave 2 we can identify individuals who update their beliefs differently because they were assigned to different vaccines. Figure 1 displays the timeline of the natural experiment.

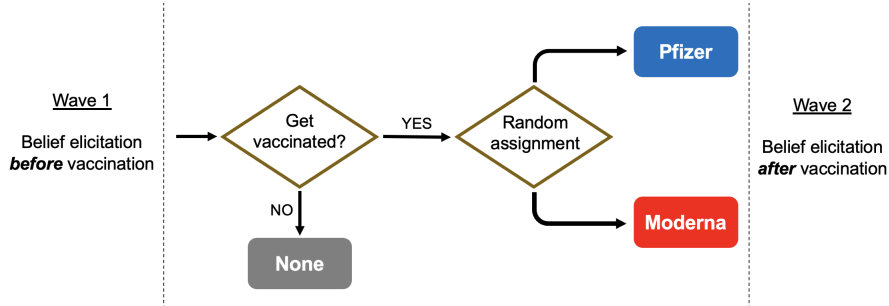


Figure 1: Timeline of the natural experiment

Reading note: UK residents could choose to get vaccinated or not, but not which vaccine they received. Pfizer and Moderna were randomly allocated. In June-July 2021 (Wave 1), we elicited beliefs about COVID-19 vaccines based on the pre-screening criterion that individuals should not have received a vaccine yet. In November-December 2021 (Wave 2), we asked the same participants to complete a second survey eliciting the same beliefs. Comparing beliefs elicited in Wave 1 and in Wave 2 allows identifying belief updating.

2.2. Identifying a causal effect of vaccine assignment on beliefs updating

Unraveling a causal link between assignment to a specific vaccine and subsequent belief updating crucially depends on a *random* assignment of vaccines. “Random” in the sense that the assignment to a vaccine was uncorrelated with observable characteristics, ex-ante beliefs and stated preferences. In the UK, vaccines were distributed subject to nationally and locally determined allocation principles, and to ensure equity in access (NHS England, 2021). No information was publicly disclosed regarding how many doses of which vaccine were delivered to a vaccination center. The same vaccination center could use different vaccines on the same day, and,

at the moment of booking an appointment, users had no information on which vaccine they would receive. On March 12th 2021, Professor Jonathan Van-Tam, Deputy Chief Medical Officer for England, made the following statement: “COVID vaccines are made at different speeds by different manufacturers. It is not possible for vaccination centers to choose the stock they are allocated and not possible for individuals to choose a vaccine.”⁶

At the beginning of Wave 2, we asked participants whether or not they could choose the vaccine they received. There was no reason to lie since participants could complete the survey and receive payment regardless of their answer to this question. Overall, we drop from the analysis 65 participants who reported they could choose their vaccine (see SM, section A1 and Table S6 for details). In the UK, patients found out the type of their vaccine a few minutes *before* (and not after) the injection: in theory, they could refuse to complete the procedure, book another appointment, and try their chances again. In the survey, some participants took the time to explain why they could choose their vaccine, but none of them reported that they refused a vaccination. Among them, 12 could do so for medical reasons, while 7 went to a vaccination center where they expected to deliver only the vaccine that they wanted (leveraging information from peers or other unofficial sources). We suspect that not all the remaining 46 participants could actually choose their vaccine, and that such a high number was due to a misunderstanding of our questionnaire. Indeed, several respondents reported that they could “choose” in a certain sense, just because they were not eligible for all vaccines. For instance, participant n.954 wrote: “I was young for the Astrazeneca so that was partly a choice as I could refuse it but other than that I didn’t get to choose”. Regrettably, this was not the intended meaning of our question. Since most participants did not leave any further comment, we might have been unable to identify many of these cases.

As we will show in the Results section, we observe no difference in the ex-ante beliefs between the PFIZER group and the MODERNA group. Although individuals’ beliefs measured in Wave 1 could be slightly different from those on the day of the assignment to the treatment, the time lapse was reasonably short: over two-thirds of the vaccinated sample report that they got their first jab within five weeks from their participation in Wave 1. In fact, the UK vaccination campaign was rolled out by age groups, and we fielded Wave 1 right after the under-30-year-old group became eligible for vaccination. Based on IP addresses, we also elicit longitude and latitude of every respondent and find no detectable geographical pattern in the assignment to a vaccine (Moran’s I is statistically null).

2.3. Final sample

Details of the study and procedures can be found in the SM, section A1. All of the participants gave their informed consent at the beginning of the study. The Research Ethics Committee of the Paris Institute of Political Studies reviewed and unanimously approved the procedures (n.2021-023). Subjects were recruited through Prolific and were informed that they may be recontacted to complete a second questionnaire a few months later. 1,285 individuals (458 males; mean age = 24.4) completed Wave 1; 951 completed Wave 2 (310 males; mean age = 23.4), and 856 (282 males, mean age = 23.3) passed the exclusions restrictions detailed in the SM, section A1. The statistical power analysis is described in the SM, section A8. Participants received £0.73 and £1.2 for their participation in Waves 1 and 2, respectively. These payments correspond to hourly rates of about £11 and are well above the average payment offered on Prolific.

Our final sample is made of 856 participants, which is relatively small but sufficient to detect a small-medium effect (see SM, section A8). 306 received no vaccine between Wave 1 and Wave 2. We refer to them as the group NONE. 457 received the vaccine Pfizer-BioNtech (“Pfizer” henceforth) and 93 received the vaccine Moderna.

⁶Source: <https://healthmedia.blog.gov.uk/2021/03/12/covid19vaccines-faqs/>.

We refer to them as the groups PFIZER and MODERNA, respectively.⁷ The NONE-PFIZER-MODERNA shares in our sample mirror the vaccination roll-out statistics in the UK population in the Fall 2021 (Rough and Powell, 2021).

In the analyses, we study how people who received MODERNA (treatment group) vs. PFIZER (control group) update their beliefs differently regarding Moderna. We focus on beliefs about Moderna because, ex-ante, only a minority would choose it, relative to Pfizer. The SM (section A2, Figures S1-S5, Table S8) reports the same analyses regarding beliefs about Pfizer’s vaccine. Note that NONE is neither a treatment nor a control since people self-selected into this group. We display the beliefs update of unvaccinated individuals with a descriptive purpose, and not for causal identification.

2.4. Measurements

In Wave 1 of the natural experiment, participants had to answer a survey made of 3 blocks of questions. In the first block, they had to state their beliefs regarding the safety and effectiveness of each of the four vaccines that were available at that time in the UK (Pfizer-BioNtech, Moderna, Oxford-AstraZeneca, Janssen). Beliefs were elicited using cardinal probability scales and ordinal relative ranking.

Cardinal beliefs measures. Vaccines’ safety was measured on a logarithmic scale, where people reported the incidence of significant side effects from 1 out of 10 (very unsafe) to 1 out of 10^7 (very safe). Options were chosen on a labelled 1-7 Likert scale. According to scientific evidence at the time of the survey, the correct answer was around 4 for both vaccines (Baden et al., 2021; Polack et al., 2020). Vaccines’ effectiveness was measured on a percentage scale, where people reported how effective they consider each vaccine, from 0% (completely ineffective against severe COVID-19 and death), to 100% (completely effective against severe COVID-19 and death). According to scientific evidence, the correct answer was around 95% for both vaccines (Bruxvoort et al., 2022; Tartof et al., 2021).

Ordinal beliefs measures. Respondents were asked to rank the four vaccines from the safest (1) to the least safe (4), and from the most effective (1) to the least effective (4). These questions enable to elicit which vaccine was considered *better*, independently of how *good* it was considered.

In the second block, participants were asked to report which vaccine they would choose to receive if they had the choice, as well as their estimated likelihood to receive each vaccine. Finally, the third block asked for demographic variables (age, gender, area of residence, health conditions, etc.).

In Wave 2, we asked the same participants to complete the same blocks of questions. We additionally asked participants: (i) to remember what they thought at the time of Wave 1; (ii) to report some details of their experience with COVID-19 and with the vaccine; (iii) to choose a piece of scientific information about vaccines that they wanted to read; and (iv) to take a short quiz designed to assess their general knowledge about COVID-19 vaccines. Table S1 in SM summarises the main variables that we measured in the two waves, and Tables S2-S5 in SM report the complete questionnaires.

3. Results

We analyse how people who received different vaccines change their beliefs, stated preferences and memories. We use the exogenous allocation of the Pfizer and Moderna vaccines to identify causal effects of the type of vaccine received on beliefs, preferences and memories. We will focus on belief updating first.

⁷Throughout the article, we will use small caps (PFIZER, MODERNA) to refer to the vaccine that individuals received, and standard typeface (Pfizer, Moderna) to refer to the vaccine that individuals expressed their beliefs, choices and memories about. Therefore, “PFIZER’s beliefs about Moderna” should be understood as the beliefs about Moderna expressed by the group of individuals who received (or will receive) Pfizer.

3.1. Before vaccination, individuals exhibit similar beliefs about Moderna vaccine.

In Wave 1, individuals do not know which vaccine they will be allocated to. Ex-ante, there is thus no reason to observe different beliefs between the PFIZER group and the MODERNA group. Importantly, this does *not* mean that the two vaccines should be considered equally good by the participants. It means that ex-ante beliefs about how good each vaccine is should not differ between the PFIZER group and the MODERNA group.

Figure 2 *left* displays the average reported beliefs about Moderna’s safety. Before receiving a vaccine, the average belief of individuals regarding Moderna is not significantly different between individuals that *will* receive PFIZER (5.27) and individuals that *will* receive MODERNA (5.30, $p = 0.851$, Student’s t-test). Similarly, there is no significant difference between the two groups regarding the average reported beliefs about the effectiveness of Moderna (see Figure 2 *right*, $Mean_{PF1} = 75.3$, $Mean_{MOD} = 77.1$, $p = 0.365$, Student’s t-test). In both graphs, the group NONE reports much lower levels of perceived safety and effectiveness. This happens because individuals decide to get vaccinated or not, and those with lower confidence in the safety and/or effectiveness of (any type of) vaccination treatment in Wave 1 self-select out of the treatment itself in the following months.

These first results show that the groups PFIZER and MODERNA hold ex-ante identical beliefs about the characteristics of the Moderna vaccine, whether it be safety or effectiveness.

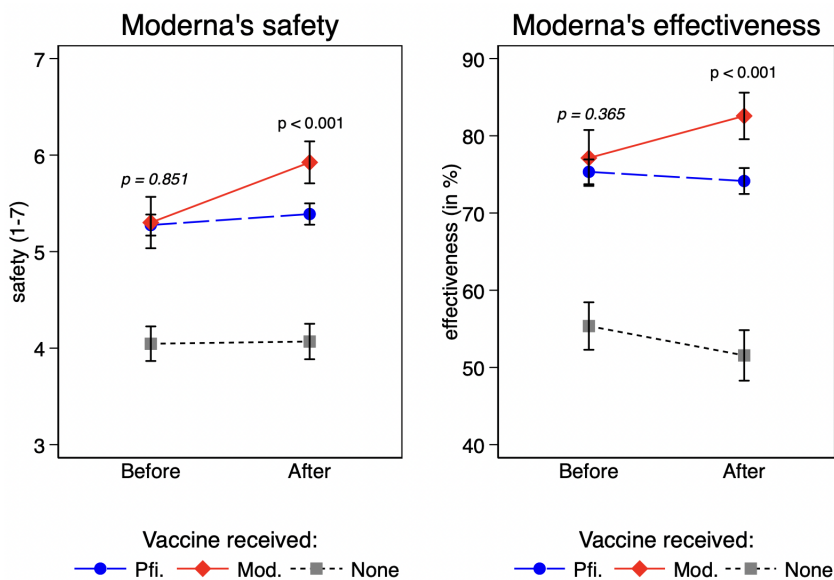


Figure 2: Beliefs about Moderna’s vaccine

Reading note: ‘Before’ corresponds to belief elicitation that happened in June-July 2021, i.e. before any individuals received their first vaccination. ‘After’ corresponds to belief elicitation that happened in November-December 2021, i.e. after some individuals received their vaccination. Safety is measured on a labelled logarithmic scale, from 1 (incidence of significant side effects = 1 out of 10) to 7 (incidence of significant side effects = 1 out of 10^7). Effectiveness is measured on a percentage scale, from 0% (completely ineffective against severe COVID-19 and death), to 100% (completely effective against severe COVID-19 and death). P-values come from Student’s t-tests. Error bars refer to 95% confidence intervals. $N = 856$, data from Waves 1 and 2, groups MODERNA, PFIZER and NONE.

3.2. After vaccination, individuals exhibit upward beliefs updating about their vaccine.

In Wave 2 (Figure 2 *left*), the average beliefs of individuals regarding Moderna’s safety is significantly higher for individuals who received MODERNA (5.92) than for individual who received PFIZER (5.39, $p < 0.001$, Student’s t-test). A similar asymmetric pattern is observed regarding the effectiveness of the vaccine (see Figure 2 *right*). *After* vaccination, the average beliefs of individuals regarding Moderna’s effectiveness is significantly

higher for individuals who received MODERNA (82.6) than for individuals who received PFIZER (74.1, $p < 0.001$, Student’s t-test).

These results are supported by the regression analyses illustrated in Table 1. We regress the posterior belief about Moderna on the treatment variable (treatment = MODERNA), the prior belief, and a set of covariates, including experiencing severe COVID-19 and severe side effects from the vaccine. When only two points in time are available, this estimation with lagged-dependent-variable adjustment is a common alternative to difference-in-differences. While difference-in-differences relies on the parallel trend assumption (which is very likely to hold, but not testable, in our two-wave panel), the causal identification of our lagged model relies on the assumption of ignorability conditional on past outcomes. That is, in the absence of treatment, the outcomes for both the treated and control groups would exhibit the same (conditional) distributions given the prior (Ashenfelter, 1978; Ding and Li, 2019).⁸

Between Wave 1 and Wave 2, people who received MODERNA update their beliefs about Moderna’s safety more positively (+0.54, $p < 0.001$; Student’s t-test) than those who received PFIZER, and this conditional on having the same initial belief, on not having experienced side effects and a set of other control variables which could plausibly affect belief updating. The same pattern is observed regarding Moderna’s effectiveness. Those who received MODERNA update their beliefs about Moderna’s effectiveness more positively (+6.32, $p < 0.001$; Student’s t-test) than those who received PFIZER, and this conditional on having the same initial belief and other covariates. More details are reported in SM, section A3.

Looking at the whole distribution of beliefs (rather than the mean) corroborates that individuals disproportionately shifted their beliefs upward about the vaccine they received. Figure 3 reports the cumulative distribution of beliefs about the safety and effectiveness of Moderna. It shows that, in Wave 2, MODERNA group’s beliefs about Moderna first-order stochastically dominate PFIZER group’s beliefs ($p < 0.001$, Somers’ D). In contrast, in Wave 1, the distributions were almost identical ($p = 0.738$ for safety and $p = 0.313$ for effectiveness, Somers’ D). Concretely, it means that a randomly selected respondent who received MODERNA is 27% more likely to have a higher belief about Moderna’s safety and 30% more likely to have a higher belief about Moderna’s effectiveness, than a randomly selected respondent who received PFIZER.

Evidence of asymmetric belief updating is also observed when considering ordinal (rather than cardinal) measures of safety and effectiveness. Figure 4 compares, among individuals who received Moderna, the fraction of them who consider Moderna safer and more effective than Pfizer. Figure 4 *left* shows that the percentage of participants who ranked Moderna as safer is much higher after vaccination than before (-0.35 , $p < 0.001$, Proportion test). Figure 4 *right* displays a similar pattern for Moderna’s effectiveness (-0.27 , $p < 0.001$, Proportion test). Said differently, while about only 1 respondent in 8 considered Moderna to be safer or more effective than Pfizer ex-ante, about half of those who received MODERNA consider it the safest and most effective option ex-post. When we replicate the analysis for beliefs about Pfizer, we find evidence of asymmetric belief updates about its safety, but not about its effectiveness (see SM, section A2). We do not have a clear explanation for this asymmetry. Section A2 in the SM discusses it further.

Finally, section A4 and Figure S6 in the SM shows that ex-post beliefs about AstraZeneca and Janssen are not significantly different between the MODERNA and PFIZER groups. There is a slight update between Wave 1 and Wave 2, but this update is common to the MODERNA and PFIZER groups. This placebo test corroborates that belief divergence about Moderna can be causally attributed to the fact of receiving that specific health treatment.

⁸Incidentally, if the ignorability condition is violated (but the parallel trend hypothesis is valid), the regression coefficient associated with the treatment variable would *under-estimate* the true effect (see Angrist and Pischke, 2009, p.184).

	<i>Dep. var.: Safety</i>			<i>Dep. var.: Effectiveness</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment = Moderna	0.54*** (0.13)	0.52*** (0.13)	1.73*** (0.54)	8.43*** (2.02)	6.32*** (1.91)	26.24*** (8.66)
Prior		0.37*** (0.04)	0.42*** (0.04)		0.36*** (0.04)	0.41*** (0.05)
Prior × Treatment			-0.23** (0.10)			-0.26** (0.11)
Socio-demographics		✓	✓		✓	✓
Experience with COVID-19		✓	✓		✓	✓
Time dummies		✓	✓		✓	✓
Geographical dummies		✓	✓		✓	✓
Constant	5.39*** (0.06)	3.43*** (0.52)	3.19*** (0.53)	74.14*** (0.83)	43.13*** (8.16)	38.84*** (8.32)
Observations	550	545	545	550	545	545
R^2	0.028	0.242	0.250	0.031	0.234	0.242

Reading note: Linear regressions of posterior belief on the treatment and other covariates. Columns (1)-(3) refer to beliefs about safety; columns (4)-(6) refer to beliefs about effectiveness. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. *Socio-demographics* include age and dummies for gender, student status, employment status, pregnancy/particular health condition. *Experience with COVID-19* includes the following covariates: vaccine quiz score, dummies = 1 if the participants report having been hospitalized because of COVID-19, having had severe side effects, having tested positive after the vaccine, having received one dose only, knowing someone who had severe side effects after the same vaccine, knowing someone who tested positive after the same vaccine, knowing someone who had severe side effects after a different vaccine, and knowing someone who tested positive after a different vaccine. *Time dummies* include weekly time dummies for the date of participation in Wave 2 and monthly time dummies for the date the last dose was received. *Geographical dummies* include dummies for the country of residence (England, Scotland, Wales or Northern Ireland). Table S7 in SM reports the full results.

Table 1: Regressions of the posterior beliefs about Moderna with lagged-dependent-variable adjustment

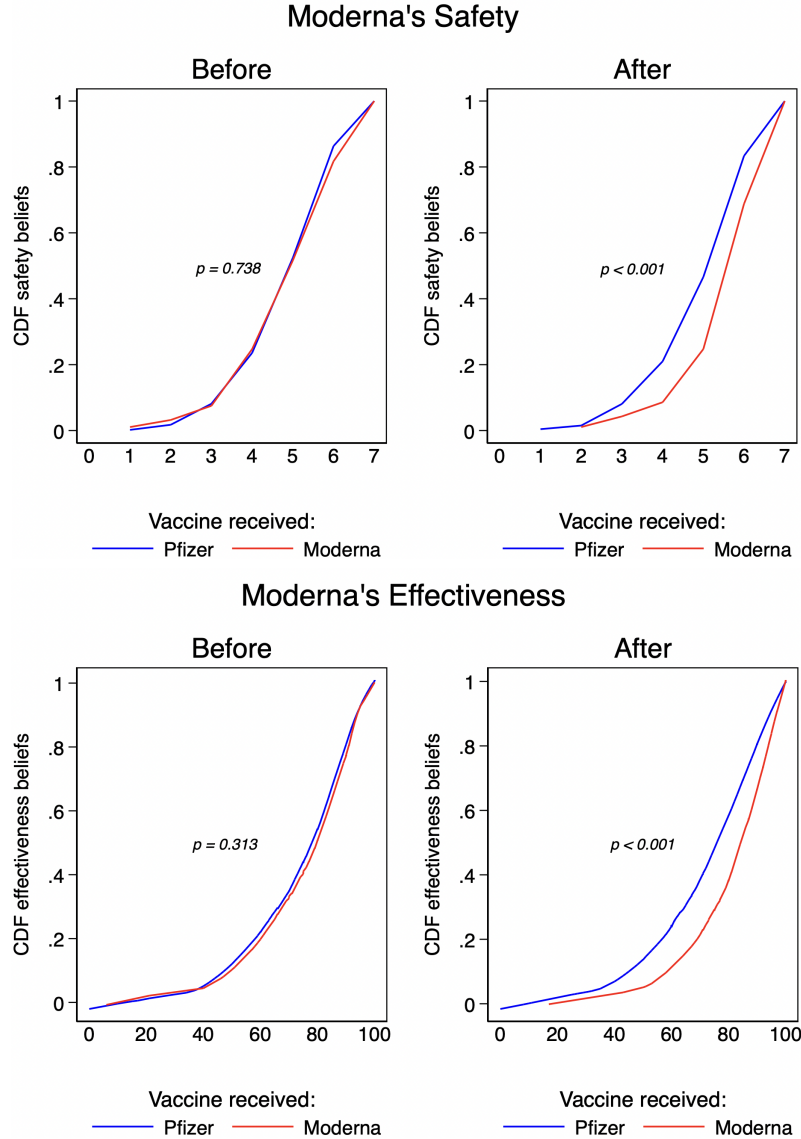


Figure 3: Cumulative density function of beliefs about Moderna

Reading note: The figure compares the cumulative distributions of beliefs of the MODERNA and PFIZER groups before and after receiving the vaccine. ‘Before’ corresponds to belief elicitation that happened in June-July 2021, i.e. before individuals received their first vaccination. ‘After’ corresponds to belief elicitation that happened in November-December 2021, i.e. after individuals received a vaccination. P-values come from Somers’ D tests. $N = 550$, data from Waves 1 and 2, groups MODERNA and PFIZER.

3.3. The asymmetry is driven by individuals experiencing a mismatch between their priors and their vaccine.

The belief updating of individuals who received MODERNA vs. PFIZER shows that the received vaccine predicts belief updating. We now investigate the role of priors *within* each treatment group. Table 1, columns (3) and (6) include an interaction term between priors and treatment. The estimated coefficient associated with the interaction term is negative and significant, indicating that receipt of Moderna impacts posterior beliefs about the safety and effectiveness of that vaccine less positively among those who had higher prior beliefs.

To clarify the role of high and low priors, we classify an individual as having *low priors* about the vaccine he will receive if he belongs to the 50% of individuals who had the lowest beliefs about its safety (resp. effectiveness)

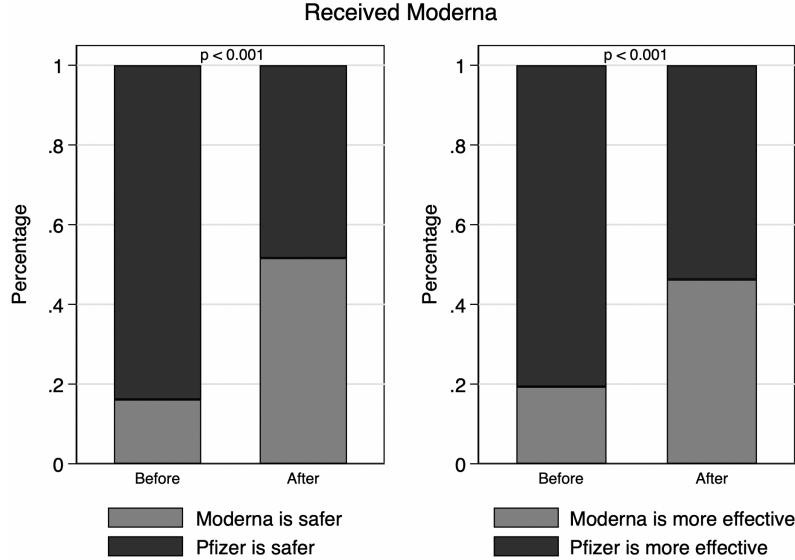


Figure 4: Relative safety and effectiveness about Moderna

Reading note: The *left* part of the figure shows the share of individuals who ranked Moderna as safer than Pfizer (or vice versa). The *right* part of the figure shows the share of individuals who ranked Moderna as more effective than Pfizer (or vice versa). ‘Before’ corresponds to belief elicitation that happened in June-July 2021, i.e. before individuals received their first vaccination. ‘After’ corresponds to belief elicitation that happened in November-December 2021, i.e. after individuals received a vaccination. P-values come from Proportion tests. $N = 93$, data from Waves 1 and 2, group MODERNA.

in Wave 1. Otherwise, the individual is classified as having *high priors*.⁹ This classification allows identifying four subgroups: individuals with *high* (1) vs. *low* (2) priors about a vaccine they will receive; and individuals with *high* (3) vs. *low* (4) priors about a vaccine they will *not* receive. Ex-post, prior beliefs *match* the received vaccine for categories (1) and (4). However, for categories (2) and (3), there is a *mismatch* between prior beliefs and the received vaccine.

Figure 5 reproduces Figure 2 separating between people who experienced a *match* or a *mismatch* between their prior beliefs and the vaccine they actually received. This decomposition highlights two patterns. First, individuals with *low* priors about their vaccine disproportionately update their beliefs upward, compared to those who did not receive that vaccine. Second, individuals with *high* priors about a vaccine they did *not* receive negatively update their beliefs about that vaccine, while we do not observe significant updating for individuals who received the vaccine they had high priors about. These patterns are observed whether we consider Moderna’s safety or effectiveness.

Overall, this decomposition shows that asymmetric belief updating is driven by individuals who are the most at risk of experiencing a discrepancy between their prior beliefs and the vaccine they received.

3.4. Individuals also shift their memories and stated preferences according to the vaccine they received.

In Wave 2, individuals were asked to recall how safe and effective they used to think each vaccine was about five months earlier. While in Wave 1 there was no difference in beliefs between the MODERNA and PFIZER groups, in Wave 2 we observe a significant difference in memories between the two groups: MODERNA recall higher beliefs than PFIZER, both about safety ($Mean_{MOD} - Mean_{PFI} = 0.38$, $p = 0.005$, Student’s t-test) and

⁹Our results are robust to alternative classifications, with the threshold set either at the lower quartile or at the upper quartile (see SM, Figure S7).

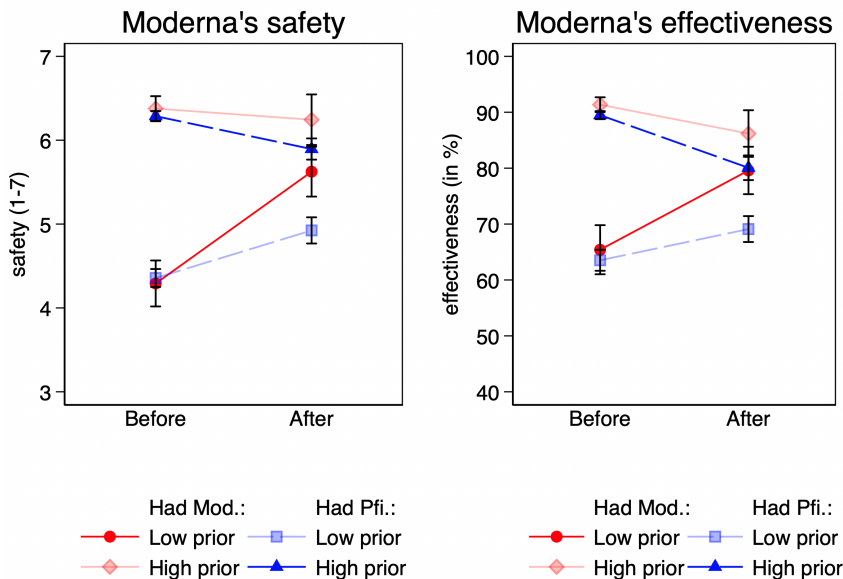


Figure 5: Beliefs about Moderna’s vaccine, by priors

Reading note: ‘Before’ corresponds to belief elicitation that happened in June-July 2021, i.e. before individuals received their first vaccination. ‘After’ corresponds to belief elicitation that happened in November-December 2021, i.e. after individuals received a vaccination. An individual is classified as having Low priors about the vaccine he will receive if he belongs to the 50% of individuals who had the lowest beliefs about its safety (resp. effectiveness) *before* vaccination. Otherwise, the individual is classified as having High priors. Error bars refer to 95% confidence intervals. $N = 550$, data from Waves 1 and 2, groups MODERNA and PFIZER.

effectiveness ($Mean_{MOD} - Mean_{PFI} = 6.03$, $p = 0.005$, Student’s t-test), as displayed in Figure 6. Ex-post memories shift in a similar fashion as beliefs (see also SM, section A9, Figure S9). This pattern is related to, but different from, hindsight bias (i.e., the tendency to recall the past as more predictable than it actually was), that we also observe in our sample (see SM, section A7, Table S9).

Regarding individuals’ hypothetical choices, we measured how they shifted following the vaccination. In Wave 1, individuals were asked which vaccine they would choose to receive if they had the choice. In Wave 2, they were asked which vaccine they would have chosen to receive if they have had the choice. Here again, we observe a strong asymmetry conditional on the received vaccine. In Wave 1, before knowing which vaccine they will receive, the percentage of participants who would choose to receive Moderna is about the same between the two groups MODERNA and PFIZER (0.01, $p = 0.633$, Proportion test). In contrast, in Wave 2, people who received Moderna are significantly more likely to report that they would have chosen to receive Moderna all along (0.26, $p < 0.001$, Proportion test). Similar effects appear when comparing the stated preference for Pfizer between groups in Wave 1 (0.05, $p = 0.320$, Proportion test) and in Wave 2 (0.30, $p < 0.001$, Proportion test). Figure 7 summarizes these findings.

4. Mechanisms

In the following sections, we explore different psychological mechanisms of how individuals may update their beliefs. We show that our results can neither be explained by conventional experience effects nor by some enhanced knowledge about the benefits of one’s own vaccine. Instead, they are consistent with the predictions of motivated beliefs as well as a model of over-inference from weak signals.

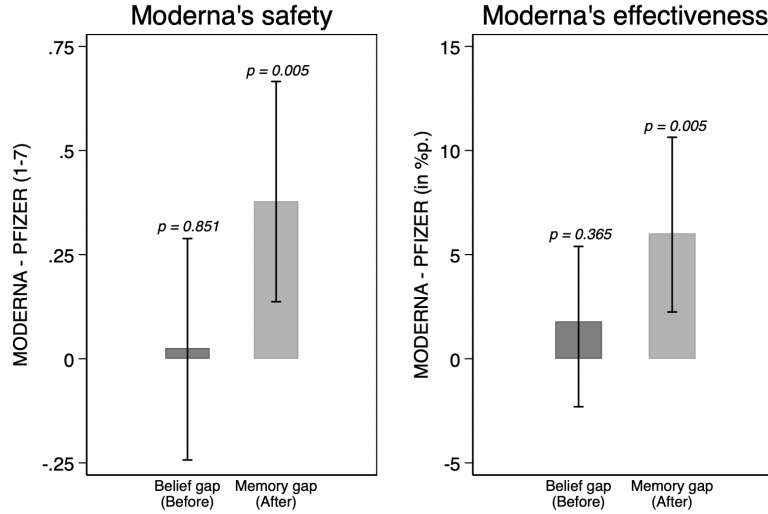


Figure 6: Difference in beliefs and memories about Moderna's vaccine

Reading note: In each figure, the left bar (i.e., belief gap) represents the difference between the *beliefs* reported before vaccination by the groups MODERNA and PFIZER. The right bar (i.e., memory gap) represents the difference between the *memories* reported after vaccination by those two groups. *Before* vaccination, there is no significant difference in beliefs about Moderna's safety and effectiveness. *After* vaccination, those who received MODERNA recall it to be safer and more effective than those who received PFIZER. P-values come from Student's t-tests. Error bars refer to 95% confidence intervals. $N = 550$, data from Waves 1 and 2, groups MODERNA and PFIZER.

4.1. Bayesian updating cannot rationalize the observed asymmetric updating pattern.

Isn't all this rational? People are likely to exploit observations from their personal experience to infer the general properties of their vaccine, as predicted by the theory of Bayesian inference. To get a sense of the extent to which Bayesian updating descriptively fits the data, we look at the size of the update among those who experienced no severe side effects (99% of the sample).¹⁰ When we make a parametric estimation of the exact posterior predicted by Bayesian updating (details of the computation are in the SM, section A5), we find it to be inconsistent with the observed results. Individuals exhibit upward belief updating that exceeds by far the update predicted by Bayesian updating.

The order of magnitude is akin to someone drawing a green ball from an urn where they expect to have one million green balls and one black ball, and subsequently thinking that there must be *two* million green balls (and one black) in the urn. In the case of long-term side effects (arguably the strongest factor of vaccine hesitancy in our sample, see ONS, 2021), there was no feedback at all, since individuals were interviewed only a few weeks after their injection.

Overall, despite the fact that people might use the outcome of their vaccination to update their beliefs, conventional Bayesian inference cannot explain the large size of the updates that we observe.

4.2. Individuals neither prioritize positive information nor are better informed about their own vaccine.

Individuals were interviewed during a period where immense amount of information about vaccines was available, and was, sometimes, contradictory. In this context, people must decide what they wish to know

¹⁰We also look at the update of the 8 individuals who declared that they experienced severe side effects (what should be considered as "severe side effects" was left to the participant's interpretation): two of them revised their beliefs *downward*, other two revised *upwards* and the remaining four did not update their beliefs.

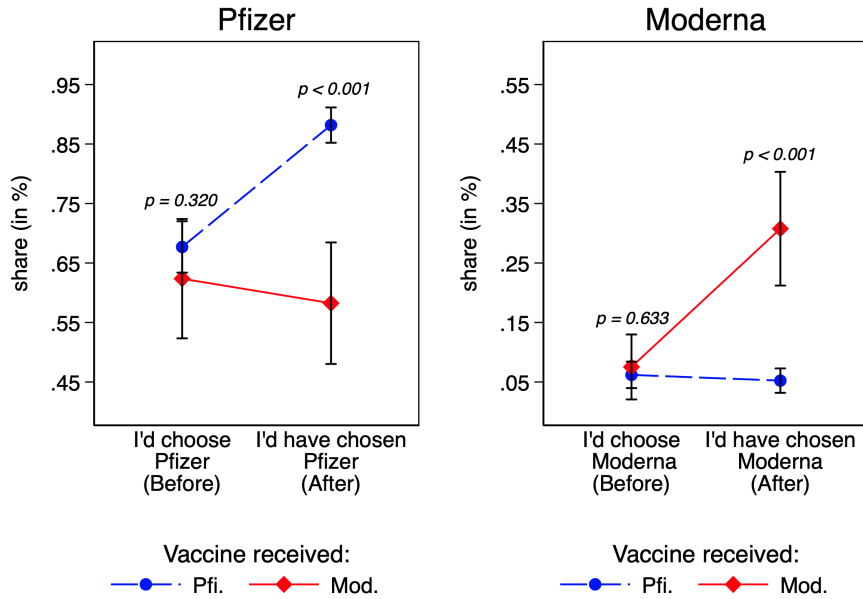


Figure 7: Stated preferences for vaccines

Reading note: *Before* vaccination, participants were asked to state among all vaccines the one that they would choose to receive if they had the choice. *After* vaccination, participants were asked to state among all vaccines the one that they would have chosen to receive if they had the choice. ‘Before’ corresponds to belief elicitation that happened in June-July 2021, i.e. before individuals received their first vaccination. ‘After’ corresponds to belief elicitation that happened in November-December 2021, i.e. after individuals received a vaccination. Note that y-axes are re-scaled for readability. P-values come from Proportion tests. Error bars refer to 95% confidence intervals. $N = 550$, data from Waves 1 and 2, groups MODERNA and PFIZER.

(see Golman et al., 2017; Sharot and Sunstein, 2020), and the vaccine they received could impact this attitude toward information.

Different elements allow us to investigate the potential role played by information seeking. First of all, at the end of our survey in Wave 2, we gave each individual the possibility to read information regarding COVID-19 vaccines. They had to choose whether they wanted to read a piece of information about the ‘advantages’ or ‘disadvantages’ of the vaccine they did or did not receive (see SM, section A6).¹¹ We indeed observe an asymmetry in preferences for information, but it is between own vs. other vaccine, rather than between favorable vs. unfavorable information. Around 90% of vaccinated individuals chose to read a piece of information related to the vaccine they received. About half chose to read about its relative advantages, and about half preferred to read about its relative disadvantages (see Figure S8 in SM). When given the choice, individuals did not actively select favorable information. However, since most scientific public information about vaccines was positive, seeking more information about one’s own vaccine could lead individuals to form more optimistic beliefs about it. If individuals were more optimistic about their vaccine because they were more informed, we should observe that beliefs about their own vaccine selectively converge toward the state of scientific knowledge. That is, we should observe negative updating for individuals who ex-ante overshoot the safety and the effectiveness of the vaccine they received. This is not the case: optimistic updating is observed even if most individuals already

¹¹The incentives to seek information right after vaccination and at the time of Wave 2 were somewhat different, as information on short-term side-effects was now redundant. However, other potential information (about medium-to-long-term side effects, e.g., infertility; about the length of the protection against severe/fatal forms of COVID-19 and against new COVID-19 variants) was still valuable at the time of Wave 2.

overshoot the safety of the vaccine they will receive (scientific evidence suggests safety to be around 4/7 on the Likert scale, i.e. one severe side effect every 10^4 vaccines, see Polack et al., 2020; Baden et al., 2021). In other words, individuals do *not* hold more accurate beliefs about their vaccine *after* receiving it than *before*. Finally, our survey in Wave 2 also included a short quiz designed to measure individuals' general knowledge about COVID-19 vaccines. Variations in the size of belief update are uncorrelated with the participants' score (Pearson correlation coefficient of 0.05 or lower), nor the score variable helps explain beliefs variations in the regressions with lagged-dependent-variable adjustment (see Table 1).

Overall, we observe information seeking behavior to correlate with the treatment (PFIZER or MODERNA), but we do not observe signs indicating that individuals prioritize positive information about their own vaccine, are better informed and/or hold more accurate beliefs about it.

4.3. People update their beliefs in a fashion consistent with motivated beliefs.

In our survey, individuals were interviewed about a topic with a clearly desirable outcome for their current and future health: the safer and the more effective their vaccine, the better. Once the individual has received a specific vaccine, there is no going back. The vaccine injection is definitive, with potentially large long-term benefits/costs. Insofar as the assignment to a specific vaccine makes its safety and effectiveness much more desirable for the patient, one might expect that this change in desires will result in more optimistic beliefs. This is the core prediction of motivated beliefs.

Motivated beliefs (Kunda, 1990; Bénabou, 2015; Zimmermann, 2020; Möbius et al., 2022), aka *wishful thinking* (Mayraz, 2011; Engelmann et al., 2019; Caplin and Leahy, 2019; Mayraz, 2019), refer to people's tendency to believe what they would like to be true. The literature has abundantly shown that people's desire for good health can affect their beliefs (see, e.g., Weinstein, 1980, 1982; Jansen et al., 2011; Hanoch et al., 2019; Brnstrm and Brandberg, 2010) and the natural experiment that we study seems to confirm it.

The three main results presented in Section 3 are in line with the predictions of motivated beliefs.¹² First, individuals who now have an interest in their vaccine being safer and more effective than they used to think, end up changing their mind in this direction. That is, diverging wishes generate diverging beliefs, as predicted by Caplin and Leahy (2019). Second, those who had high priors about a vaccine that they did *not* receive deflate their beliefs about that vaccine, while those who had low priors about their vaccine update their beliefs overly optimistically. These two patterns are in line with, respectively, a *sour grape effect* and *sweet lemon effect*,¹³ i.e., two strategies to reduce cognitive dissonance (Kay et al., 2002), that have potentially far-reaching economic consequences (Dalton et al., 2016). Third, individuals who received Moderna tend to recollect their past beliefs as more favorable to Moderna than they actually were, consistently with self-serving motives (Zimmermann, 2020; Müller, 2022), in particular avoiding regret. Regret avoidance is intimately related to the sour grape effect, as someone recalling that they thought well of Moderna is also someone who is justifying their past self for making a responsible decision (see Sugden, 1985, for a discussion). Finally, we observe that many individuals shift their stated preferences according to the vaccine they received, consistently with adaptive preferences

¹²In the previous section, we presented no evidence of motivated information seeking. By itself, this does not undermine the motivated nature of belief formation, but it sets a limit to its scope (i.e., through information processing only, and not information seeking).

¹³The sour grape effect refers to people's tendency to devalue something they desire but cannot have; the sweet lemon effect refers to the converse, i.e. the tendency to positively reappraise something that is attainable or attained, even if it was initially undesirable. The name "sour grape" comes from Aesop's fable "The Fox and the Grapes", popularized by La Fontaine.

(Elster, 1983).¹⁴

Overall, individuals update beliefs, memory and stated preferences in a fashion that is overly favorable for their current situation.

4.4. Results can be rationalized by alternative non-Bayesian mechanisms.

We find that people’s beliefs react strongly to their assignment to a vaccine, although their personal experience is poorly informative about unlikely events. This pattern might seem at odds with the well-established and well-replicated finding that humans tend to *under-infer* from signal(s), compared to the Bayes benchmark (see the review by Benjamin, 2019). Recently, however, Augenblick et al. (2021) has rightly pointed out that this literature has looked almost exclusively at highly informative signals. When weak signals are considered instead, people tend to over-react to new information relative to Bayes theorem, both in the lab and in the field. They propose a model of cognitive imprecision, in which people are unsure about the informativeness of the signal they receive, and therefore inflate how much they can learn from poorly informative events. Our results about belief updating can be accommodated within this non-Bayesian framework.

The source of information, i.e. first-hand experience, might play a role too. In a series of experiments, Conlon et al. (2022) document a general tendency to over-infer personal information compared to other sources. In all their treatments, Conlon et al. (2022) find that people better recall the information from their own experience than some equally well-observed and informative signals from others. The salience of the vaccination experience could also exacerbate over-inference, as predicted by diagnostic expectations (Gennaioli and Shleifer, 2010; Bordalo et al., 2018, 2020). According to this model, people over- or under-estimate probabilities according to how easily information comes to mind. In the case of the personal experience with the vaccine, we can reasonably assume that patients use their personal experience as highly representative when accessing information from memory.

Diagnostic expectations could also explain the memory asymmetry that we observe, but a more general explanation is possible. Even if participants cannot recall exactly what they used to think about vaccines, they probably have a general idea, that they adjust based on the relevant information they have available at the moment. Their current beliefs are among this relevant information. This heuristic process could explain why we observe a recall bias in the direction of inflated memories for one’s own vaccine, on top of inflated beliefs. A similar process could explain the bias in stated preferences.

Overall, the pattern of beliefs updating that we document could be explained by some non-Bayesian mechanisms, whereby individuals learned too much from their own experience.

5. Conclusions

Our study identifies the causal effect of vaccine assignment on beliefs updating, stated preferences and memory distortions. We showed that, *before* vaccination, individuals exhibit similar beliefs and similar stated preferences about the different vaccines. *After* vaccination however, they exhibit upward beliefs updating about the safety and the effectiveness of the vaccine they received. When asked to recall their past opinions, they tend to wrongly remember that they thought well of their vaccine all along. When asked what they would have chosen to receive, they are more likely to cite their vaccine than those who received another one.

¹⁴In practice, it is very hard to disentangle to what extent either the change in preferences trickled down on beliefs or the other way round. Elster (1983, p.124) notices that the difference is so subtle that in the French version of La Fontaine’s tale about sour grapes the fox is deluded in his perception of the vermilion grapes, and wrongly believes that they are green. In the English version, he wrongly believes that they are sour, i.e., a matter of preferences rather than beliefs.

Overall, our results are consistent with a pattern of motivated beliefs, where individuals seem to update beliefs to convince themselves that they are in the best state of the world. Receiving a vaccine is likely to impact beliefs via the change in desires, over and above the purely informational content of the post-vaccination experience and the effect of skewed information seeking. These results are also consistent with a cognitive model of over-inference, where people learn too much from their experience and therefore end up being overly optimistic about their own vaccine. While the former interpretation has the advantage of accounting for the psychological needs for self-reassurance that probably occurred in this specific context, the latter has the advantage of being more parsimonious, as it relies on purely cognitive mechanisms. Importantly, there is little reason to think that motivational and cognitive mechanisms are mutually exclusive, while it is plausible that they both intervene in the formation of beliefs (Melnikoff and Strohminger, 2023; Gilovich, 2008, p.80).

In designing our study, we decided to focus on beliefs as the main variable of interest, and we carefully constructed the online questionnaire accordingly: we asked individuals about their beliefs at the beginning of the survey, about their memories just after and about their hypothetical choice later on. Beliefs, memories and preferences are deeply intertwined, and our design does not allow to pinpoint to what extent, say, motivated memory bent current beliefs or the other way round. Such mediation analyses have been explored elsewhere (Bordalo et al., 2022) and are beyond the scope of this study. Nevertheless, all reported effects can be causally associated to the administration of a specific vaccine.

Optimistic belief updating can be either harmful or beneficial. On the one hand, optimistic beliefs about COVID-19 vaccines could act as a “belief trap” (Scheffer et al., 2022) potentially leading to harmful behaviors (e.g. non-compliance with safety rules). On the other hand, they may protect individuals’ well-being by reducing fear and anxiety, especially given the irreversible nature of vaccine injection (Jefferson et al., 2017). More empirical evidence is needed to establish the precise psychological mechanism and assess the relative costs and benefits. Yet, our results can inform the debate about whether individuals should be able to choose which health treatment they receive (including which COVID-19 booster jab, see Kramer et al. 2021).¹⁵ The natural experiment that we investigate clearly shows that individuals form optimistic beliefs about the treatment they receive, even when they cannot choose it.

¹⁵Many other factors matter, of course. While giving vaccination choice respects the principle of patient self-determination and may increase overall vaccination acceptance, accommodating individual vaccine preferences would exacerbate current inequities in vaccine administration and potentially cast doubt on the fact that each authorized vaccine works. Additionally, choice of the vaccine may become a target of misinformation campaigns from vaccine companies or social media.

References

- Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*, Princeton University Press.
- Ashenfelter, O. (1978). Estimating the effect of training programs on earnings, *The Review of Economics and Statistics* **60**(1): 47–57.
- Ashworth, M., Thunström, L., Cherry, T. L., Newbold, S. C. and Finnoff, D. C. (2021). Emphasize personal health benefits to boost COVID-19 vaccination rates, *Proceedings of the National Academy of Sciences* **118**(32): e2108225118.
- Augenblick, N., Lazarus, E. and Thaler, M. (2021). Overinference from weak signals and underinference from strong signals, arXiv preprint arXiv:2109.09871.
URL: <https://drive.google.com/file/d/1irPVAmdcz8zKVzqplHkvw1wdnAN2H7Qh/view>
- Baden, L. R., El Sahly, H. M., Essink, B., Kotloff, K., Frey, S., Novak, R., Diemert, D., Spector, S. A., Rouphael, N., Creech, C. B. et al. (2021). Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine, *New England journal of medicine* **384**(5): 403–416.
- Bar-Hillel, M. and Budescu, D. (1995). The elusive wishful thinking effect, *Thinking & Reasoning* **1**(1): 71–103.
- Beasley, R. K. and Joslyn, M. R. (2001). Cognitive dissonance and post-decision attitude change in six presidential elections, *Political psychology* **22**(3): 521–540.
- Bénabou, R. (2015). The economics of motivated beliefs, *Revue d'économie politique* **125**(5): 665–685.
- Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases, *Handbook of Behavioral Economics: Applications and Foundations 1* **2**: 69–186.
- Benjamin, D. J., Heffetz, O., Kimball, M. S. and Rees-Jones, A. (2012). What do you think would make you happier? What do you think you would choose?, *American Economic Review* **102**(5): 2083–2110.
- Bordalo, P., Burro, G., Coffman, K. B., Gennaioli, N. and Shleifer, A. (2022). Imagining the future: memory, simulation and beliefs about COVID, NBER Working Paper 30353.
URL: <https://www.nber.org/papers/w30353>
- Bordalo, P., Gennaioli, N., Ma, Y. and Shleifer, A. (2020). Overreaction in macroeconomic expectations, *American Economic Review* **110**(9): 2748–2782.
- Bordalo, P., Gennaioli, N. and Shleifer, A. (2018). Diagnostic expectations and credit cycles, *The Journal of Finance* **73**(1): 199–227.
- Bradley, G. W. (1978). Self-serving biases in the attribution process: A reexamination of the fact or fiction question., *Journal of personality and social psychology* **36**(1): 56.
- Brnstrm, R. and Brandberg, Y. (2010). Health risk perception, optimistic bias, and personal satisfaction, *American Journal of Health Behavior* **34**(2): 197–205.
- Bruvoort, K. J., Sy, L. S., Qian, L., Ackerson, B. K., Luo, Y., Lee, G. S., Tian, Y., Florea, A., Takhar, H. S., Tubert, J. E. et al. (2022). Real-world effectiveness of the mRNA-1273 vaccine against COVID-19: Interim results from a prospective observational cohort study, *The Lancet Regional Health-Americas* **6**: 100134.
- Campos-Mercade, P., Meier, A. N., Schneider, F. H., Meier, S., Pope, D. and Wengström, E. (2021). Monetary incentives increase COVID-19 vaccinations, *Science* **374**(6569): 879–882.
- Caplin, A. and Leahy, J. V. (2019). Wishful thinking, NBER Working Paper 25707.
URL: https://www.nber.org/system/files/working_papers/w25707/w25707.pdf
- Conlon, J. J., Mani, M., Rao, G., Ridley, M. and Schilbach, F. (2022). Failing to learn from others, NBER Working Paper 28844.
URL: <https://gautam-rao.com/pdf/Not%20Learning%20from%20Others.pdf>

- Cox, J. C., Green, E. P. and Hennig-Schmidt, H. (2016). Experimental and behavioral economics of healthcare, *Journal of Economic Behavior & Organization* **131**: A1–A4.
- Dalton, P. S., Ghosal, S. and Mani, A. (2016). Poverty and aspirations failure, *The Economic Journal* **126**(590): 165–188.
- Ding, P. and Li, F. (2019). A bracketing relationship between difference-in-differences and lagged-dependent-variable adjustment, *Political Analysis* **27**(4): 605–615.
- Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself, *American Economic Journal: Microeconomics* **3**(2): 114–138.
- Elster, J. (1983). *Sour grapes*, Cambridge University Press.
- Engelmann, J., Lebreton, M., Schwardmann, P., van der Weele, J. J. and Chang, L.-A. (2019). Anticipatory anxiety and wishful thinking, *American Economic Review* (forthcoming).
URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20191068&from=f>
- Eyster, E., Li, S. and Ridout, S. (2021). A theory of ex-post rationalization, arXiv preprint:2107.07491.
URL: <https://arxiv.org/pdf/2107.07491.pdf>
- Festinger, L. (1962). Cognitive dissonance, *Scientific American* **207**(4): 93–106.
- Gadkari, A. S. and McHorney, C. A. (2012). Unintentional non-adherence to chronic prescription medications: How unintentional is it really?, *BMC health services research* **12**(1): 1–12.
- Galizzi, M. M. and Wiesen, D. (2018). Behavioral experiments in health economics, Oxford Research Encyclopedia of Economics and Finance.
URL: <https://doi.org/10.1093/acrefore/9780190625979.013.244>
- Ganguly, A. and Tasoff, J. (2017). Fantasy and dread: The demand for information and the consumption utility of the future, *Management Science* **63**(12): 4037–4060.
- Gennaioli, N. and Shleifer, A. (2010). What comes to mind, *The Quarterly Journal of Economics* **125**(4): 1399–1433.
- Gilovich, T. (2008). *How we know what isn't so*, Simon and Schuster.
- Golman, R., Hagmann, D. and Loewenstein, G. (2017). Information avoidance, *Journal of Economic Literature* **55**(1): 96–135.
- Hanoch, Y., Rolison, J. and Freund, A. M. (2019). Reaping the benefits and avoiding the risks: unrealistic optimism in the health domain, *Risk Analysis* **39**(4): 792–804.
- Horne, R., Chapman, S. C., Parham, R., Freemantle, N., Forbes, A. and Cooper, V. (2013). Understanding patients' adherence-related beliefs about medicines prescribed for long-term conditions: a meta-analytic review of the Necessity-Concerns framework, *PloS one* **8**(12): e80633.
- Horne, R. and Weinman, J. (1999). Patients' beliefs about prescribed medicines and their role in adherence to treatment in chronic physical illness, *Journal of psychosomatic research* **47**(6): 555–567.
- Jamieson, K. H., Romer, D., Jamieson, P. E., Winneg, K. M. and Pasek, J. (2021). The role of non-COVID-specific and COVID-specific factors in predicting a shift in willingness to vaccinate: A panel study, *Proceedings of the National Academy of Sciences* **118**(52): e2112266118.
- Jansen, L. A., Appelbaum, P. S., Klein, W. M., Weinstein, N. D., Cook, W., Fogel, J. S. and Sulmasy, D. P. (2011). Unrealistic optimism in early-phase oncology trials, *Irb* **33**(1): 1.
- Jefferson, A., Bortolotti, L. and Kuzmanovic, B. (2017). What is unrealistic optimism?, *Consciousness and cognition* **50**: 3–11.
- Kahan, D. M. (2012). Ideology, motivated reasoning, and cognitive reflection: An experimental study, *Judgment and Decision making* **8**: 407–24.

- Kaplan, R. M. and Milstein, A. (2021). Influence of a COVID-19 vaccine’s effectiveness and safety profile on vaccination acceptance, *Proceedings of the National Academy of Sciences* **118**(10): e2021726118.
- Kay, A. C., Jimenez, M. C. and Jost, J. T. (2002). Sour grapes, sweet lemons, and the anticipatory rationalization of the status quo, *Personality and Social Psychology Bulletin* **28**(9): 1300–1312.
- Kramer, D. B., Opel, D. J., Parasidis, E. and Mello, M. M. (2021). Choices in a crisis—individual preferences among SARS-CoV-2 vaccines, *New England Journal of Medicine* **384**(17): e62.
- Krizan, Z. and Windschitl, P. D. (2009). Wishful thinking about the future: Does desire impact optimism?, *Social and Personality Psychology Compass* **3**(3): 227–243.
- Kunda, Z. (1990). The case for motivated reasoning., *Psychological bulletin* **108**(3): 480.
- Levin, J. M., Bukowski, L. A., Minson, J. A. and Kahn, J. M. (2023). The political polarization of COVID-19 treatments among physicians and laypeople in the United States, *Proceedings of the National Academy of Sciences* **120**(7): e2216179120.
- Lord, C. G., Ross, L. and Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence, *Journal of personality and social psychology* **37**(11): 2098.
- Mahmud, S., Mohsin, M., Khan, I. A., Mian, A. U. and Zaman, M. A. (2021). Knowledge, beliefs, attitudes and perceived risk about COVID-19 vaccine and determinants of COVID-19 vaccine acceptance in Bangladesh, *PloS one* **16**(9): e0257096.
- Mayraz, G. (2011). Wishful thinking, Available at SSRN 1955644.
URL: <http://dx.doi.org/10.2139/ssrn.1955644>
- Mayraz, G. (2019). Priors and desires: A Bayesian model of wishful thinking and cognitive dissonance, University of Sydney WP.
URL: <https://mayraz.com/papers/PriorsAndDesires.pdf>
- Melnikoff, D. and Strohminger, N. (2023). Bayesianism and wishful thinking are compatible, *Nature Human Behavior* **4**: 1258–1264.
URL: <https://osf.io/preprints/psyarxiv/yhmvw>
- Möbius, M. M., Niederle, M., Niehaus, P. and Rosenblat, T. S. (2022). Managing self-confidence: Theory and experimental evidence, *Management Science* **68**(11): 7793–7817.
- Mullainathan, S. and Washington, E. (2009). Sticking with your vote: Cognitive dissonance and political attitudes, *American Economic Journal: Applied Economics* **1**(1): 86–111.
- Müller, M. (2022). Selective memory around big life decisions, preprint.
URL: <https://app.box.com/s/ai7srlzy648yar6dbq9xpdtvcwhrgvdt>
- NHS England (2021). COVID-19 vaccination sites standard operating procedure, *Technical report*, NHS England.
URL: <https://www.england.nhs.uk/coronavirus/covid-19-vaccination-programme/guidance-for-vaccination-centres/>
- ONS (2021). Coronavirus vaccine hesitancy in younger adults: June 2021, *Technical report*, ONS.
URL: <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthandwellbeing/articles/coronavirusvaccinehesitancyinyoungeradults/june2021>
- Oster, E., Shoulson, I. and Dorsey, E. (2013). Optimal expectations and limited medical testing: Evidence from Huntington disease, *American Economic Review* **103**(2): 804–30.
- Palan, S. and Schitter, C. (2018). Prolific.ac—A subject pool for online experiments, *Journal of Behavioral and Experimental Finance* **17**: 22–27.

- Polack, F. P., Thomas, S. J., Kitchin, N., Absalon, J., Gurtman, A., Lockhart, S., Perez, J. L., Pérez Marc, G., Moreira, E. D., Zerbini, C. et al. (2020). Safety and efficacy of the BNT162b2 mRNA COVID-19 vaccine, *New England journal of medicine* **383**(27): 2603–2615.
- Rough, E. and Powell, T. (2021). Coronavirus: COVID-19 vaccine roll-out frequently asked questions, Commons Library Research Briefing.
URL: <https://commonslibrary.parliament.uk/research-briefings/cbp-9081/>
- Scheffer, M., Borsboom, D., Nieuwenhuis, S. and Westley, F. (2022). Belief traps: Tackling the inertia of harmful beliefs, *Proceedings of the National Academy of Sciences* **119**(32): e2203149119.
- Schneider, F. H., Campos-Mercade, P., Meier, S., Pope, D., Wengström, E. and Meier, A. N. (2023). Financial incentives for vaccination do not have negative unintended consequences, *Nature* **613**: 526–533.
- Schwardmann, P., Tripodi, E. and Van der Weele, J. J. (2022). Self-persuasion: Evidence from field experiments at international debating competitions, *American Economic Review* **112**(4): 1118–1146.
- Sharot, T., Korn, C. W. and Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality, *Nature Neuroscience* **14**(11): 1475–1479.
- Sharot, T. and Sunstein, C. R. (2020). How people decide what they want to know, *Nature Human Behaviour* **4**(1): 14–19.
- Sugden, R. (1985). Regret, recrimination and rationality, *Theory and decision* **19**: 77–99.
- Suzuki, T. (2019). Choice set dependent performance and post-decision dissonance, *Journal of Economic Behavior & Organization* **163**: 24–42.
- Tartof, S. Y., Slezak, J. M., Fischer, H., Hong, V., Ackerson, B. K., Ranasinghe, O. N., Frankland, T. B., Ogun, O. A., Zamparo, J. M., Gray, S. et al. (2021). Effectiveness of mRNA BNT162b2 COVID-19 vaccine up to 6 months in a large integrated health system in the USA: a retrospective cohort study, *The Lancet* **398**(10309): 1407–1416.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events, *Journal of personality and social psychology* **39**(5): 806.
- Weinstein, N. D. (1982). Unrealistic optimism about susceptibility to health problems, *Journal of behavioral medicine* **5**(4): 441–460.
- Weinstein, N. D., Marcus, S. E. and Moser, R. P. (2005). Smokers’ unrealistic optimism about their risk, *Tobacco control* **14**(1): 55–59.
- Zimmermann, F. (2020). The dynamics of motivated beliefs, *American Economic Review* **110**(2): 337–61.