ROCKWOOL Foundation Berlin
Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES 236/26

Beyond the Facts: Familiarity, Credibility, and Persuasion

Evangelos Dioikitopoulos, Rigissa Megalokonomou, Tommaso Sartori, Yves Zenou

www.rfberlin.com

SEPTEMBER 2026

Beyond the Facts: Familiarity, Credibility, and Persuasion

Authors

Evangelos Dioikitopoulos, Rigissa Megalokonomou, Tommaso Sartori, Yves Zenou

Reference

JEL Codes: C93, D83, I12, D91

Keywords: randomized experiment, persuasion, belief formation, source credibility, messenger familiarity, social media influencers, virtual reality, engagement, health communication, vaccination

Recommended Citation: Evangelos Dioikitopoulos, Rigissa Megalokonomou, Tommaso Sartori, Yves Zenou (2026): Beyond the Facts: Familiarity, Credibility, and Persuasion. RFBerlin Discussion Paper No. 236/26

Access

Papers can be downloaded free of charge from the RFBerlin website: https://www.rfberlin.com/discussion-papers

Discussion Papers of RFBerlin are indexed on RePEc: https://ideas.repec.org/s/crm/wpaper.html

Disclaimer

Opinions and views expressed in this paper are those of the author(s) and not those of RFBerlin. Research disseminated in this discussion paper series may include views on policy, but RFBerlin takes no institutional policy positions. RFBerlin is an independent research institute.

RFBerlin Discussion Papers often represent preliminary or incomplete work and have not been peer-reviewed. Citation and use of research disseminated in this series should take into account the provisional nature of the work. Discussion papers are shared to encourage feedback and foster academic discussion.

All materials were provided by the authors, who are responsible for proper attribution and rights clearance. While every effort has been made to ensure proper attribution and accuracy, should any issues arise regarding authorship, citation, or rights, please contact RFBerlin to request a correction.

These materials may not be used for the development or training of artificial intelligence systems.

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com

Beyond the Facts: Familiarity, Credibility, and Persuasion*

Evangelos Dioikitopoulos Rigissa Megalokonomou Tommaso Sartori Yves Zenou

August 31, 2026

Abstract

Public information campaigns often rely on credible expert messengers to influence people's beliefs. We study how messenger identity and communication technology shape persuasion in a large-scale randomized experiment on vaccination attitudes among young adults in Greece. Participants receive an identical public-health message that varies only in who delivers it—a medical scientist or a social media influencer—and in how it is delivered—through immersive virtual reality or conventional video. Exposure raises trust in vaccines and trust in the healthcare system. Scientists are rated substantially more credible than influencers yet are no more persuasive on average; the two groups' effectiveness rests on different characteristics—the scientists' on perceived reliability and knowledgeability, the influencers' on personal appeal, prior familiarity, and reputation. We find that persuasion is greater under virtual reality, consistent with stronger engagement. Changes in beliefs persist for at least six months, and initially hesitant participants report greater willingness to vaccinate. A model in which engagement, prior uncertainty, and effective credibility enter persuasion multiplicatively rationalizes these patterns.

Keywords: randomized experiment, persuasion, belief formation, source credibility, messenger familiarity, social media influencers, virtual reality, engagement, health communication, vaccination

JEL Classifications: C93, D83, I12, D91

*Evangelos Dioikitopoulos, Athens University of Economics and Business, Department of Economics, Greece (edioik@aueb.gr); Rigissa Megalokonomou, Monash University, Department of Economics, Australia, IZA, and CESifo (rigissa.megalokonomou@monash.edu); Tommaso Sartori, Monash University, Department of Economics, Australia (tommaso.sartori@monash.edu); Yves Zenou, Monash University, Department of Economics, Australia (yves.zenou@monash.edu).
We are grateful to Massimiliano Bratti, Sebastiano Della Lena, Claudio Labanca, Stelios Michalopoulos, Elias Papaioannou, Andrea Prat, and participants at seminars at Monash University and the University of Milan, as well as at the 2026 Conference on Research on Economic Theory and Econometrics, for their precious comments and feedback. We would also like to thank Ilias Aretos, Ilias Balatsiouras, Chloe Bouka, Alexandra Chairopoulou, Konstantina Kakkavou, Stelios Kalantzakis, Theodore Kourtalis, Alexandra Moutsatsou, Kelly Pantou, Maria Papaggeli, Nikos Papavasileiou, Angelos Rammos, Stylianos Rousalis, Georgia Samartzi, Panagiotis Souli, Chrysa Thanou, Nikos Yfantis, for their excellent research assistance. Rigissa Megalokonomou acknowledges research support from Shovel-ready, Monash University, Monash Establishment Grant (No: 1755774). The RCT has also been pre-registered at the AEA RCT Registry with RCT ID number AEARCTR-0016954, published on October 13, 2025, prior to the start of data collection on December 1, 2025 (https://doi.org/10.1257/rct.16954-1.0); Appendix B.1 maps the registration to the paper. This project has been reviewed by the Office of Research Ethics at the Athens University of Economics and Business (clearance number: $EY\Delta$/5-2025) and by the Monash University Human Research Ethics Committee (project ID: 49408).

1 Introduction

Governments and institutions often need to communicate information effectively on issues of public importance. Whether the information is believed depends on its content, its messenger, and its medium. On the choice of messenger, the standard approach emphasizes credibility and thus favors the most expert and trustworthy source available (Lupia and McCubbins 1998; Druckman 2001; Gentzkow and Shapiro 2006). However, behavioral theories of belief formation suggest a broader answer. Persuasion also depends on the engagement a message commands (DellaVigna and Gentzkow 2010) and on the associations evoked by the person delivering it (Mullainathan, Schwartzstein, and Shleifer 2008; Bordalo, Gennaioli, and Shleifer 2020), which implies that an audience's prior relationship with the messenger may matter alongside expertise. This distinction has become increasingly relevant. Expert communication now competes with content from social media influencers, whose influence can rest on familiarity and parasocial connections rather than formal credentials (Horton and Wohl 1956; K. Freberg, Graham, McGaughey, and L. A. Freberg 2011; Jin, Muqaddam, and Ryu 2019). At the same time, new communication technologies such as virtual reality (VR) offer new ways to deliver information and enhance the persuasiveness of a message (Herrera, Bailenson, Weisz, Ogle, and Zaki 2018; Andries, Bursztyn, Chaney, Djourelova, and Imas 2026). We provide experimental evidence on how messenger characteristics and communication technology shape persuasion when the information itself remains unchanged.

We study this question in the context of vaccination attitudes and beliefs. The pairing of messenger choice and vaccination is an old one: In 1956, at the height of the American polio campaign, public-health officials enlisted Elvis Presley—vaccinated in front of the press before appearing on *The Ed Sullivan Show*—betting that his fame would convince hesitant teenagers where experts' advice had not (Mawdsley 2016). The campaign appeared to vindicate that strategy, as vaccination rates for polio among American teenagers subsequently increased sharply. Seven decades later, however, the underlying question remains: How much does who delivers public-health information matter for persuasion? Vaccines are among the safest and most effective public-health technologies ever developed (World Health Organization and UNICEF 2024). However, routine vaccination coverage remains below pre-pandemic levels in many countries, while public satisfaction with healthcare systems has declined across much of the OECD in the post-pandemic period (World Health Organization 2025; OECD 2025; OECD 2024). Across many advanced economies, including the United States and other OECD countries, declining vaccination rates have been accompanied by renewed outbreaks of vaccine-preventable diseases, such as measles, despite widespread access to healthcare and vaccines (Centers for Disease Control and Prevention 2025a; Hellenic National Public Health Organization (EODY) 2025). Similar concerns extend to other immunization programs, and the Greek setting of our experiment reflects these broader trends, with uptake declining in patterns similar to those observed elsewhere in the OECD and the United States (Gountas et al. 2024; OECD 2025).

We implement a large-scale randomized experiment with more than 1,200 young adults in Greece. Treated participants receive an identical public-health message describing declining vaccination rates and explaining the role of vaccines in disease prevention and public-health protection. Participants assigned to the control group are exposed to a neutral message unrelated to vaccines. The design then varies two features of communication while holding the informational content fixed. First, the message is delivered either through a conventional two-dimensional video or through an immersive virtual reality (VR) environment. Second, the message is delivered by either a medical scientist or a popular social media influencer, with both male and female speakers in each category. This design allows us to separately examine the role of communication technology and messenger identity in shaping beliefs about vaccines and the healthcare system. We complement the survey outcomes with eye-tracking data recorded during the VR sessions and with two follow-up surveys, one administered approximately two months after the intervention and the second after at least six months had elapsed since the intervention for all participants.

The experiment delivers four main findings on how information generates persuasion. First, exposure to the public-health message increases trust in vaccines and trust in the healthcare system. The increase in trust in vaccines is especially meaningful, amounting to 28 percent of a standard deviation relative to baseline levels. Second, communication technology matters even when the information itself is unchanged: Delivering the message through immersive virtual reality produces belief changes approximately 11 percent of a standard deviation larger than standard video, consistent with greater engagement with the information. Third, persuasive effectiveness depends on the messenger and varies substantially across individual speakers. The two scientists generate similar effects, but the two influencers differ markedly, with the largest effect generated by the influencer with the strongest prior reputation. Finally, the effect on trust in vaccines shows remarkable persistence: It remains substantial two months after the intervention and is detectable and robust more than six months later. Initially hesitant participants also report an increased willingness to vaccinate at the two-month follow-up.

The larger effect of VR is consistent with immersive delivery generating greater engagement with the information. The mechanism evidence points to different sources of persuasive effectiveness across speakers. Scientists' effects are closely associated with perceived reliability and knowledgeability, which account for much of their advantage relative to influencers in the mediation analysis. The female influencer generates positive effects despite limited prior recognition, with personal appeal playing a prominent role. The male influencer's effect, by contrast, is not explained by contemporaneous evaluations but is partly associated with pre-existing familiarity and reputation. Differences in visual attention across speakers are imprecisely estimated and provide limited explanatory power. Social media personalities can therefore be persuasive despite lacking formal expertise, drawing instead on personal appeal and familiarity.

To explain these empirical results, we develop a model of belief formation. The model belongs to the class of behavioral theories in which how information is processed and remembered, rather than information alone, governs belief formation (Mullainathan, Schwartzstein, and Shleifer 2008; Bordalo, Gennaioli, and Shleifer 2020). The central insight is that persuasion is governed by three multiplicative forces: Engagement, Bayesian weight, and effective credibility. Immersive communication technology (VR) drives the first force; prior uncertainty governs the second; intrinsic source credibility governs the third; and parasocial familiarity shifts the third in a way that can equalize the average persuasive impact of experts and influencers. The model delivers closed-form expressions for belief revision and persistence that map directly onto the experimental treatments.

This paper first contributes to the literature on persuasion and belief formation. In the framework of DellaVigna and Gentzkow (2010), persuasion depends on exposure, how individuals process a message, and the credibility of its source; behavioral theories further emphasize selective processing and associative memory in shaping how beliefs respond to information (Mullainathan, Schwartzstein, and Shleifer 2008; Bordalo, Gennaioli, and Shleifer 2020). Existing evidence on source effects has emphasized credibility and perceived bias (Lupia and McCubbins 1998; Druckman 2001; Gentzkow and Shapiro 2006). Our design holds information content fixed while independently varying the messenger and communication medium, allowing us to study these dimensions of persuasion separately from message reach or diffusion. We show that the same information is more persuasive when delivered through an immersive medium, consistent with greater engagement with the message. At the same time, the effectiveness of different messengers reflects different sources of credibility. For scientists, reliability and knowledgeability appear central; for influencers, personal appeal and pre-existing familiarity and reputation can sustain persuasive effectiveness despite the absence of formal expertise. In particular, a messenger's prior reputation can compensate for a lack of perceived expertise. The model formalizes these mechanisms, with engagement, prior uncertainty, and effective credibility entering persuasion multiplicatively.

Second, we contribute to the applied literature on health communication. A large literature studies nudges and reminders (Milkman et al. 2021; Dai et al. 2021; Serra-Garcia and Szech 2023) and health-information provision (Dupas 2011; Ashworth, Thunström, Cherry, Newbold, and Finnoff 2021; Bartoš, Bauer, Cahlíková, and Chytilová 2022; Alsan and Eichmeyer 2024; Delavande, Shahab, Younas, and Zafar 2026). We provide, to our knowledge, the first causal evidence comparing social media influencers with medical experts as public-health messengers under identical informational content (Alatas, Chandrasekhar, Mobius, Olken, and Paladines 2024; Fahey et al. 2024; O'Brien, Ganjigunta, and Dhillon 2024; De Veirman, Cauberghe, and Hudders 2017).[1] We also provide one of the first applications of immersive virtual reality to public-health communication (Herrera, Bailenson, Weisz, Ogle, and Zaki 2018; de-Juan-Ripoll et al. 2018; Chen, Huang, McIntosh, Shepherd, and Zhao 2025; Andries, Bursztyn, Chaney, Djourelova, and Imas 2026; Lacle-Melendez, Silva-Medina, and Bacca-Acosta 2025), finding that immersion amplifies immediate belief change but shows little evidence of increasing its persistence. By re-contacting participants approximately two months after exposure—a horizon rarely observed in this literature (Alsan and Eichmeyer 2024; Bartoš, Bauer, Cahlíková, and Chytilová 2022)—we provide evidence on the durability of belief change: A substantial share of the initial increase in vaccine trust remains detectable at both follow-ups.

Choosing a messenger and choosing a communication technology are not logistical afterthoughts of an information campaign; they are persuasion decisions in their own right. In selecting a messenger, a public-health authority should weigh not only credentials but the audience's prior relationship with that messenger: Familiarity can substitute for expertise, and the most credentialed voice is not always the most persuasive one. The delivery technology is a second lever, one that operates independently of the messenger.

[1] Existing evidence on influencers in health communication largely concerns reach and diffusion rather than causal effects on beliefs: Alatas, Chandrasekhar, Mobius, Olken, and Paladines (2024) show that celebrity endorsement increases the diffusion of pro-vaccination content; Fahey et al. (2024) find that influencer-based outreach engages hard-to-reach populations in HIV-prevention research; and O'Brien, Ganjigunta, and Dhillon (2024) document wellness influencers amplifying vaccine-skeptical discourse during the COVID-19 pandemic. The trusted communicators in Delavande, Shahab, Younas, and Zafar (2026) are locally respected community members rather than social media personalities.

The optimal design of a campaign therefore depends on audience-specific priors, familiarity distributions, and the available communication technologies—not on content alone. In information environments where scientific facts compete with familiar online voices, these choices, as much as the facts themselves, can determine whether the facts are believed.

2 Background and Experimental Design

2.1 Background on Vaccination in Greece

Vaccination is among the most cost-effective public-health interventions for reducing infectious disease transmission and improving population health. The World Health Organization (WHO) estimates that global immunization efforts have saved at least 154 million lives over the last fifty years, and continues to identify routine vaccination as a core component of disease prevention (World Health Organization and UNICEF 2024; World Health Organization 2025). Despite this strong scientific consensus, vaccination coverage remains incomplete and has declined in a number of high-income countries in recent years (World Health Organization 2025; OECD 2025). Existing evidence shows that individuals often hold beliefs about vaccine safety and effectiveness that are weakly aligned with scientific evidence and that those beliefs can play an important role in vaccination decisions (Bartoš, Bauer, Cahlíková, and Chytilová 2022; Alsan and Eichmeyer 2024).

Greece, where we collected our data, fits squarely within this broader pattern of declining vaccination coverage and renewed epidemiological risk. Panel A of Figure A1 presents recent WHO estimates of first-dose measles vaccination coverage across selected countries and aggregates.[2] The figure shows a decline in coverage following the pandemic across the EU27 average, the OECD average, Australia, and Greece, with all series, including the United States, remaining below the 95 percent threshold generally associated with herd immunity. In the United States, Centers for Disease Control and Prevention (CDC) data report measles, mumps, and rubella (MMR) vaccination coverage of 92.5 percent among kindergartners in the 2024–2025 school year, again below the herd-immunity threshold (Centers for Disease Control and Prevention 2025b). Both Greece and the United States have experienced renewed measles outbreaks in recent years (Centers for Disease Control and Prevention 2025a; Hellenic National Public Health Organization (EODY) 2025).[3]

Human papillomavirus (HPV) vaccination also remains below target levels. Panel B of Figure A1 reports HPV vaccination coverage among girls by age 15 using WHO estimates for the United States, the OECD average, and the EU27 average, while the estimate for Greece is taken from Doulou et al. (2026), who reconstruct national coverage using prescription data. In 2024, coverage reached approximately 77 percent in the United States and 75 percent across the OECD, compared with 68 percent for the EU27 average and an estimated 64 percent in Greece. Although Greece extended eligibility for publicly funded HPV vaccination to both girls and boys in 2023, pandemic-related disruptions required substantial catch-up

[2] All figures and tables whose numbers start with A, B, or C can be found in the Online Appendix.

[3] In particular, the data available on the website of the CDC (Centers for Disease Control and Prevention 2025a) show that 2025 recorded the highest number of measles cases in the United States since 1992, while the number of confirmed cases reported in 2026 had already surpassed the 2025 total by August of that year.

efforts through 2024 (Gountas et al. 2024).

The evidence points to a common feature of vaccination in high-income settings: While vaccines are widely available, uptake remains lower than what would be desirable. Understanding the factors that shape vaccination decisions therefore remains important even in settings with well-developed healthcare systems. Among these factors, a growing body of evidence points to the role of beliefs and information shaping health behavior (Alsan and Eichmeyer 2024; Dupas 2011). This motivates a focus not only on the information individuals receive, but also on how that information is communicated and perceived. In this study, we examine how the identity of the messenger shapes the effectiveness of vaccine-related information.

The Greek setting provides a useful environment in which to study persuasion and belief updating in public-health communication. As we documented above, vaccination coverage remains below target levels, while social media constitute an important source of information among the young population from which our experimental sample is drawn. These features make the choice of who communicates vaccine-related information particularly relevant and motivate our experimental design, which examines how messenger identity and communication technology jointly shape attitudes toward vaccination.

2.2 Experimental Design

The experiment examines whether the effectiveness of information depends on both messenger identity and the medium through which it is delivered. For messenger identity, we contrast two prominent sources in contemporary communication environments: Medical professionals and social media influencers. This comparison captures a central tension between the perceived expertise associated with medical professionals and the parasocial trust associated with influencers. For the communication medium, we contrast a standard two-dimensional video with an immersive virtual reality (VR) experience.[4] Across both dimensions, the message content is held constant, which allows us to isolate how individuals update their beliefs and attitudes toward vaccination in response to these variations.

All treated participants are exposed to the same public-health message, which focuses on declining vaccination rates, the importance of HPV vaccination for younger cohorts, and the role of vaccination in protecting both individual and public health. The script was developed in consultation with a team of medical and clinical doctors to ensure that the final content was scientifically accurate, clinically appropriate, and communicated in a clear and accessible manner. The full text of the message (in English translation) is reported in Figure A2.[5] The message is delivered either by a medical doctor or a Greek social media influencer. For each messenger, we include both a male and a female speaker. All videos were recorded in the same, neutral, environment using a teleprompter, to minimize differences in both environment and message duration. After giving a short introduction of themselves, consisting of their name, surname, and occupation, speakers read an identical script to ensure that only the identity of the messenger varies across conditions.[6]

[4] The VR environment was developed using Unity, a real-time 3D development platform, by a dedicated software development team supported by a specialist VR programmer. The environment was piloted extensively prior to data collection to ensure stability, consistent rendering across headsets, and a seamless participant experience.

[5] Unlike Figure A1, which reports vaccination coverage for the first dose of the measles vaccine, the message focuses on uptake of the second dose.

[6] The speakers recorded the video in the same room at the Athens University of Economics and Business, where a professional

The public-health message was delivered through two distinct communication media: A standard two-dimensional video or an immersive VR experience. This design allows us to compare the effectiveness of delivering the same informational content through two different communication modalities while holding the message and speaker constant. The resulting contrast captures the overall effect of the VR implementation—including the immersive virtual environment and the broader delivery experience—relative to the embedded video condition, rather than isolating the independent effects of immersion (Herrera, Bailenson, Weisz, Ogle, and Zaki 2018). After recruiting participants through a centralized university email and flyers, we randomly assigned them to one of three groups: A control group, a video treatment group, or a VR treatment group. Within the video and VR treatment groups, the identity of the speaker is further randomized across the four profiles (doctor vs. influencer, male vs. female). To ensure clean identification, each experimental session included participants from a single treatment arm: Either the control group, the video group, or the VR group. We first assigned each session to one of the three arms and then allowed participants to book only sessions corresponding to their assigned treatment. As a result, participants could view and register only for sessions corresponding to their assigned treatment arm (Figure B2 provides a screenshot of the schedule structure for the second week of data collection). Thus, assignment to the information treatment (information vs. control) and the delivery medium (VR vs. video) varied across sessions. Within the information-treatment sessions, however, the identity of the information provider (scientist vs. influencer, and speaker gender) was randomized at the individual level.[7]

The survey begins with a baseline questionnaire that collects demographic characteristics, prior beliefs about vaccines and healthcare, and information on media consumption and exposure to social media influencers. After the baseline questions, participants in the video treatment group view the assigned video embedded within the survey platform. Participants assigned to the VR treatment are instructed to wear a VR headset and, after typing in their individual identifier, they enter the VR environment. More precisely, they find themselves in a virtual public park environment, sitting on a bench. This part of the experiment acts as a brief acclimatization period, allowing participants to adjust to the VR display and become familiar with the virtual environment and basic interactions before the information treatment begins.[8] This acclimatization sequence is identical for all VR participants regardless of treatment assignment. Participants are then instructed to pick up a phone located next to them, unlock it, and then play the video that appears on the screen. We designed these steps in order to make the experience feel as realistic as possible. The VR experience lasts approximately three to four minutes, with varying durations across the four speaker conditions (see Figure A3). Participants in the control group are shown a neutral video unrelated to health topics, with a duration comparable to that of the videos containing the public-health message.[9] Each session was supervised by three to four research assistants, drawn from a team of 17, who followed a common written protocol across treatment arms.

cameraman filmed the speakers reading the message from a teleprompter, using the same white background.

[7]Participation was incentivized with a €5 voucher redeemable at a local cafeteria on the university campus and entry into a raffle for various Apple devices. We provide a picture of one of the vouchers in Figure B3.

[8]To familiarize participants with the experiment-specific virtual environment, the acclimatization period was designed to reduce potential confounding arising from heterogeneous prior experience with VR technology. This ensured that any treatment effects were less likely to reflect differences in participants' familiarity with the medium itself.

[9]The control video consists of a short clip about London. It is intended to match the treatment videos in format and duration, while it remains unrelated to health or vaccination topics.

Following exposure to the treatment or the control video, all participants complete a post-treatment survey that re-measures beliefs about vaccines and healthcare, with treated participants also receiving questions to elicit perceptions of the speaker (e.g., credibility, knowledge, and charisma). We also collect measures of willingness to engage in discussions on a range of topics, which include vaccination and broader public issues. All participants are re-contacted approximately two months after attending the laboratory experiment to complete the preregistered follow-up survey, allowing us to assess the persistence of treatment effects over time. We subsequently conduct an additional, non-preregistered follow-up after all participants have been at least six months removed from the experiment. Because the experiment was conducted in two waves, this second follow-up occurs approximately six months after treatment for participants in the second wave and eight to nine months after treatment for those in the first wave.[10]

Overall, the design consists of one control condition and eight treatment cells, defined by the cross of four speaker identities and two communication modes. Because we keep the informational content constant across all treatment conditions, any differences in outcomes can be attributed to variation in the messenger or in the means of communication. We provide additional details on the experimental protocol and implementation in Appendix B.3. The experiment was pre-registered in the AEA RCT Registry before data collection began (Dioikitopoulos, Megalokonomou, Sartori, and Zenou 2025); Appendix B.1 reproduces the registered design and outcomes, maps them to the analyses in the paper, and discloses all deviations. In Figure A4, we report participation by day and session (i.e., time slot within the day). Session intensity was higher in the first wave (up to 11 sessions per day) and reduced in the February–March period (up to six sessions per day) to accommodate the university lab availability. In total, we scheduled 197 laboratory sessions (67 control sessions, 64 video sessions, and 66 VR sessions): 146 during the first wave and 51 during the second. Out of the 197, 23 were canceled due to low enrollment. In the first wave, eight of the scheduled sessions had no participants; in the second wave, 15 sessions had no participants. The daily median number of participants per session ranged between four and seven.

2.3 Defining an Influencer

One key aspect of our design concerns what distinguishes the two types of speakers: Experts (scientists) and social media influencers. While the former are relatively straightforward to define—we selected two highly reputable medical doctors (specialized in infectious diseases) within the Greek medical community—characterizing an influencer requires greater conceptual precision. Broadly speaking, an influencer is someone capable of shaping the attitudes or behaviors of others. However, such influence may arise through different channels, including expertise, authority, or public recognition. In this study, we focus specifically on the latter.

The persuasive power of public recognition has long been exploited in marketing. A large literature shows that celebrity endorsements can shape consumer attitudes, purchasing decisions, and brand perceptions (Knoll and Matthes 2017). Their influence is not confined to commercial settings—the Elvis Presley

[10]Participants are asked to provide their email address at the end of the baseline survey for the purpose of follow-up contact. Email addresses are stored separately from survey responses to preserve anonymity, and are used solely for the purpose of re-contacting participants for the follow-up survey.

episode described in the introduction is the classic example (Mawdsley 2016). The emergence of social media has transformed this model by creating a new class of public figures whose influence derives primarily from online content creation rather than traditional celebrity. These social media influencers have become central to modern marketing strategies because of their ability to affect consumer attitudes and behavior (De Veirman, Cauberghe, and Hudders 2017).

One explanation for this influence comes from the concept of parasocial relationships, introduced by Horton and Wohl (1956). Through repeated exposure, audiences develop one-sided relationships with media personalities, creating feelings of familiarity, trust, and emotional connection despite having no reciprocal interaction. Social media arguably strengthens this mechanism by providing more frequent, personalized, and interactive content, allowing influencers to become a regular presence in the daily lives of their followers.

Our objective, however, is not to compare scientists and influencers in their natural communication environments. In practice, influencers differ from experts not only in popularity but also in communication style, emotional expression, production quality, and audience engagement. We therefore standardize the message, script, recording environment, and production quality across speakers, so that these dimensions do not vary systematically across experimental conditions. This design substantially narrows the set of differences between speakers relative to real-world settings, allowing the estimated treatment effects to be interpreted as arising from the selected speakers rather than from differences in message content or production environment. While speakers may still differ in individual characteristics such as tone of voice, cadence, confidence, or charisma, these features are inherent to the messenger and cannot be manipulated in our design. To better understand their potential role, our survey elicits detailed post-treatment evaluations of each speaker, which guide our understanding of the mechanisms that may underlie the observed treatment effects.

After standardizing several aspects of delivery, such as the script and the recording environment, influencers primarily differ from scientists in the social capital they bring to the interaction. The two influencers selected for the experiment each have audiences exceeding 250,000 followers and a substantially larger public profile than the scientists. The scientists are highly reputable infectious disease specialists chosen specifically for their professional expertise while having virtually no media presence, making prior recognition by participants unlikely. Consequently, our comparison is designed to identify whether identical health information is more persuasive when delivered by these social media personalities with a broad online following or by an expert whose credibility derives from professional qualifications rather than public recognition.

3 Data and Construction of Variables

We collected data from young adults at the Athens University of Economics and Business in Greece through a laboratory experiment conducted in two waves, in December 2025 and between February and March 2026, followed by two online surveys administered approximately two months and at least six months after the intervention, respectively.

3.1 Data

The data come from four sources: A survey administered during the laboratory session, two online follow-up surveys, the first administered approximately two months after treatment and the second approximately six to nine months after treatment, and eye-tracking records from the VR headsets. The combined dataset covers 1,245 participants. Of these, 1,218 (97.8 percent) completed the laboratory survey. Complete English translations of both questionnaires, with a question-by-question breakdown, are provided in Appendix B.4.1 (laboratory survey), Appendix B.4.2 (first follow-up survey), and Appendix B.4.3 (second follow-up survey).

The laboratory survey opens with a baseline module that records demographic characteristics, media habits, and beliefs relevant to the experiment. The primary outcomes are elicited twice—before and after exposure to the information treatment or the control condition—allowing us to study belief updating after receiving (or not) the information treatment. Both outcome batteries use a 0–10 scale, where 0 means "not at all" and 10 "completely." The first battery measures Trust in Vaccines: Participants evaluate the extent to which vaccines are "safe for the health of the individual," "effective at preventing severe diseases and pandemics," "important for one's own health," and "important for public health." The second measures Trust in the Healthcare System: Participants report their trust in "conventional medicine and surgery," "healthcare professionals," and "public health authorities." The items appear in Block 3 (Appendix B.4.1, Block 3, Questions 23–24) and Block 6 (Appendix B.4.1, Block 6, Questions 32–33).[11]

For participants exposed to the public-health message, the post-treatment module also records perceptions of the speaker and the message. Participants provide an overall evaluation of each on a scale from −5 ("very negative") to 5 ("very positive"), and rate the speaker on seven characteristics—charismatic, persuasive, reliable, knowledgeable, sympathetic, attractive, and passionate—each on a 0–10 scale, where 0 means "not at all" and 10 "extremely." These ratings underlie the perceived speaker quality outcomes (Block 7, Questions 35–36, Appendix B.4.1, Block 7).

The follow-up surveys re-elicit the main trust outcomes and collect information on subsequent vaccination attitudes and behavior. The first follow-up was administered approximately two months after treatment for participants in both experimental waves, while the second was administered approximately six months after treatment for participants in the second wave and eight to nine months after treatment for those in the first wave. Both surveys record whether participants became more willing to vaccinate and whether they received additional vaccinations after the experimental session (both measured as yes/no responses).[12] Sixty-three percent of laboratory participants completed the first follow-up survey, while 26 percent completed the second follow-up survey.

For participants assigned to the VR treatment condition, the headsets' integrated eye-tracking allows us to observe where participants direct their visual attention throughout the experience. From these records

[11] Beyond the main outcomes, the post-treatment survey also measures participants' willingness to discuss a set of topics with friends and family. Participants are asked: "After completing this questionnaire, how likely are you to start discussions about the following topics?" Responses are collected on a 0–10 scale ranging from "not at all likely" to "very likely." Topics include vaccines, public health, social media, artificial intelligence, and politics. These questions correspond to Block 6, Question 34 in Appendix B.4.1, Block 6.

[12] The main trust outcomes correspond to Block 1, Questions 2–3 of Appendix B.4.2, Block 1 for the first follow-up and Appendix B.4.3, Block 1 for the second follow-up. Questions on willingness to vaccinate and realized vaccination behavior correspond to Questions 4–5 in the respective questionnaires.

we construct three measures: Overall exposure time, time spent looking at the speaker, and the share of gaze directed toward versus away from the speaker.[13] Eye-tracking records can be matched to survey responses for all 378 participants in the VR treatment arm who completed the laboratory survey.

3.2 Construction of the Primary Variables

Several variables in the survey are inherently multi-dimensional, so the analysis relies on summary measures rather than a large set of individual outcomes. Using these indices allows us to aggregate conceptually related survey items into summary measures of Trust in Vaccines and Trust in the Healthcare System (Eichengreen, Saka, and Aksoy 2024; Bazzi, Fiszbein, and Gebresilasse 2020). Throughout the analysis, we use these composite outcomes in place of the individual survey items.

For the two main outcome families, related to trust in vaccines and the healthcare system, we construct summary indices using principal component analysis (PCA). Specifically, we estimate a two-component PCA using pre-treatment responses only and apply an orthogonal varimax rotation to improve the interpretability of the resulting components. The resulting scoring rule is then applied to the pre-treatment, post-treatment, and follow-up responses to construct comparable indices across survey waves. The first component primarily captures vaccination-related attitudes and is interpreted as Trust in Vaccines, while the second primarily captures institutional trust and is interpreted as Trust in the Healthcare System.

Table A1 reports the rotated component coefficients together with the variance explained by each retained component and Cronbach's α for the corresponding sets of survey items. The two retained components explain approximately 78 percent of the total variation in the seven survey items: The vaccine component accounts for 49 percent of the variance, while the healthcare-system component accounts for 29 percent. Reliability is high throughout. Cronbach's α, which measures the internal consistency of a multi-item scale (ranging from 0 to 1, with higher values indicating greater reliability), exceeds 0.90 for the vaccine construct and ranges between 0.75 and 0.80 for the healthcare-system construct across all measurement points. The rotated coefficients display a clear simple structure. Vaccination-related items are primarily associated with the vaccine component and receive negligible weight in the healthcare-system component, whereas institutional trust items are primarily associated with the healthcare-system component and receive negligible weight in the vaccine component. This separation indicates that the two retained components capture distinct dimensions of trust rather than a single underlying construct.

We also construct an index of Perceived Speaker Quality using PCA on the seven post-treatment evaluations of the speaker: charismatic, persuasive, reliable, knowledgeable, sympathetic, attractive, and passionate. Table A2 reports the component weights for the first principal component. All seven items receive positive and relatively even weights on this component, with weights ranging from 0.34 to 0.40, indicating that no single characteristic disproportionately drives the index. The first principal component explains 63 percent of the total variation, and the seven items exhibit high internal consistency, with a Cronbach's α of 0.90.

[13]The VR sessions were conducted using Meta Quest Pro headsets, which include integrated eye-tracking sensors capable of recording users' gaze direction within the virtual environment. The device combines immersive VR capabilities with real-time tracking of eye movements, which allows us to construct measures of visual attention during the experimental experience.

4 Sample Characteristics

This section describes the characteristics of the participant sample, which consists of 1,245 individuals. Table 1 summarizes the main demographic variables, which we collect in the first part of the survey. Overall, non-completion rates are low, with 1,218 participants (97.8 percent) completing the laboratory survey. Completion rates are 98.8 percent in the control group, 99.1 percent in the video group, and 95.5 percent in the VR group.[14] Accordingly, the descriptive evidence presented below is based on participants who completed the laboratory survey. The sample consists of young adults enrolled at university, with an average age of approximately 20 years. The gender composition is balanced (51 percent female), and the vast majority of participants (around 95 percent) are undergraduate students. Treatment assignment is broadly balanced, with similar shares of participants allocated to the control and video conditions (34 percent) and a slightly smaller share assigned to the VR condition (31 percent).[15]

A central feature of the sample is the prominence of social media in participants' information environment. We first examine whether participants know and trust the male and the female influencers. The upper panel in Table A3 reports those statistics. The male influencer featured in the experiment is substantially more recognized than the female influencer, with average familiarity scores of 3.24 and 1.84, respectively, on a five-point scale.[16] This asymmetry provides useful variation for studying how prior familiarity interacts with message reception. In terms of participants' perceptions of the influencers, we observe a modest difference in overall evaluations between the male and female influencer, with the former being perceived more positively, although average opinions of both are slightly negative.[17]

As shown in the lower panel of Table A3, most respondents report substantial time spent online. Around 40 percent of the participants spend between two and three hours per day on social media, and roughly 70 percent report spending more than two hours daily. Participants indicate that they use social media primarily to maintain social connections or for entertainment—in response to the question "For which of the following purposes do you primarily use social media?"—while only a relatively small share (around 10 percent) report using it mainly to follow news and current events or professional networking. Nevertheless, responses to the subsequent survey question reveal that social media constitutes the dominant source of information for most participants. In particular, more than 80 percent of respondents, when asked "Where do you usually get most of your news about what is happening in the world?", identify social media

[14] Completion rates differ significantly across treatment arms (joint $p = 0.008$), driven by a 3.4-percentage-point lower completion rate in the VR group relative to the control group ($p = 0.010$); completion does not differ between the video and control groups ($p = 0.760$).

[15] Differences in the number of participants initially assigned across arms reflect ex ante laboratory scheduling parameters, as VR sessions required specialized equipment, longer time slots, and dedicated staff.

[16] Familiarity is measured in Question 21 of the survey (Appendix B.4.1, Block 2). Participants were presented with a list of Greek content creators and asked, "How familiar are you with them?" Responses were recorded on a five-point scale ranging from 1 ("I do not know them at all") to 5 ("Extremely familiar"). The list included both influencers who appeared as speakers in the experiment and other influencers who did not, so as not to reveal in advance which individuals participants might subsequently encounter. The names and social-media handles of these content creators are anonymized in the version of the questionnaire reported in Appendix B.4.1.

[17] Overall evaluations are measured in Question 22 of the survey (Appendix B.4.1, Block 2). Participants were presented with the same list of Greek content creators as in Question 21 and asked, "Thinking about the content produced by each of the following creators, what is your overall opinion of them?" Responses were collected using a slider ranging from -2 ("Very Negative") to 2 ("Very Positive"). As in Question 21, the names and social-media handles of these content creators are anonymized in the version of the questionnaire reported in Appendix B.4.1.

platforms as their primary source of news.

We then turn to beliefs about vaccines and the healthcare system measured at baseline. Table A4 reports these measures in their original survey scale, ranging from 0 ("not at all") to 10 ("completely"). At baseline, participants express generally favorable views toward vaccines and scientific institutions, assigning relatively high scores across all four vaccine-related dimensions: Effectiveness, safety, importance for personal health, and importance for public health. Average responses range from 7.07 to 7.99 on a 0–10 scale, with the strongest support observed for the role of vaccines in protecting public health. Trust in doctors and modern medicine exceeds 7.5 out of 10 on average. In contrast, trust in public health authorities is noticeably lower, just below 6 points. Consistent with this pattern, only a small fraction of respondents report that they would automatically refuse a vaccine recommendation from health authorities (3 percent). Instead, most participants indicate that vaccine acceptance depends either on the specific vaccine (58 percent) or on their doctor's recommendation (20 percent). The reliance on a doctor's recommendation for a substantial subset of respondents, alongside the broader differences in institutional trust, suggests that messenger identity and source credibility play a meaningful role in how individuals respond to health information.

Additional insights come from the eye-tracking data collected during the VR sessions. Table A5 reports descriptive measures of participants' behavior within the VR environment, including total exposure time, time spent watching the speaker once the message begins, and the number of seconds during which participants keep their gaze directed toward vs. away from the speaker. The difference between total time in the VR environment and time spent watching the speaker primarily reflects the acclimatization exercise completed before the message, as well as the time participants took to pick up the virtual phone and initiate the video.[18] The first panel reports overall descriptive statistics, while Panels A through D provide separate breakdowns by speaker profile. Figure A3 complements these descriptives by presenting the median values of the same measures graphically across speakers. Several patterns emerge. First, message duration differs substantially across speakers. Influencers deliver the message more quickly than scientists: Average time spent watching the speaker is below 130 seconds for both influencers (Panels A and B of Table A5), compared with approximately 150 seconds for the female scientist and close to 138 seconds for the male scientist (Panels C and D). Second, despite these shorter exposure times, influencers attract greater visual attention from participants. Both Table A5 and Figure A3 show that participants spend a larger share of time looking directly at influencer speakers rather than at the surrounding virtual environment. A comparison across speakers further highlights how exposure length and visual attention need not coincide. In absolute terms, the male influencer and female scientist receive the highest average total gaze time, at approximately 56 seconds each. However, differences in message duration make gaze time difficult to compare directly across speakers. When gaze time is benchmarked against the duration of the message delivery—as illustrated by the median patterns in Figure A3—participants spend approximately 44 percent

[18] Participants could technically pause the video once it had started, although this occurred for only five participants in the full VR sample. These isolated pauses account for the unusually high maximum values of time spent watching the speaker reported in Table A5, which can therefore exceed the duration of the corresponding prerecorded video. For these participants, the additional time generated by the pause remains part of the total observation window used when calculating the share of time spent looking at the speaker.

of their exposure looking at the male influencer, compared with 38 percent for the female influencer, 37 percent for the female scientist, and 32 percent for the male scientist. Overall, the longer duration of the scientists' messages does not translate into greater visual attention, highlighting the distinction between exposure length and sustained attention.

5 Estimation Strategy

This section outlines the empirical specifications used to estimate treatment effects and clarifies the objects of identification. We first present the main specification, which captures the overall effect of exposure to the information treatment and its components. We then consider a set of nested specifications that isolate specific margins of interest. Finally, we report balancing exercises.

5.1 Treatment Dimensions and Estimation

The experimental design varies along four dimensions. Participants may receive the information treatment (IT), defined as exposure to the public-health message, or be assigned to the control group. Conditional on receiving the information treatment, three additional dimensions are randomized: (i) The communication medium (VR vs. standard video), (ii) the identity of the speaker (scientist vs. influencer), and (iii) the gender of the speaker (male vs. female).

We estimate separate specifications for the distinct experimental comparisons. We first pool all information-treatment arms to estimate the average effect of receiving the public-health message relative to the control condition:

$$Y_{is} = \alpha + \beta\,\mathrm{IT}_s + \rho Y^0_{is} + \theta X_i + \delta_{\mathrm{pref}(i)} + \delta_{\mathrm{day}(s)} + \delta_{\mathrm{order}(s)} + \epsilon_{is}, \tag{1}$$

where Y_{is} denotes the post-treatment outcome of participant i in session s, and Y^0_{is} is the corresponding pre-treatment (baseline) measure. The binary indicator IT_s takes the value one if the participant is exposed to the public-health message (VR or video) and zero in the control group. The main treatment effect is captured by β, which shows the overall effect of receiving the information treatment relative to the control condition.

The vector X_i includes participant characteristics including demographic background, socioeconomic status, political orientation, health-related behaviors, and social connectedness. We include a set of categorical indicators for participant gender (male/female/other or undisclosed), study level (undergraduate/graduate/postgraduate or other), parental education for both father and mother (high school/professional qualification/bachelor's degree/master's degree/doctorate), political ideology (left/center-left/center/center-right/right/no interest in politics), smoking frequency (never/rarely/often/every day), and exercise frequency (never/no more than once a week/more than once a week/every day or almost every day), and propensity to vaccinate (as measured using participants' stated willingness to accept a vaccine recommended by a public health authority: Yes/no/it depends on the vaccine/it depends on my doctor's recommendation). We additionally control for social connectedness, captured by both the number of close friends reported by the participant and the number of participants who reported the participant as one of their close friends, as well as session participation intensity, measured by the total number of participants within the session. The specification also includes fixed effects for prefecture of origin $\delta_{\text{pref}(i)}$, session day $\delta_{\text{day}(s)}$, and the order of the session throughout the day $\delta_{\text{order}(s)}$, in order to capture unobserved heterogeneity across geographic areas, days of the week, and times of day. Standard errors are clustered at the session level.

Medium Effects (VR vs. Video). To isolate the effect of the communication medium, we restrict the sample to treated participants ($\text{IT}_s = 1$) and estimate:

$$\begin{aligned} Y_{is} = \alpha + \gamma\, \text{VR}_s + \kappa\, \text{SC}_{is} + \lambda(\text{SC}_{is} \times \text{Male}_{is}) + \mu\, \text{Male}_{is} \\ + \rho Y_{is}^0 + \theta X_i + \delta_{\text{pref}(i)} + \delta_{\text{day}(s)} + \delta_{\text{order}(s)} + \epsilon_{is}, \end{aligned} \quad \text{for IT}_s = 1. \tag{2}$$

In this specification, VR_s equals one if the public-health message is delivered through virtual reality and zero if it is delivered through a standard video. The indicator SC_{is} equals one when the speaker is a scientist and zero when the speaker is an influencer, while $Male_{is}$ equals one for male speakers and zero for female speakers. The coefficient γ identifies the causal effect of delivering the message through VR rather than video, holding speaker identity constant.

Speaker Identity (Scientist vs. Influencer). We again restrict our attention to treated participants ($\text{IT}_s = 1$) and estimate:

$$\begin{aligned} Y_{is} = \alpha + \kappa\, \text{SC}_{is} + \mu\, \text{Male}_{is} + \gamma\, \text{VR}_s + \phi\, (\text{VR}_s \times \text{Male}_{is}) \\ + \rho Y_{is}^0 + \theta X_i + \delta_{\text{pref}(i)} + \delta_{\text{day}(s)} + \delta_{\text{order}(s)} + \epsilon_{is}, \end{aligned} \quad \text{for IT}_s = 1. \tag{3}$$

Here, κ captures the average difference between scientist and influencer speakers, controlling for the communication medium. The coefficient μ captures the difference between male and female speakers under video delivery, and ϕ allows this difference to vary with the communication medium.

Full Speaker Characterization. Finally, to characterize the differences across the four speaker profiles, we estimate the same specification as in equation (2), with the sample remaining restricted to treated participants ($\text{IT}_s = 1$). Now we focus on the coefficients associated with speaker identity and gender. Table 4 reports the equivalent reparametrization of this specification with three profile indicators (male influencer, female scientist, and male scientist), whose coefficients equal μ, κ, and $\kappa + \mu + \lambda$, respectively. The four profiles correspond to the interaction between speaker identity (scientist or influencer) and gender (male or female), with the female influencer as the omitted category. The coefficient μ captures the difference between the male and female influencer, while κ captures the difference between the female scientist and the female influencer. The interaction term λ allows the effect of speaker gender to differ between scientists and influencers. Together, these coefficients allow us to compare the four speaker profiles, conditional on the communication medium and the other covariates.

Across all specifications, the identifying variation comes from the randomized assignment of treatment conditions. Equation (1) identifies the average effect of exposure to the public-health message relative to the control condition, pooling across communication media and speaker identities. Equation (2) identifies the effect of delivering the message through immersive virtual reality rather than a standard video, holding the informational content and speaker characteristics constant. Equation (3) instead identifies differences attributable to messenger identity, comparing scientists and influencers, while the speaker identity and gender coefficients in equation (2) allow us to compare the four individual speaker profiles, holding the communication medium fixed.

5.2 Balance and Randomization Checks

Our identification strategy relies on the random assignment of the treatment across groups. This requires that participants' predetermined characteristics are balanced across treatment arms. The balancing regressions take the following form:

$$T_{ig} = \alpha_g + \theta_g\, X_i + \delta_{\text{pref}(i)} + \delta_{\text{day}(s)} + \delta_{\text{order}(s)} + \varepsilon^g_{is}, \tag{4}$$

where T_{ig} is an indicator for treatment assignment in comparison g (e.g., information treatment vs. control, VR vs. video, or scientist vs. influencer). X_i denotes a single predetermined participant characteristic. We estimate equation (4) separately for each baseline variable, including demographics (gender, age, study level, parental education, political ideology), habits (smoking behavior, exercise frequency, social media use), and baseline beliefs (trust in vaccines and the healthcare system). Our main specification includes prefecture, session-day, and session-timing fixed effects, and standard errors are clustered at the session level.

Table 2 provides a summary of the balancing exercises. Under columns (1)–(3), described as "Group Assignment," we report the share of statistically significant differences across baseline variables when comparing the three main groups (Control, Video, and VR). Out of 42 tests per group, only a small number are significant at conventional thresholds. The VR group exhibits somewhat more imbalance, with four coefficients significant at the 5 percent level, while no systematic pattern emerges across the other groups. Under columns (4)–(6), described as "Treatment Assignment," we report analogous summaries for the key pairwise comparisons underlying our design: information treatment vs. control (4), VR vs. video (5), and scientist vs. influencer (6). The information-treatment comparison is particularly well balanced: none of the coefficients are significant at the 1 percentor 5 percent levels, and four are significant at the 10 percent level. The VR-versus-video comparison shows somewhat less balance, with five of 42 coefficients significant at the 5 percent level. The scientist-versus-influencer comparison also contains a small number of significant differences, including one at the 1 percent level. Given the number of tests conducted, some differences are expected by chance. Tables A6–A9 report the full set of balancing regressions, including unadjusted comparisons that do not condition on session timing. These show that the few significant differences are scattered and not concentrated in any specific group of variables. Importantly, baseline beliefs about vaccines and the healthcare system are well balanced across treatment arms. Some isolated differences arise in behavioral variables such as smoking and exercise, but these are not systematic and do not suggest meaningful imbalance.[19]

Overall, the unadjusted baseline balance and the stability of our estimates across specifications are consistent with successful randomization. Observable characteristics are comparable across treatment conditions, which supports the causal interpretation of the effects.

6 Empirical Results

In this section, we provide empirical evidence on how individuals respond to information. We focus on three related questions. First, does exposure to information affect beliefs? Second, does the effectiveness of information depend on the medium through which it is delivered? Third, does the identity of the messenger influence how information is received? Throughout, the organizing question is whether persuasion tracks the perceived credibility of the source, as credibility-based accounts of source effects would suggest, or whether other features of the messenger become relevant.

We begin by estimating the effect of exposure to the information treatment on our primary outcomes—trust in vaccines and trust in the healthcare system—establishing whether information alone is sufficient to shift beliefs. We then study whether these effects differ between virtual reality and standard video delivery, examining the role of the communication medium. Next, we turn to the identity of the messenger, comparing scientist and influencer speakers, before providing a more granular analysis across individual speaker profiles. We study both immediate effects and persistence over time using data from two follow-up surveys. We further explore heterogeneity in responses across participant characteristics. We also ask whether the intervention extends beyond participants' own beliefs to their willingness to discuss these topics with others. Section 7 then draws on post-treatment evaluations, eye-tracking records, and pre-treatment familiarity measures to examine the mechanisms underlying these effects.

6.1 Information, Medium, and Speaker Identity

We begin with the first question: Does information matter? In particular, can providing information on vaccines shift individuals' beliefs? We start with a graphical analysis of our main outcomes, summarized by two latent constructs: Trust in Vaccines and Trust in the Healthcare System (see Section 3.2). Panel A and Panel C of Figure 1 show the distribution of these outcomes before exposure to the information treatment. Consistent with the balancing evidence presented earlier, there are no systematic differences across groups at baseline. Panel B and Panel D of Figure 1 report post-treatment outcomes. A clear shift emerges for participants exposed to the information treatment, with the largest changes observed in the VR condition relative to standard video. The regression estimates confirm the graphical evidence. Panel A of Table 3 reports coefficients representing the effect of receiving the information treatment, based on equation (1). Columns (1) and (2) focus on trust in vaccines, columns (3) and (4) on trust in the healthcare system, and

[19] Although realized session timing is technically post-assignment for the information and communication-medium comparisons, Appendix Tables A6–A9 provide little evidence of imbalance in observed participant characteristics. Consistent with this, the estimated treatment effects change only negligibly when session-level controls are included, providing further reassurance that conditioning on realized session characteristics does not meaningfully affect the experimental comparisons. For the scientist-versus-influencer comparison, assignment occurred within sessions, making session timing strictly pre-assignment.

the final columns report a placebo outcome—response time, measured as the number of seconds required to complete the response block—which we would not expect to respond systematically to the treatment.[20]

Exposure to the information treatment leads to a statistically significant increase in trust in vaccines of approximately 0.28 standard deviations (significant at the 1 percent level). The inclusion of controls (comparing columns (1) and (2)) leaves the coefficient essentially unchanged, suggesting that the effect is not driven by compositional differences across participants. The effect on trust in the healthcare system is smaller but still economically and statistically significant, at 0.21 standard deviations. To provide a sense of the magnitude of these effects in the original survey scales, Appendix Table A10 reports the corresponding treatment effects for each of the individual items underlying the two indices. For the vaccination-related measures, the treatment increases responses by roughly one-half of a point (around 5 percent of the scale range) across all four items, indicating that the effect is broad-based rather than driven by a single dimension. For the healthcare-system measures, the estimated increases range from roughly one-fifth to three-fifths of a point (approximately 2–6 percent of the scale range), with the largest improvement observed for trust in public health authorities. Because the pooled information-treatment estimate captures the average effect of the public-health campaign across both delivery modes, we also directly compare the video treatment with the video control in Table A11. Holding the delivery medium fixed allows us to evaluate the effect of the informational campaign itself. This same-medium comparison yields positive and statistically significant effects on both outcomes, with estimated increases of 0.209 points for trust in vaccines and 0.175 points for trust in the healthcare system (approximately 0.22 and 0.16 standard deviations, respectively). Thus, the campaign successfully shifts beliefs even in the absence of the broader VR delivery package.

Having established that information matters, we next turn to the role of the communication medium. A growing literature suggests that immersive experiences, such as virtual reality, can have stronger behavioral and attitudinal effects by increasing engagement and emotional involvement (Herrera, Bailenson, Weisz, Ogle, and Zaki 2018; Lacle-Melendez, Silva-Medina, and Bacca-Acosta 2025; Andries, Bursztyn, Chaney, Djourelova, and Imas 2026). We test whether this holds in our setting.[21] Panel B of Table 3 isolates the role of the communication medium by comparing VR and standard video among treated participants. We find that VR delivery leads to larger effects on both outcomes. For trust in vaccines, the estimated effect is approximately 0.11 standard deviations (0.105 index points) in our preferred specification, while for trust in the healthcare system it is around 0.10 standard deviations. These effects are statistically significant and robust to the inclusion of controls. We find no significant effects on the placebo outcome.[22]

[20] Response time is intended as a placebo outcome rather than a strict falsification test. In principle, we cannot rule out that exposure to the message may affect subsequent response speed through greater engagement or attention. We therefore interpret these estimates cautiously.

[21] Strictly speaking, the estimated "VR effect" should not be interpreted as the causal effect of VR immersion alone. The VR treatment is a bundled intervention that combines headset use, a virtual environment, an acclimatization task, interaction with virtual objects, researcher assistance, novelty, and a different overall time spent in the experimental environment. The estimand therefore captures the effect of the complete VR implementation package relative to viewing the same content in an embedded survey video, rather than the isolated effect of immersion. This challenge is common to experimental evaluations of VR interventions, as immersion is typically delivered alongside a broader set of experiential and procedural features.

[22] Table A12 reports treatment effects at the level of individual survey items underlying the main outcome indices for the VR vs. video comparison. While the magnitude of the effects varies across components, estimates are generally positive across

We then examine the broader role of the identity of the messenger. Panel C of Table 3 compares scientist and influencer speakers, holding the informational content constant. The estimated differences are small and not statistically significant across both main outcomes, although the point estimates are consistently positive for scientists. This indicates that, on average, expert messengers do not outperform influencers in shifting beliefs in this context.[23]

A more nuanced picture emerges when we examine heterogeneity across individual speakers. Table 4 reports a full breakdown by speaker identity, with the female influencer as the reference group. Relative to this baseline, all other speakers generate positive effects on trust in vaccines, and these effects are statistically significant. The largest effect is observed for the male influencer, while both scientists also generate sizable and comparable effects, with the male scientist slightly outperforming the female scientist. While average differences between scientists and influencers are limited, substantial heterogeneity exists across individual speakers. This highlights the importance of moving beyond broad categories of messenger type and considering more granular dimensions of speaker characteristics.[24]

Finally, we address the fact that our main outcomes are latent constructs estimated from multiple survey items rather than directly observed variables. In our baseline analysis, the principal component scores are treated as fixed outcomes. However, because the PCA weights used to construct these indices are themselves estimated from the data, conventional standard errors may not capture the additional sampling variation introduced by this first-step estimation (Pagan 1984; Murphy and Topel 2002). We therefore implement a session-level bootstrap procedure that replicates the full estimation process: Within each bootstrap draw, we re-estimate the principal components, reconstruct the latent indices, and re-run the treatment-effect regressions. This approach preserves the experimental design while allowing inference to reflect both uncertainty in the construction of the latent outcomes and uncertainty in the estimated treatment effects. We show in Table A17 that the resulting bootstrap standard errors are very similar to our baseline estimates for the information-treatment comparison. For the communication-medium and individual-speaker comparisons, standard errors are somewhat larger. While the point estimates and qualitative patterns remain identical, this yields less precise inference—most notably, the VR effect on trust in the healthcare system is no longer statistically significant at conventional levels. Overall, our directional and economic conclusions remain, even after accounting for the uncertainty arising from the construction of the PCA indices.

Our results indicate that what is communicated matters, but so does how, and by whom. Both scientists and social media influencers are able to shift beliefs, while delivering the same information through VR—a medium that previous research suggests can heighten engagement (Herrera, Bailenson, Weisz, Ogle, and Zaki 2018; Andries, Bursztyn, Chaney, Djourelova, and Imas 2026; Chen, Huang, McIntosh, Shepherd, and Zhao 2025)—produces larger effects than standard video. At the same time, persuasiveness varies

items, consistent with the aggregate results.

[23]Table A13 reports treatment effects at the level of individual survey items underlying the main outcome indices for the scientist vs. influencer comparison. These specifications mirror the main results but replace the composite measures with their constituent components. Differences between scientists and influencers are small and not statistically significant across items.

[24]Tables A11, A14, and A15 report analogous item-level estimates where the baseline group is the control condition, rather than comparisons across treated categories. These tables provide a complementary view of the results and confirm that the main patterns are not driven by the aggregation of outcomes into indices. Table A16 shows that the results are robust when the latent outcomes are constructed using common factor analysis rather than PCA.

across individual speakers. The two scientists generate remarkably similar effects, whereas the effects of the two influencers differ substantially: The male influencer produces an effect that is larger in magnitude than those of either scientist, although not statistically distinguishable from them, while the female influencer is considerably less persuasive. In the next section, we examine whether these differences can be traced to specific speaker characteristics.

6.2 Persistence of the Effects

We next examine the persistence of the treatment effects using data from the two follow-up surveys. We first re-contacted participants approximately two months after the intervention and re-measured beliefs about vaccines and the healthcare system, as well as vaccination intentions and self-reported vaccination behavior. We subsequently conducted a second follow-up after all participants had been at least six months removed from the intervention. As reported in Section 3.1, 63 percent of laboratory participants completed the first follow-up survey, while 26 percent completed the second. As a preliminary step, Table A18 tests whether treatment assignment predicts participation in either follow-up survey, allowing us to assess selective attrition. Panel A reports results for the first follow-up and Panel B for the second. We find no evidence that treatment status affects participation in either survey, suggesting that attrition is not systematically related to treatment assignment.

Table 5 reports the effects of exposure to the information treatment, relative to the control group, at the first follow-up (Panel A) and the second follow-up (Panel B). To assess the sensitivity of these estimates to selective attrition, we also report Lee bounds (D. S. Lee 2009) for each coefficient (attrition is approximately 37 percent at the first follow-up and 74 percent at the second).[25] Panel A shows that the effect on trust in vaccines persists two months after the intervention. The estimated effect is 0.240, and the Lee bounds remain positive and of similar magnitude, ranging from 0.213 to 0.264. Relative to the immediate post-treatment effect estimated among participants observed at the first follow-up, approximately 90 percent of the initial effect remains. Panel B provides evidence that the effect remains detectable considerably further from the intervention. Once all participants had been at least six months removed from treatment, the estimated effect on trust in vaccines is 0.184 standard deviations and statistically significant at the 5 percent level. Although smaller than at the first follow-up, it represents approximately 62 percent of the immediate effect estimated for this sample. The Lee bounds range from 0.121 to 0.214 and remain entirely above zero, providing evidence that the persistence of the effect is robust to differential attrition under the Lee-bounds assumptions.[26] The pattern is different for trust in the healthcare system. The estimated effect is substantially smaller and statistically indistinguishable from zero at the first follow-up. At the

[25] Lee bounds (D. S. Lee 2009) provide lower and upper bounds on treatment effects in the presence of potentially selective attrition, under a monotonicity assumption on selection into the observed follow-up sample. When follow-up rates differ across treatment groups, the outcome distribution of the group with the higher follow-up rate is trimmed so that its effective follow-up rate matches that of the group with the lower rate. Trimming observations from opposite tails of the outcome distribution yields worst-case lower and upper bounds on the treatment effect. The resulting interval therefore captures the range of treatment effects consistent with selective attrition under the monotonicity assumption.

[26] To assess persistence holding sample composition fixed, we compare each follow-up estimate with the immediate post-treatment effect estimated on the corresponding follow-up sample. The immediate post-treatment coefficient on trust in vaccines is 0.268 among participants observed at the first follow-up and 0.297 among those observed at the second follow-up. The corresponding follow-up coefficients of 0.240 and 0.184 therefore represent approximately 90 and 62 percent of the respective immediate effects.

second follow-up, the point estimate is 0.174 standard deviations and statistically significant at the 10 percent level, with Lee bounds ranging from 0.127 to 0.234. We interpret this latter result cautiously, particularly given the substantially lower participation rate in the second follow-up. More generally, while the intervention produces a persistent shift in trust in vaccines, we find considerably weaker evidence of downstream behavioral responses. Effects on self-reported increases in willingness to vaccinate and subsequent vaccination are generally small and statistically indistinguishable from zero. Thus, persistent changes in reported trust do not necessarily translate into measurable changes in vaccination behavior over the time horizon we observe.

Because much of the potential policy value of vaccination campaigns lies in reaching individuals who are initially hesitant, Table A19 examines persistence separately by participants' baseline propensity to vaccinate. Panel A focuses on participants with a low baseline propensity to vaccinate, defined as those who reported that their willingness to receive a recommended vaccine would depend on the specific vaccine or that they would never accept a recommended vaccine. The treatment effect on trust in vaccines is substantially larger than in the full sample, corresponding to approximately 34 percent of the control group's standard deviation at follow-up. The estimated effect on trust in the healthcare system is also larger, although it remains statistically insignificant. More importantly, exposure to the information treatment increases the probability that participants report having become more willing to vaccinate by roughly 9 percentage points, corresponding to nearly a 30 percent increase relative to the control group. The Lee bounds for this effect range from 0.093 to 0.126, indicating that the increase in willingness among initially hesitant participants is extremely robust to selective attrition. We do not, however, detect a statistically significant effect on actual vaccination uptake over the two-month follow-up period. Panel B performs the complementary exercise for participants with a high baseline propensity to vaccinate, defined as those who reported that they would always receive a recommended vaccine or would do so if advised by their doctor. The estimated effect on trust in vaccines is much smaller than among participants with a low baseline propensity to vaccinate, and we find no evidence of changes in willingness to vaccinate or subsequent vaccination uptake. Overall, these result suggest that the clearest persistent attitudinal responses, as well as the increase in reported willingness to vaccinate, are concentrated among participants who were initially more hesitant.

In Table A20, we examine whether persistence differs across the communication medium (Panel A) and messenger identity (Panel B). Most differences along these dimensions are small and statistically insignificant. One exception is a five-percentage-point increase in self-reported vaccination uptake for the VR arm relative to the video arm, statistically significant at the 1 percent level. This result, however, is sensitive to selective attrition: The corresponding Lee bounds range from -0.015 to 0.057 and therefore include zero. Panel B similarly provides no evidence of systematic differences between scientists and influencers at follow-up. For example, the Lee bounds for the difference in trust in vaccines range from -0.049 to 0.048.

These follow-up results speak to a natural concern with laboratory-measured attitudes: that post-treatment responses partly reflect experimenter demand rather than genuine belief change. Three features of the evidence are difficult to reconcile with a demand-driven interpretation. First, the follow-up surveys were completed online, outside the experimental environment, approximately two months after the laboratory session for the first follow-up and at least six months after the intervention for the second. Despite this temporal and physical distance from the experiment, a substantial share of the initial effect on trust in vaccines persists: nearly 90 percent at the first follow-up and around 60 percent at the second. Second, demand offers no obvious account of the differential persistence across outcomes: a participant inclined to satisfy the experimenter should inflate trust in vaccines and in the healthcare system alike, whereas persistence is concentrated precisely in the outcome targeted by the message's content—a pattern more consistent with consolidation of the treated belief than with response behavior. Third, among initially hesitant participants, the first follow-up also reveals an increase in reported willingness to vaccinate, measured two months after and outside the experimental setting. This pattern also connects to recent evidence that the form in which information is delivered governs how durably it is encoded: Graeber, Roth, and Zimmermann (2024) show that narratives persist in memory while statistics fade, and our finding that the intervention's effects outlast the session is consistent with encoding depth—rather than momentary compliance—driving measured persistence.

Overall, the follow-up evidence suggests that the information intervention generates a lasting increase in trust in vaccines and that this persistence is difficult to attribute solely to experimenter demand. The effect remains detectable at the first follow-up and, although attenuated, is still present at the second follow-up, more than six months later. The responses are particularly pronounced among individuals with a lower initial propensity to vaccinate, for whom the improvement in beliefs is also accompanied by an increase in reported willingness to vaccinate at the first follow-up. We find no robust evidence, however, that these changes translate into actual vaccination behavior over the period we observe.

6.3 Heterogeneity in Treatment Effects

We next examine whether treatment effects vary across participant characteristics and pre-existing beliefs. We consider heterogeneity along standard demographic dimensions (gender, parental education, field of study, political orientation, and social network size), as well as along information-environment variables, including social media use, type of use, reliance on social media as a source of news, trust in social media, and baseline beliefs about vaccines and the healthcare system. We conduct these exercises across the main treatment dimensions.

Figure 2 presents a graphical analysis of heterogeneity in the effect of the information treatment (treatment vs. control) for our two main outcomes. Several patterns emerge. First, women appear more responsive to the information treatment than men. Second, the treatment effect is stronger among participants who report being uninterested in politics. In contrast, we find no systematic differences by parental education, field of study, or network size when splitting the sample at the median. More tentative differences emerge when considering the information environment and baseline beliefs. Some evidence suggests that treatment effects vary with pre-treatment patterns of social media use, although the patterns differ across outcomes. For trust in the healthcare system, the estimated effect is larger among participants who primarily use social media for entertainment. For trust in vaccines, the estimated effect is larger among those who report social media as their main source of news. We do not observe meaningful differences when splitting by time spent on social media or by trust in social media. A more consistent pattern emerges with respect to baseline beliefs. The treatment tends to be more effective among individuals with lower initial levels of trust. For trust in vaccines, the estimated effect declines across terciles of the baseline distribution, while for trust in the healthcare system the largest estimates are observed among participants in the bottom tercile. While bounded survey scales inherently limit the scope for upward revision among participants with high baseline trust, this pattern suggests that information provision is particularly effective among initially more skeptical individuals. We also examine heterogeneity along the other treatment dimensions. Figures A5 and A6 report analogous analyses for the VR vs. video comparison and for scientist vs. influencer speakers. Overall, differences across subgroups are limited. One exception is that the VR treatment appears somewhat more effective for students outside economics-related majors, although the differences are modest and statistically imprecise. More broadly, the similarity of treatment effects across students from economics and non-economics disciplines suggests that the main findings are not driven by a particular field of study within the university.[27]

As the heterogeneity analysis considers a relatively large number of subgroup comparisons, Tables A21–A23 report omnibus interaction tests together with Benjamini–Hochberg false-discovery-rate adjusted q-values (Benjamini and Hochberg 1995). The adjustment controls the expected proportion of false discoveries among the interaction effects declared statistically significant. The correction substantially attenuates the evidence for treatment effect heterogeneity. For the information treatment, heterogeneity by baseline trust in vaccines remains statistically significant after the adjustment ($q < 0.001$), heterogeneity by political ideology remains significant only at the 10 percent level ($q = 0.050$), while differences by gender and baseline trust in the healthcare system remain only marginally significant ($q = 0.076$). In contrast, none of the heterogeneity tests for the VR vs. video comparison or for the scientist vs. influencer comparison survives the multiple-testing correction. We therefore view the subgroup analysis as exploratory and interpret the patterns outside these dimensions with appropriate caution.

Finally, Table A24 examines whether treatment effects differ according to whether the speaker shares the participant's gender. For trust in vaccines, assignment to a same-gender speaker is associated with a modest reduction in treatment effects, particularly within the scientist condition. However, the estimated homophily effects do not differ significantly between scientists and influencers, and no comparable patterns emerge for trust in the healthcare system. Overall, the evidence provides little support for the view that gender matching plays a substantively important role in shaping the effectiveness of either type of messenger.

6.4 Willingness to Discuss

Beyond participants' own beliefs, we examine whether the intervention affects their willingness to discuss vaccines and related topics with friends and family. Tables A25 and A26 report the results. Panel A of Table A25 shows that exposure to the information treatment increases the likelihood of discussing vaccines, healthcare authorities, artificial intelligence, and politics. Panel B of Table A25 reports the differences

[27] The analysis by field of study excludes 325 participants for whom information on their university department could not be collected.

between the VR and video condition, showing that participants receiving the information through VR statistically significantly increase their willingness to discuss artificial intelligence and social media, but not vaccination or healthcare topics. This finding suggests that the immersive nature of the VR experience triggers broader curiosity about the technological medium itself, beyond the specific health content of the message. Panel C of Table A25 shows no meaningful differences between scientist and influencer speakers, although point estimates suggest a lower likelihood of discussing social media and artificial intelligence following exposure to scientists; these differences are not statistically significant. The breakdown by speaker identity in Table A26 does not reveal additional systematic patterns, as estimates are generally small and imprecise relative to the baseline category (female influencer).

7 Mechanisms: Visual Attention, Speaker Characteristics, and Familiarity

The results so far show that persuasion depends not only on exposure to information, but also on how the information is delivered and, to some extent, on who delivers it. The larger effect of VR relative to standard video is consistent with the idea that the immersive environment generates greater engagement with the information. We next examine other mechanisms that may contribute to differences in persuasion across speakers. First, we ask whether participants who direct greater visual attention toward the speaker exhibit larger changes in beliefs. Second, we examine whether differences in persuasiveness are related to participants' evaluations of speaker characteristics, including knowledge, reliability, passion, and charisma. Finally, we consider the role of pre-existing familiarity and reputation, particularly for social media influencers, for whom participants may enter the experiment with established perceptions of the speaker.

7.1 Visual Attention

We begin by examining whether differences in visual attention help explain variation in persuasion across speakers. Using the eye-tracking data collected during the VR sessions, we test whether participants who direct greater visual attention toward the speaker exhibit larger changes in beliefs. The sample is therefore restricted to participants assigned to the VR treatment.

Descriptive evidence suggests that visual attention varies across speaker profiles.[28] In absolute terms, the male influencer and female scientist receive the highest average gaze time. However, because message duration differs across speakers, absolute gaze time is not directly comparable across speakers. Relative to exposure time, participants direct the largest share of their gaze toward the male influencer, followed by the female influencer, the female scientist, and the male scientist. Thus, the longer duration of the scientists' messages does not translate into greater visual attention. Regression estimates in Columns 5 and 6 of Table 6 are consistent with these descriptive patterns but are imprecisely estimated. Differences between

[28] Table A5 provides the corresponding descriptive statistics. Figure A3 reports descriptive differences in exposure length and gaze allocation across speakers.

scientists and influencers are not statistically significant at the aggregate level, and estimates comparing individual speakers are similarly imprecise.

7.2 The Role of Speaker Characteristics

We next examine whether differences in persuasion across speakers depend on how participants perceive their characteristics. We begin with two broad measures: Participants' overall impression of the speaker and of the message. These variables capture general evaluations of how positively the communicator and the content are perceived. Columns 1 to 4 of Table 6 report the corresponding results. Panel A compares scientists and influencers. The differences are large and precisely estimated for impressions of the speaker: Scientists are evaluated substantially more positively, with effect sizes of approximately 0.39 standard deviations. In contrast, there is no meaningful difference in impressions of the message itself, which suggests that participants perceive the informational content similarly regardless of the messenger. Panel B provides a breakdown by speaker identity. The highest ratings are attributed to the male scientist, followed by the female scientist, both significantly outperforming the female influencer (the reference category). By contrast, the male influencer receives lower evaluations of the speaker. Differences in impressions of the message are more muted, although the male scientist generates a slightly more positive evaluation.

To summarize participants' evaluations across a broader set of speaker characteristics, we construct an index of perceived speaker quality using principal component analysis, as described in Section 3.2. This index aggregates ratings of whether the speaker is charismatic, persuasive, reliable, knowledgeable, sympathetic, attractive, and passionate. The results are reported in Table 6 (columns 7 and 8) and further decomposed in Table 7. At the aggregate level displayed in Panel A, scientists are rated substantially higher than influencers, with differences of roughly 0.40 standard deviations. The breakdown by speaker identity, however, reveals a more nuanced pattern. While the male scientist drives much of the positive differential, the male influencer is evaluated less favorably than both scientists and the female influencer across several dimensions. Panel B further shows that scientists outperform influencers primarily on attributes associated with reliability and expertise, including reliability, knowledgeability, and persuasiveness, whereas the female influencer scores more highly on attributes such as attractiveness and passion. The male influencer, by contrast, performs relatively poorly across most perceived speaker characteristics. These differences reveal distinct bundles of perceived attributes across speakers, which we next examine in relation to their persuasiveness.

To investigate this possibility, we conduct a mediation analysis reported in Tables A27 and A28. For each characteristic, we decompose the difference in persuasive effectiveness between each speaker and the female influencer into direct and indirect effects. The indirect effect combines the effect of assignment to a given speaker, relative to the female influencer, on the corresponding speaker evaluation with the association between that evaluation and the post-treatment outcome, conditional on speaker assignment. Because the indirect effect is constructed as the product of these two estimated coefficients, we obtain its standard error and statistical significance from 1,000 bootstrap replications, resampling at the session level. The direct effect captures the remaining difference in the outcome between the two speakers after conditioning on the corresponding evaluation, with standard errors clustered at the session level. We consider each characteristic separately.[29] Because speaker assignment is randomized but post-treatment speaker evaluations are not, we interpret these decompositions as suggestive evidence on the mechanisms underlying the treatment effects rather than as identifying causal effects of the mediators themselves.

The results for Trust in Vaccines in Table A27 reveal markedly different patterns across speakers. For both scientists, differences in persuasive effectiveness relative to the female influencer are strongly associated with differences in reliability and knowledgeability. The indirect effects through reliability are approximately 0.11 points on the Trust in Vaccines index for both scientists, while those through knowledgeability are approximately 0.14 points. Once either dimension is accounted for, the remaining direct effects are close to zero. Attractiveness and passion, by contrast, generate negative indirect effects for both scientists, consistent with the female influencer's relative advantage on these dimensions partially offsetting the scientists' advantage in perceived reliability and expertise. The results suggest that the scientists' persuasive effectiveness is closely associated with attributes related to reliability and expertise, while attributes related to personal appeal work in favor of the female influencer.

The male influencer presents a notably different pattern. No measured evaluation accounts for his persuasive advantage relative to the female influencer. Instead, several indirect effects are negative, particularly those associated with charisma, attractiveness, and passion. Conditioning on these dimensions therefore increases, rather than reduces, the remaining direct effect. For example, the total difference of approximately 0.15 points on the Trust in Vaccines index increases to approximately 0.25 points after conditioning on passion. The male influencer's greater persuasive effectiveness thus occurs despite, rather than because of, his evaluations along many of the dimensions captured by our survey. This result motivates examining whether his effectiveness is instead associated with characteristics that predate the experimental interaction, particularly prior familiarity and reputation.

The corresponding decomposition for Trust in the Healthcare System, reported in Table A28, is broadly consistent with this pattern but less informative because there is little speaker-level heterogeneity in the total treatment effects for this outcome. Reliability and knowledgeability again generate positive indirect effects for both scientists, while attractiveness and passion generally operate in the opposite direction. However, the total differences between the scientists and the female influencer are small and statistically indistinguishable from zero. We therefore interpret these results as broadly consistent with the patterns observed for Trust in Vaccines, rather than as separate evidence of distinct mediation channels.

Overall, the evidence points to substantial differences in the characteristics associated with persuasion across speakers. For both scientists, persuasive effectiveness is closely associated with perceived reliability and knowledgeability, while attributes related to personal appeal work relatively more in favor of the female influencer. These same evaluations, however, do little to explain the greater persuasive effect of the male influencer. Indeed, his effect is larger despite relatively unfavorable evaluations across many of the characteristics measured after treatment. This pattern motivates us to turn to characteristics that predate the experimental interaction, and in particular to prior familiarity and reputation.

[29] We exclude perceived persuasiveness from the mediation analysis because it is conceptually too close to the outcome of interest: Unlike the other speaker evaluations, it directly captures participants' perceptions of the speaker's ability to persuade and therefore provides limited information about the underlying attributes through which persuasion may occur.

7.3 Familiarity and Reputation

We examine whether pre-existing familiarity and reputation help explain differences in persuasive effectiveness across influencers using participants' pre-treatment measures of familiarity with and overall opinion of the two influencers. If persuasive effectiveness depends partly on the reputation individuals attach to familiar public figures—in line with the persuasion framework of DellaVigna and Gentzkow (2010)—these pre-existing perceptions may help explain why otherwise similar messages generate different responses across speakers. The data provide initial support for this interpretation: We observe a strong positive correlation (0.45) between familiarity and overall opinion for the male influencer, while the corresponding correlation for the female influencer is close to zero.[30] Figure 3 further explores this relationship by plotting regression-adjusted levels of trust in vaccines among participants assigned to each influencer across categories of familiarity (Panel A) and overall opinion or reputation (Panel B). For the male influencer, predicted trust is substantially higher among participants who report at least some familiarity with him or a positive opinion of him than among those who report none, with little additional variation across higher categories. This is consistent with a minimum level of recognition or favorable reputation being associated with greater effectiveness of the message. We find no comparable relationship for the female influencer, for whom predicted trust remains broadly similar across categories of familiarity and reputation.

In Table A29, we further examine whether prior familiarity is associated with the difference in persuasive effectiveness between the two influencers. Restricting the analysis to the influencer treatments and to participants for whom familiarity with and prior opinion of the assigned influencer are observed, the male influencer increases Trust in Vaccines by approximately 0.15 points on the index relative to the female influencer in column (1). This estimate mirrors the differential effect of the male influencer reported in column (2) of Table 4, with the small difference reflecting the sample restriction used in this exercise. Column (2) shows that conditioning on prior familiarity reduces this difference to approximately 0.10 points, an attenuation of about 31 percent. By comparison, column (3) shows that controlling for participants' prior opinion of the assigned influencer produces a smaller attenuation, of approximately 14 percent. Column (4), which includes both controls, reduces the difference to approximately 0.10 points, close to the estimate obtained with familiarity alone. Because both familiarity and prior opinion capture perceptions formed before the experimental interaction, this exercise differs from the mediation analysis above: Rather than representing indirect effects of speaker assignment, the attenuation indicates that pre-existing differences in familiarity and reputation are associated with part of the difference in persuasive effectiveness between the two influencers.

These results suggest that pre-existing familiarity is associated with part of the male influencer's relative effectiveness, with the evidence on prior reputation pointing in the same direction. Post-treatment trust is higher among participants who are at least somewhat familiar with him or hold a favorable prior opinion, although additional increases in familiarity or reputation do not translate into progressively larger

[30] Familiarity is measured using the survey question "Below you will see a list of Greek content creators. How familiar are you with them?", with responses collected on a Likert scale ranging from "I do not know them at all" to "Extremely familiar." Overall opinion is based on the question "Thinking about the content produced by each of the following creators, what is your overall opinion of them?", with responses ranging from "Very negative" to "Very positive." These measures are discussed in Section 4.

effects. At the same time, the female influencer generates positive treatment effects despite being substantially less familiar to participants. The mechanism evidence suggests that the credibility underlying persuasion can derive from different speaker characteristics. For scientists, persuasion is closely associated with perceived reliability and knowledgeability; for the female influencer, attributes related to personal appeal appear relatively more important; and for the male influencer, pre-existing familiarity and reputation play a role despite relatively unfavorable evaluations. Thus, effective credibility need not originate in expertise alone: It may also derive from personal appeal or from a reputation established before the persuasive interaction. Given that VR likely increases persuasion by generating greater engagement with the information, these patterns provide the empirical motivation for the model developed in Section 8.

8 A Model of Persuasion via Engagement, Credibility, and Familiarity

The experiment documented in the paper establishes five empirical regularities that we seek to explain within a single theoretical framework.

R1. *Information matters.* Exposure to the public-health message significantly increases trust in vaccines (by roughly 0.28 standard deviations) and, to a lesser extent, trust in the healthcare system.

R2. *The medium matters.* The same message delivered through immersive virtual reality (VR) generates approximately 0.11 of a standard deviation more belief change than standard video delivery.

R3. *Messenger identity matters, but expertise alone does not explain it.* Medical experts are perceived as substantially more reliable and knowledgeable, yet social media influencers achieve comparable average persuasion effects.

R4. *Visual attention varies across speakers but provides limited explanatory power.* Eye-tracking data collected within the VR arm show influencers capture a larger share of gaze time, although the messenger-specific gaze differential is imprecisely estimated (Table 6) and is not itself tested as a mediator of belief revision.

R5. *Effects persist.* A substantial portion of the initial belief revision remains detectable at the two-month and six-month follow-ups, consistent with the intervention producing durable belief change rather than a transient response.

The model is built around three primitives: an *engagement-enhancing technology* that converts exposure into absorbed information; a *credibility-familiarity trade-off* that determines how an agent weights the absorbed signal; and a *memory-decay structure* that governs persistence.

8.1 Environment

8.1.1 Agents and the State of the World

There is a continuum of agents $i \in [0, 1]$. The state of the world $\theta \in \mathbb{R}$ represents the "true quality" of vaccines—a latent variable that aggregates scientific evidence on vaccine safety, efficacy, and importance for public health. A higher value of θ means that vaccines are safer, more effective, or more socially valuable. Agents do not observe θ directly; they hold beliefs about it.

Prior Beliefs. At the beginning of the experiment (time $t = 0$, pre-treatment), agent i holds a prior

$$\theta \sim \mathcal{N}(\mu_0, \sigma_0^2), \tag{5}$$

where $\mu_0 \in \mathbb{R}$ is the prior mean (the agent's baseline trust in vaccines) and $\sigma_0^2 > 0$ is the prior variance (reflecting uncertainty). Heterogeneity across agents in μ_0 or σ_0^2 will generate the heterogeneous treatment effects documented in the empirical analysis above; we suppress the i subscript when working with a representative agent.

8.1.2 The Public-Health Message

A messenger $m \in \{S, I\}$ (S for *scientist*, I for *influencer*) delivers a fixed informational content: the scientific consensus on vaccine quality. We assume the consensus coincides with the true state, $\theta^* = \theta$; the message is then a noisy signal of the state,

$$x_m = \theta + \varepsilon_m, \qquad \varepsilon_m \sim \mathcal{N}(0, \sigma_m^2), \tag{6}$$

so that the likelihood used in updating is $x_m \mid \theta \sim N(\theta, \sigma_m^2)$. Crucially, θ is identical across all treatment arms—it is the same script read by scientists and influencers alike. The messenger identity affects σ_m^2 (how precisely the signal is decoded), not its content.[31]

Assumption 1 (Credibility Ordering) *Scientists are endowed with higher intrinsic credibility than influencers, so that the agent assigns a lower interpretation noise to the scientist's signal:*

$$\sigma_S^2 < \sigma_I^2.$$

This assumption captures the empirical finding that scientists are rated significantly higher on reliability, knowledgeability, and persuasiveness (Table 7).

8.1.3 Engagement as a Filter

Not all of the signal x_m reaches the agent. Before the signal is processed, the agent allocates engagement intensity $a \in [0, 1]$, which acts as an absorption coefficient. The *absorbed signal* is

$$\tilde{x}_m(a) = a x_m + (1 - a)\mu_0 + \eta, \qquad \eta \sim \mathcal{N}(0, \sigma_\eta^2), \tag{7}$$

[31] Nothing of substance changes if the consensus is instead a noisy but unbiased proxy for the state, so that the messenger's signal is $x_m = \theta + u + \varepsilon_m$ with $u \sim N(0, \sigma_u^2)$ independent of ε_m and all other shocks. Since $u + \varepsilon_m \sim N(0, \sigma_m^2 + \sigma_u^2)$, the conditional distribution retains the form $x_m \mid \theta \sim N(\theta, \sigma_m^2 + \sigma_u^2)$: agents continue to update beliefs about the true state θ, and every expression in Sections 8.2–8.4 is unchanged after replacing σ_m^2 by $\sigma_m^2 + \sigma_u^2$.

where η is residual noise due to inattention and μ_0 is the agent's prior mean. In expectation, when $a = 0$ the agent absorbs none of the informational content and the absorbed signal reduces to the prior plus non-informative noise, $E[\tilde{x}_m(0)] = \mu_0$; realized posteriors can still differ from the prior at $a = 0$ due to the realization of η, which is uninformative about θ by construction. When $a = 1$, the agent absorbs the full signal plus idiosyncratic noise η.

In equation (7), we assume agents process the absorbed signal naively: they treat the absorbed signal $\tilde{x}_m(a)$ as an unbiased signal of the state with variance $\tilde{\sigma}_m^2(a) = a^2\sigma_m^2 + \sigma_\eta^2$, and do not correct for the attenuation of the unabsorbed content. This is precisely the behavioral-inattention formulation of Gabaix (2014) and Gabaix (2019), in which the perceived signal is anchored on a default, with the default here equal to the prior mean μ_0. Inattention is precisely the failure to register content one did not absorb, so the agent cannot undo an attenuation whose magnitude she does not perceive.[32] The behavioral justification is immediate: inattention means the agent is unaware of the portion of the message she did not absorb; correcting for attenuation would require knowing exactly how much content was missed, which contradicts the premise of inattention.

The economically relevant object throughout is the *engagement multiplier*

$$F(a) \;\equiv\; \lambda^{(m)}(a)\cdot a \;=\; \frac{\sigma_0^2\, a}{\sigma_0^2 + a^2\sigma_m^2 + \sigma_\eta^2}, \qquad F_I = F^{(I)}(a_I),\; F_S = F^{(S)}(a_S), \tag{8}$$

where the Bayesian weight on the new signal is

$$\lambda^{(m)}(a) \;=\; \frac{\sigma_0^2}{\sigma_0^2 + a^2\sigma_m^2 + \sigma_\eta^2} \;\in\; (0,1). \tag{9}$$

Assumption 2 (Moderate Signal Noise) *For each messenger $m \in \{S, I\}$, the interpretation noise satisfies $\sigma_m^2 < \sigma_0^2 + \sigma_\eta^2$.*

Under Assumption 2, $F(a)$ is increasing on $[0,1]$: greater engagement raises the absorbed informational content faster than it raises absorbed noise, so the engagement multiplier is strictly increasing.

Communication Medium and Exogenous Engagement. The key insight is that the communication medium $\ell \in \{\text{Video}, \text{VR}\}$ *shifts engagement intensity exogenously*, independently of the agent's choice. Immersive VR creates a richer sensory environment that can generate greater engagement with the information, even for agents who would otherwise be partially distracted.

[32] If instead agents fully understand the absorption technology, they can debias the signal:

$$\hat{z}_m \;=\; \frac{\tilde{x}_m(a) - (1-a)\mu_0}{a} \;=\; x_m + \frac{\eta}{a}, \qquad \text{Var}(\hat{z}_m \mid \theta) = \sigma_m^2 + \frac{\sigma_\eta^2}{a^2},$$

which is an unbiased signal whose precision is increasing in a. Standard updating then yields $\mathbb{E}[\Delta\mu] = \gamma_m^i\, \hat{\lambda}(a)\,(\theta - \mu_0)$ with $\hat{\lambda}(a) = \sigma_0^2/(\sigma_0^2 + \sigma_m^2 + \sigma_\eta^2/a^2)$ strictly increasing in a. All qualitative results go through in this variant as well, with $F(a)$ replaced by $\hat{\lambda}(a)$ and no additional assumption needed.

Assumption 3 (Medium Ordering) *VR generates strictly higher engagement intensity than video, all else equal:*

$$a_{VR} > a_{Video},$$

where a_ℓ denotes the medium-specific component of engagement.

We write total engagement intensity as

$$a = a_\ell + \alpha_m, \tag{10}$$

where $a_\ell \in (0,1)$, $\ell \in \{\text{Video}, \text{VR}\}$, is the medium-induced component and $\alpha_m \in [0, 1 - a_\ell]$, $m \in \{S, I\}$, is a messenger-specific component capturing the agent's intrinsic interest in the messenger.

Messenger-Specific Engagement. Messenger-specific characteristics may also affect engagement with the information, which we capture through α_m. We allow influencers to generate greater messenger-specific engagement than scientists:

$$\alpha_I > \alpha_S. \tag{11}$$

For influencers, this engagement may reflect characteristics such as familiarity, personal appeal, and parasocial connections. Scientists, by contrast, benefit from greater perceived reliability and knowledgeability, which the model captures through lower signal noise σ_m^2.

The visual-attention measures collected in the VR arm provide complementary evidence on differences in gaze allocation across speakers, but do not directly measure the broader engagement parameter α_m.

8.2 Belief Updating

Belief updating proceeds in two stages. In the first stage the agent processes the absorbed signal through a credibility-weighted filter. In the second stage, a familiarity adjustment modifies the effective weight placed on the messenger's message.

Stage 1: Bayesian Update on the Absorbed Signal. Given the absorbed signal $\tilde{x}_m(a)$, the variance $\tilde{\sigma}_m^2(a) = a^2\sigma_m^2 + \sigma_\eta^2$, and the prior $\mathcal{N}(\mu_0, \sigma_0^2)$, standard Bayesian updating yields a posterior mean

$$\mu_1^{(m)}(a) \;=\; \frac{\tilde{\sigma}_m^2(a)\,\mu_0 \;+\; \sigma_0^2\,\tilde{x}_m(a)}{\sigma_0^2 + \tilde{\sigma}_m^2(a)}, \tag{12}$$

with posterior precision

$$\tau_1^{(m)}(a) \;=\; \frac{1}{\sigma_0^2} + \frac{1}{\tilde{\sigma}_m^2(a)}. \tag{13}$$

Since the Bayesian weight on the new signal $\lambda^{(m)}(a)$ is defined in (9), we can write

$$\mu_1^{(m)}(a) \;=\; (1 - \lambda^{(m)}(a))\,\mu_0 \;+\; \lambda^{(m)}(a)\,\tilde{x}_m(a). \tag{14}$$

Equation (14) makes the mechanism transparent: the posterior is a convex combination of the prior mean and the absorbed signal, with weight $\lambda^{(m)}(a)$ on the signal, which is decreasing in a.

Stage 2: Familiarity-Adjusted Credibility Weighting. Agents do not take the signal at face value; they discount it by a credibility factor that depends jointly on the perceived expertise of the messenger and on their prior familiarity with that messenger. Familiarity generates *parasocial trust* (Horton and Wohl 1956), which partially offsets the credibility disadvantage of influencers.

Define the *effective credibility* of messenger m for agent i as

$$\gamma_m^i = \underbrace{\bar{\gamma}_m}_{\text{intrinsic credibility}} + \underbrace{\delta r_m^i f_m^i}_{\text{familiarity bonus}}, \tag{15}$$

where $\bar{\gamma}_m \in (0,1)$ is the population-average intrinsic credibility of messenger type m, $f_m^i \in [0,1]$ is agent i's prior familiarity with messenger m (measured on a normalized scale), $r_m^i \in [0,1]$ is agent i's (normalized, non-negative) prior reputation of, or opinion about, messenger m, and $\delta > 0$ is the marginal return to familiarity in trust formation. Empirically, f_m^i corresponds to the familiarity question (Question 21) rescaled linearly from its 1–5 range to $[0,1]$, and r_m^i to the overall-opinion question (Question 22) rescaled linearly from its -2 to 2 range to $[0,1]$, so that a neutral opinion maps to $r_m^i = 1/2$.

Assumption 4 (Bounded Credibility) $\delta \leq \min\{1-\bar{\gamma}_S,\, 1-\bar{\gamma}_I\}$.

Since $r_m^i f_m^i \in [0,1]$, Assumption 4 implies $\gamma_m^i = \bar{\gamma}_m + \delta r_m^i f_m^i \leq \bar{\gamma}_m + \delta \leq 1$ for every $m \in \{S, I\}$ and every agent i, so that, together with $\lambda^{(m)}(a) \in (0,1)$, we have $\gamma_m^i \lambda^{(m)}(a) \in (0,1)$: the familiarity-adjusted posterior in equation (16) remains a strict convex combination of the prior mean and the absorbed signal for every messenger and every agent. Unlike a bound on δ relative to $\bar{\gamma}_S - \bar{\gamma}_I$, Assumption 4 does not restrict the sign of $\gamma_I^i - \gamma_S^i$: for an influencer with sufficiently high familiarity f_I^i and reputation r_I^i, effective credibility can exceed that of the scientist, i.e., $\gamma_I^i > \gamma_S^i$. Familiarity can therefore not only close but, for high-familiarity, high-reputation agents, overturn the credibility gap: parasocial trust does not merely substitute for expertise one-for-one, but can dominate it. In particular, familiarity dominating for high-f_I^i (influencer) agents lines up with Table 4 and Section 7.3.

The familiarity-adjusted posterior mean is

$$\hat{\mu}_1^{(m,i)}(a) \;=\; (1 - \gamma_m^i\, \lambda^{(m)}(a))\, \mu_0 \;+\; \gamma_m^i\, \lambda^{(m)}(a)\, \tilde{x}_m(a). \tag{16}$$

The *belief revision*—the change in beliefs produced by the intervention—is

$$\Delta\mu^{(m,i)}(a) \;\equiv\; \hat{\mu}_1^{(m,i)}(a) - \mu_0 \;=\; \gamma_m^i\, \lambda^{(m)}(a)\, [\tilde{x}_m(a) - \mu_0]. \tag{17}$$

Taking expectations over ε_m and η, and noting that $\mathbb{E}[\tilde{x}_m(a)] = a(\theta - \mu_0) + \mu_0$, we obtain the *expected belief revision*:

$$\mathbb{E}\big[\Delta\mu^{(m,i)}(a)\big] \;=\; \gamma_m^i \lambda^{(m)}(a) a(\theta - \mu_0). \tag{18}$$

Equation (18) is the central equation of the model. It decomposes the expected belief revision into three multiplicative terms: (i) effective credibility γ_m^i, which depends on messenger identity and familiarity; (ii) the Bayesian weight $\lambda^{(m)}(a)$, which is increasing in prior uncertainty and decreasing in engagement (more absorption imports more message noise)—so that engagement raises persuasion only through the multiplier $F(a) = a\,\lambda^{(m)}(a)$, which is increasing in a under Assumption 2; and (iii) engagement intensity a, which is driven by both the medium and the messenger's ability to attract engagement.

8.3 Main Results

We now use equation (18) to derive the five empirical regularities as formal propositions.

8.3.1 R1: Information Matters

We have a first result:[33]

Proposition 1 (Information Shifts Beliefs) *Suppose $\theta > \mu_0$ (the true vaccine quality exceeds the agent's prior, as is empirically plausible given documented vaccine hesitancy). Then, for any $m \in \{S, I\}$, any medium ℓ, and any $a > 0$:*

$$\mathbb{E}[\Delta\mu^{(m,i)}(a)] > 0.$$

Moreover, the belief revision is increasing in prior uncertainty σ_0^2 and vanishes as $\sigma_0^2 \to 0$.

The heterogeneity result from the experiment—that treatment effects are largest among initially skeptical individuals—follows immediately: agents with lower μ_0 face a larger gap $\theta - \mu_0$ and therefore experience a larger expected belief revision. Also, more uncertain agents (higher σ_0^2) assign greater weight to the new signal.

The smaller effect on trust in the healthcare system (relative to trust in vaccines) can be rationalized by allowing for a two-dimensional state space $\boldsymbol{\theta} = (\theta_V, \theta_H)$, where θ_V is vaccine quality and θ_H is healthcare system quality. The signal x_m is primarily informative about θ_V, with a weaker cross-informational content for θ_H. Formally, the message is a direct signal of θ_V but only an indirect (and therefore noisier) signal of θ_H, yielding a smaller Bayesian weight on the healthcare dimension.

8.3.2 R2: The Medium Matters

Proposition 2 (VR Amplifies Belief Revision) *For a given messenger m and messenger-specific engagement α_m, VR delivery generates a strictly larger expected belief revision than video delivery:*

$$\mathbb{E}[\Delta\mu^{(m,i)}(a_{VR} + \alpha_m)] \;>\; \mathbb{E}[\Delta\mu^{(m,i)}(a_{Video} + \alpha_m)].$$

The VR premium operates through *engagement*, not through the informational content of the message. Since $F(a)$ is increasing in a, any technology that mechanically raises engagement intensity—even without changing the signal—will amplify belief revision. This is consistent with the experimental finding that VR and video deliver identical scripts but produce different persuasion effects.

[33] All proofs of propositions can be found in Online Appendix C.

8.3.3 R3: Scientists and Influencers Achieve Comparable Average Persuasion

The experiment finds that, on average, scientists and influencers produce statistically indistinguishable belief changes, despite scientists being rated as substantially more credible. Proposition 3 shows that the credibility disadvantage of influencers can be exactly offset by their familiarity advantage.

Proposition 3 (Credibility–Familiarity Parity) *Fix a medium $\ell \in \{Video, VR\}$. The expected belief revisions of a scientist and an influencer are equal for agent i if and only if*

$$\gamma_I^i\, F(a_I) = \gamma_S^i\, F(a_S), \tag{19}$$

i.e., if and only if

$$\delta\big(r_I^i f_I^i - r_S^i f_S^i\big) = \bar{\gamma}_S - \bar{\gamma}_I + \gamma_S^i\, \frac{F_S - F_I}{F_I} \tag{20}$$

holds.

The parity condition (19) has three distinct forces. First, the influencer's effective credibility γ_I^i is boosted by high familiarity f_I^i, which is empirically plausible for the well-known male influencer (familiarity score 3.24 on a five-point scale vs. 1.84 for the female influencer). Second, the influencer's higher messenger-specific engagement α_I raises the engagement product $\lambda^{(I)}(a_I)\cdot a_I$. Third, the scientist's lower noise $\sigma_S^2 < \sigma_I^2$ raises $\lambda^{(S)}$ for any given a. Parity obtains when the first two forces together offset the third. At the population level, these forces happen to approximately balance, yielding the near-zero average scientist–influencer treatment effect differential reported in Panel C of Table 3 in the paper. Individual-level heterogeneity around this mean—documented in Figure 3, which plots regression-adjusted post-treatment trust levels against prior familiarity and reputation—then reflects variation across agents in f_m^i. Table 4 documents a complementary form of heterogeneity: fixed differences across the four speaker profiles attributable to delivery and presentation rather than to any single agent's familiarity.

Corollary 1 (Familiarity Heterogeneity) *Suppose $\theta > \mu_0$. For influencer I, the expected belief revision satisfies*

$$\frac{\partial\, \mathbb{E}\big[\Delta\mu^{(I,i)}(a_I)\big]}{\partial f_I^i} = \delta\, r_I^i\, F(a_I)\,(\theta - \mu_0) \;\geq\; 0,$$

which is strictly positive if and only if $r_I^i > 0$. The model therefore predicts a familiarity gradient that is steep for a messenger with a positive prior reputation and flat for a messenger with a neutral or negative prior reputation.

Corollary 1 gives the partial effect of familiarity holding reputation r_I^i fixed. Because familiarity and reputation are themselves correlated in the data (0.45 for the male influencer; ≈ 0 for the female influencer), the total gradient traced out in Figure 3 also reflects how reputation co-moves with familiarity: $dE[\Delta\mu]/df_I^i = \delta F(a_I)(\theta-\mu_0)\,\big[r_I^i + f_I^i\,\partial r_I^i/\partial f_I^i\big]$. Corollary 1 is therefore *qualitatively consistent with* rather than a direct, parameter-free prediction of the steep-vs.-flat contrast in Figure 3: it correctly signs the effect ($r_I^i > 0$ implying a positive gradient) but the observed steepness reflects the joint movement of familiarity and reputation, not familiarity alone. We also note the empirical pattern in Figure 3 more closely resembles a level shift from no familiarity to some familiarity than a sustained linear slope; a model in which $r_I^i(f_I^i)$ itself rises steeply from zero and then flattens would reproduce this shape more precisely.

8.3.4 R4: Engagement and Persuasion

Write $\Delta_m \equiv \mathbb{E}[\Delta\mu^{(m,i)}(a_m)]$, $F_m \equiv F(a_m)$, and (using the reputation-weighted familiarity) $\gamma_m^i = \bar{\gamma}_m + \delta r_m^i f_m^i$. The exact decomposition of the scientist–influencer gap is

$$\Delta_I - \Delta_S = (\theta - \mu_0)\Big[\underbrace{(\gamma_I^i - \gamma_S^i)F_I}_{\text{credibility–familiarity channel}} + \underbrace{\gamma_S^i(F_I - F_S)}_{\text{engagement channel}} \Big]. \tag{21}$$

We have the following result:

Proposition 4 (Engagement Does Not Fully Explain Persuasion) *Suppose influencers generate greater messenger-specific engagement than scientists ($\alpha_I > \alpha_S$). Then:*

(a) *The influencer's engagement advantage raises the engagement multiplier $F(a_I) = \lambda^{(I)}(a_I)\cdot a_I$ relative to the counterfactual with $\alpha_I = \alpha_S$;*

(b) *Set $r_I^i f_I^i = r_S^i f_S^i = 0$. Persuasion parity requires $F^{(I)}(a_I) = (\bar{\gamma}_S/\bar{\gamma}_I)\, F^{(S)}(a_S)$. Since $F^{(I)}$ is strictly increasing on $[0,1]$ with maximum $F^{(I)}(1) = \sigma_0^2/(\sigma_0^2+\sigma_\eta^2+\sigma_I^2)$, define the engagement-only threshold $\bar{\alpha} \equiv \big(F^{(I)}\big)^{-1}\Big(\frac{\bar{\gamma}_S}{\bar{\gamma}_I}\, F^{(S)}(a_S)\Big) - a_\ell - \alpha_S$, whenever $\frac{\bar{\gamma}_S}{\bar{\gamma}_I}\, F^{(S)}(a_S) \le F^{(I)}(1)$. If $\frac{\bar{\gamma}_S}{\bar{\gamma}_I} F^{(S)}(a_S) > F^{(I)}(1)$, no feasible engagement level closes the gap with $f_I^i = 0$, and either (i) $\delta r_I^i f_I^i > 0$ or (ii) a lower σ_I^2 is required. Otherwise, engagement alone suffices whenever $\alpha_I - \alpha_S \ge \bar{\alpha}$.*

(c) *Decompose the influencer–scientist revision gap as in (21). Assume $F_I \ge F_S$. Then, the engagement channel is uniformly bounded:*

$$\gamma_S^i\,(F_I - F_S) \;\le\; \gamma_S^i\,\frac{\sigma_0^2}{\sigma_0^2 + \sigma_\eta^2}\,(\alpha_I - \alpha_S).$$

Hence the credibility–familiarity channel exceeds the engagement channel in absolute value whenever

$$\Big|\,\delta\big(r_I^i f_I^i - r_S^i f_S^i\big) - \big(\bar{\gamma}_S - \bar{\gamma}_I\big)\Big|\,F_I \;>\; \gamma_S^i\,\frac{\sigma_0^2}{\sigma_0^2 + \sigma_\eta^2}\,(\alpha_I - \alpha_S). \tag{22}$$

The eye-tracking evidence in Table 6 provides a complementary measure of visual attention within the VR arm. Differences in gaze allocation across speakers are imprecisely estimated and provide limited evidence that visual attention explains differences in persuasion. This is consistent with the model, in which the broader engagement channel enters belief revision multiplicatively with effective credibility and prior uncertainty. Differences in engagement can therefore coexist with differences in effective credibility arising from familiarity and reputation.

In part (c), the left-hand side of (22) is strictly increasing in f_I^i (for $r_I^i > 0$) beyond the point $\delta r_I^i f_I^i = \bar{\gamma}_S - \bar{\gamma}_I + \delta r_S^i f_S^i$, while the right-hand side does not depend on f_I^i. Therefore, provided δr_I^i is large enough that (22) can hold at $f_I^i = 1$, there exists a threshold $\bar{f} \in (0,1)$ such that the credibility–familiarity channel dominates for all $f_I^i > \bar{f}$. Engagement remains a necessary conduit ($F(0) = 0$), but among audiences with high familiarity with and favorable prior views of the influencer, it is not the primary driver of the influencer–scientist gap. Thus, the model can rationalize the greater persuasive effectiveness observed among participants with high familiarity with and favorable prior views of the male influencer without requiring differences in visual attention to account for the speaker gap. We emphasize that $\bar{f}$ is derived from the model's structure and is not calibrated against the paper's point estimates of $\bar{\gamma}_m$, δ, σ_0^2, or σ_η^2, none of which is separately identified by the experimental design. The argument therefore establishes that the empirical pattern is *rationalized* by the model for an economically plausible range of parameters, not that the model *estimates* which channel dominates in this sample.

8.3.5 R5: Persistence of Belief Revision

We now extend the model to a dynamic setting to address the persistence finding. Let time be indexed by $t \in \{0,1,2\}$, where $t = 0$ is the pre-treatment baseline, $t = 1$ is immediately post-treatment, and $t = 2$ is the two-month follow-up.

Memory Decay. After the treatment session, agents do not actively re-process the signal. Instead, the post-treatment belief $\hat{\mu}_1^{(m,i)}$ depreciates toward the pre-treatment prior at rate $\rho \in (0,1)$, subject to a *consolidation effect*: a fraction $\kappa_\ell \in (0,1)$ of the belief revision becomes consolidated into long-term memory. We allow κ to differ across arms, reflecting potential differences in the durability with which information is encoded, and thus to depend on the medium $\ell \in \{\text{Video}, \text{VR}\}$.

Formally, define the *consolidated belief shift* as

$$\hat{\mu}_1^{c,(m,i)} \;=\; \mu_0 + \kappa_\ell \, \Delta\mu^{(m,i)}(a), \tag{23}$$

where $\kappa_\ell \in (0,1)$ reflects the fraction of the belief revision that is durably encoded. The remaining fraction $(1-\kappa_\ell)$ is subject to decay. The belief at time $t = 2$ is

$$\hat{\mu}_2^{(m,i)} \;=\; \hat{\mu}_1^{c,(m,i)} + (1-\rho)\,[\hat{\mu}_1^{(m,i)} - \hat{\mu}_1^{c,(m,i)}], \tag{24}$$

which simplifies to

$$\hat{\mu}_2^{(m,i)} \;=\; \mu_0 + [\kappa_\ell + (1-\rho)(1-\kappa_\ell)]\,\Delta\mu^{(m,i)}(a). \tag{25}$$

Proposition 5 (Persistence)

(a) The belief revision at $t = 2$ satisfies

$$\mathbb{E}[\hat{\mu}_2^{(m,i)} - \mu_0] \;=\; \underbrace{[\kappa_\ell + (1-\rho)(1-\kappa_\ell)]}_{\equiv \Pi_\ell \,\in\, (0,1)}\,\mathbb{E}[\Delta\mu^{(m,i)}(a)].$$

The persistence multiplier $\Pi_\ell \in (0,1)$ satisfies $\Pi_\ell > \kappa_\ell$ and $\Pi_\ell < 1$, so belief revision partially but not fully decays.

(b) *The ratio of follow-up to immediate effects is Π_ℓ. Because Π_ℓ depends on the medium but not on the messenger, the ranking of messenger effects within a medium is preserved at follow-up, whereas the ranking of medium effects (and of messenger–medium combinations) can change.*

(c) *Trust in vaccines persists more than trust in the healthcare system if the consolidation parameter κ_ℓ is higher for vaccine attitudes, i.e., $\kappa_{\ell,V} > \kappa_{\ell,H}$. This can be rationalized by the greater personal relevance of vaccine trust for the target population.*

The ratio of follow-up to immediate treatment effects identifies the composite persistence multiplier $\Pi_\ell = \kappa_\ell + (1-\rho)(1-\kappa_\ell)$, but not the consolidation and decay parameters separately: any pair (κ_ℓ, ρ) on the locus $\kappa_\ell + (1-\rho)(1-\kappa_\ell) = \Pi_\ell$ is observationally equivalent given two measurement waves. Comparing our preferred specifications, the estimated VR advantage for trust in vaccines changes from 0.105 immediately after treatment (Table 3) to -0.083 at the two-month follow-up (Table A20): the VR advantage does not just shrink after two months, it disappears or flips sign, which is consistent with Proposition 5(b).

8.4 Comparative Statics

We now collect the main comparative statics implied by the central equation (18).[34]

Proposition 6 (Comparative Statics) *The expected belief revision $\mathbb{E}[\Delta\mu^{(m,i)}(a)]$ is:*

(i) ***Increasing in prior uncertainty σ_0^2:*** *Agents who are more uncertain about vaccine quality revise their beliefs more upon exposure to information.*

(ii) ***Increasing in engagement intensity** a **(under Assumption 2):*** *This effect operates through the engagement multiplier $F(a)$, which is in turn increasing in the richness of the communication medium ($a_{VR} > a_{Video}$) and in messenger-specific engagement ($\alpha_I > \alpha_S$).*

(iii) ***Increasing in effective credibility γ_m^i:*** *It is increasing in both intrinsic credibility $\bar{\gamma}_m$ and in familiarity f_m^i.*

(iv) ***Decreasing in baseline trust μ_0:*** *Agents who already hold high prior beliefs (close to θ) have little room for upward revision.*

(v) ***Strictly decreasing in signal noise σ_m^2:*** *Lower noise (higher credibility) increases $\lambda^{(m)}(a)$ directly, but may not be the binding constraint when familiarity f_m^i is high.*

These comparative statics are directly consistent with the heterogeneity patterns documented in the paper: women show larger effects than men (interpretable as higher σ_0^2 or higher familiarity-driven responsiveness); politically disengaged individuals show larger effects (consistent with higher prior uncertainty); and agents with lower baseline vaccine trust show the largest revisions (comparative static (iv)).

[34]The results stated in Proposition 6 are straightforward, and we therefore omit their proofs.

8.5 Relation to the Persuasion Literature

Our framework is closely related to the Bayesian persuasion approach of Kamenica and Gentzkow (2011) and the empirical persuasion framework of DellaVigna and Gentzkow (2010). Like DellaVigna and Gentzkow (2010), we emphasize that persuasion is determined jointly by exposure, engagement, credibility, and audience characteristics. Our contribution is to model these four forces within a unified multiplicative structure (equation (18)) that allows us to understand each channel from the experimental variation.

Relative to the sender-receiver game of Kamenica and Gentzkow (2011), our model is one of exogenous messages (the script is fixed) with endogenous processing. This is appropriate for public-health campaigns where the policymaker controls content but not the audience's engagement with the information or prior beliefs. The novel element is the interaction between the communication medium (which shifts engagement) and messenger familiarity (which generates a credibility bonus), operating through the same multiplicative structure.

8.6 Relation to Models of Source Credibility

A large literature—including Lupia and McCubbins (1998), Druckman (2001), and Gentzkow and Shapiro (2006)—models source credibility as determining the precision of the signal an agent receives. Our model nests this approach via the noise parameter σ_m^2, but adds two features: (i) engagement intensity as a separate amplifier, which can compensate for lower credibility when the messenger generates greater engagement; and (ii) familiarity as a credibility-augmenting channel, which captures the parasocial trust mechanism central to influencer communication (Horton and Wohl 1956).

8.7 Implications for Public-Health Policy

The model has two main policy implications. First, the medium choice is a high-leverage design decision. Since $F(a)$ is strictly increasing and, under Assumption 2, (weakly) concave in a over $[0, 1]$—strictly concave whenever $\sigma_m^2 > 0$—moving from video to VR shifts the belief-revision function upward, but with diminishing marginal returns to additional engagement. This implies that the largest gains from immersive technology accrue where baseline engagement is low; as VR becomes cheaper, the model favors targeting low-engagement segments of the audience over further intensifying engagement among already-engaged participants. Second, the optimal messenger choice depends on the audience's familiarity distribution. For audiences with low familiarity ($f_I^i \approx 0$), the expert's intrinsic credibility advantage dominates and experts should be preferred *provided* the influencer's engagement advantage remains below the threshold $\bar{\alpha}$ identified in Proposition 4(b); if the influencer generates an engagement advantage in excess of $\bar{\alpha}$, engagement alone can offset the credibility gap even in the absence of familiarity, and the influencer should be preferred instead.

9 Conclusion

Our theoretical and empirical analyses lead to three main conclusions. First, information itself matters: Even a short informational intervention meaningfully shifts beliefs, and the effect on trust in vaccines largely persists for at least six months. Second, how information is delivered matters. The same message delivered through immersive virtual reality produces larger effects than standard video, consistent with immersive delivery generating greater engagement with the information. Third, the identity of the messenger matters, but perceived expertise alone does not determine persuasive effectiveness (Lupia and McCubbins 1998; Druckman 2001; Gentzkow and Shapiro 2006): Scientists are consistently rated as more reliable and knowledgeable, yet influencers persuade just as effectively on average.

The mechanism evidence helps explain this apparent disconnect. For scientists, the mediation analysis indicates that perceived reliability and knowledgeability account for much of their persuasive effectiveness relative to influencers. Personal appeal, including attractiveness and passion, instead works in favor of the female influencer. The male influencer presents a different pattern: Despite generating the strongest persuasive effect, he receives relatively unfavorable evaluations across most measured characteristics, none of which explains his greater effectiveness. Instead, his advantage is partly associated with pre-existing familiarity and reputation. Participants familiar with him—who generally hold a favorable prior opinion—exhibit higher post-treatment trust, and accounting for prior familiarity attenuates the difference in persuasive effectiveness between the two influencers.

These results contribute to a broader literature on persuasion, source credibility, and information transmission (Lupia and McCubbins 1998; Druckman 2001; Gentzkow and Shapiro 2006; DellaVigna and Gentzkow 2010). Our findings suggest that persuasive effectiveness cannot be reduced to a single dimension of messenger quality. Consistent with persuasion frameworks in which source credibility shapes the weight individuals place on information (DellaVigna and Gentzkow 2010), our evidence suggests that effective credibility can derive from different messenger characteristics. For scientists, it is closely associated with perceived reliability and knowledgeability; for influencers, personal appeal and, for established public figures, pre-existing familiarity and reputation can provide alternative sources of persuasive influence. The same message may therefore be effective for different reasons depending on who delivers it and the prior relationship between the messenger and the audience.

This distinction is especially relevant in contemporary media environments, where individuals are repeatedly exposed to information through social media personalities and online platforms. Familiar public figures may enter a communication setting with an established relationship with their audience, allowing prior exposure and reputation to shape how their messages are received. At the same time, our results caution against interpreting familiarity as a necessary condition for influencer effectiveness: The less familiar female influencer also generates positive effects despite lacking the same degree of prior recognition. Familiarity and reputation may therefore strengthen persuasion without being sufficient, or necessary, to generate it.

A back-of-envelope calculation puts these results in policy terms. Conventional video delivers belief change at a cost of a few cents per person, while immersive delivery currently costs roughly thirty times more per unit of belief change—a multiple that has fallen by two-thirds since our data collection and continues to decline. Immersive delivery therefore passes a cost-benefit test precisely where our heterogeneity results locate the largest returns: Initially hesitant audiences, in settings such as schools and clinics where people are already assembled. On the messenger margin, an influencer's fee embeds organic distribution to a familiar audience, so the choice between messengers reduces to reach and familiarity rather than credibility—the same conclusion the model delivers analytically.[35]

More broadly, the results highlight the importance of communication design in public-health interventions (Dupas 2011; Bartoš, Bauer, Cahlíková, and Chytilová 2022; Alsan and Eichmeyer 2024). The effectiveness of information depends not only on what is communicated, but also on who communicates it and through which medium. Our findings further suggest that selecting effective messengers may require looking beyond conventional ideas of expertise or credibility: The attributes that make one messenger effective need not be those that make another effective. Understanding these interactions may become increasingly important as governments and institutions attempt to communicate scientific information in fragmented and highly mediated information environments.

References

Alatas, V., A. G. Chandrasekhar, M. Mobius, B. A. Olken, and C. Paladines (2024). "Do Celebrity Endorsements Matter? A Twitter Experiment Promoting Vaccination in Indonesia". In: *The Economic Journal* 134.659, pp. 913–933.

Alsan, M. and S. Eichmeyer (2024). "Experimental Evidence on the Effectiveness of Nonexperts for Improving Vaccine Demand". In: *American Economic Journal: Economic Policy* 16.1, pp. 394–414.

Andries, M., L. Bursztyn, T. Chaney, M. Djourelova, and A. Imas (2026). "In their Shoes: Empathy Through Information". In: *The Quarterly Journal of Economics*, forthcoming.

Ashworth, M., L. Thunström, T. L. Cherry, S. C. Newbold, and D. C. Finnoff (2021). "Emphasize Personal Health Benefits to Boost COVID-19 Vaccination Rates". In: *Proceedings of the National Academy of Sciences* 118.32, e2108225118.

Bartoš, V., M. Bauer, J. Cahlíková, and J. Chytilová (2022). "Communicating Doctors' Consensus Persistently Increases COVID-19 Vaccinations". In: *Nature* 606.7914, pp. 542–549.

Bazzi, S., M. Fiszbein, and M. Gebresilasse (2020). "Frontier Culture: The Roots and Persistence of "Rugged Individualism" in the United States". In: *Econometrica* 88.6, pp. 2329–2368.

Benjamini, Y. and Y. Hochberg (1995). "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 57.1, pp. 289–300.

Bordalo, P., N. Gennaioli, and A. Shleifer (2020). "Memory, Attention, and Choice". In: *The Quarterly Journal of Economics* 135.3, pp. 1399–1442.

Centers for Disease Control and Prevention (2025a). *Measles Cases and Outbreaks*. https://www.cdc.gov/measles/data-research/index.html.

[35] Appendix B.2 details the calculation.

Centers for Disease Control and Prevention (2025b). *Vaccination Coverage and Exemptions among Kindergartners*. URL: https://www.cdc.gov/schoolvaxview/data/index.html.

Chen, Z., W. Huang, P. McIntosh, I. Shepherd, and X. Zhao (2025). "Fear on the Plank? Virtual but More Real!" SSRN Working Paper. URL: https://ssrn.com/abstract=5269016.

Dai, H., S. Saccardo, M. A. Han, L. Roh, N. Raja, S. Vangala, H. Modi, S. Pandya, M. Sloyan, and D. M. Croymans (2021). "Behavioural Nudges Increase COVID-19 Vaccinations". In: *Nature* 597.7876, pp. 404–409.

De Veirman, M., V. Cauberghe, and L. Hudders (2017). "Marketing Through Instagram Influencers: The Impact of Number of Followers and Product Divergence on Brand Attitude". In: *International Journal of Advertising* 36.5, pp. 798–828.

Delavande, A., A. Shahab, J. Younas, and B. Zafar (2026). "Improving Childhood Immunization through Maternal Support: Evidence from a Randomized Intervention in Pakistan". In: *Journal of Development Economics* 182, p. 103808.

DellaVigna, S. and M. Gentzkow (2010). "Persuasion: Empirical Evidence". In: *Annual Review of Economics* 2.1, pp. 643–669.

Dioikitopoulos, E., R. Megalokonomou, T. Sartori, and Y. Zenou (Oct. 13, 2025). *Beyond the Facts: Who Changes Minds—Professors or Influencers?* AEA RCT Registry, AEARCTR-0016954. https://doi.org/10.1257/rct.16954-1.0.

Doulou, K., F. Bacopoulou, A. Michos, A. Tsolakidis, K. Mathioudakis, D. Zografopoulos, and G. Kourlaba (2026). "HPV Vaccination Coverage Among Children and Adolescents in Greece Using National Prescription Data". In: *Vaccine* 70, p. 128026.

Druckman, J. N. (2001). "On the Limits of Framing Effects: Who can Frame?" In: *The Journal of Politics* 63.4, pp. 1041–1066.

Dupas, P. (2011). "Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya". In: *American Economic Journal: Applied Economics* 3.1, pp. 1–34.

Eichengreen, B., O. Saka, and C. G. Aksoy (May 2024). "The Political Scar of Epidemics". In: *The Economic Journal* 134.660, pp. 1683–1700.

Fahey, O. B., D. J. Carr, E. M. Ocean, V. Nittas, C. P. Kane, P. M. Monti, and T. B. Wray (2024). "Black Social Media Influencers Engage Higher Percentages of Black Gay and Bisexual Men in Online Outreach for HIV Prevention Research Relative to Paid Ads". In: *Digital Health* 10, p. 20552076241253758.

Freberg, K., K. Graham, K. McGaughey, and L. A. Freberg (2011). "Who Are the Social Media Influencers? A Study of Public Perceptions of Personality". In: *Public Relations Review* 37.1, pp. 90–92.

Gabaix, X. (Nov. 2014). "A Sparsity-Based Model of Bounded Rationality". In: *The Quarterly Journal of Economics* 129.4, pp. 1661–1710.

— (2019). "Behavioral Inattention". In: *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 2. Elsevier, pp. 261–343.

Gentzkow, M. and J. M. Shapiro (2006). "Media Bias and Reputation". In: *Journal of Political Economy* 114.2, pp. 280–316.

Gountas, I., A. Skroumpelos, G. Trimis, N. Markoulatos, U. Sabale, S. Salomonsson, K. Saxena, L. Poughias, and A. Karokis (2024). *HPV Vaccination Deficit Among Girls in Greece: Where We Are Now?* DOI: 10.1016/j.jval.2024.10.1374.

Graeber, T., C. Roth, and F. Zimmermann (2024). "Stories, Statistics, and Memory". In: *The Quarterly Journal of Economics* 139.4, pp. 2181–2225.

Hellenic National Public Health Organization (EODY) (2025). *Epidemiological Data for Measles in Greece, 2004–2024*. Tech. rep.

Herrera, F., J. Bailenson, E. Weisz, E. Ogle, and J. Zaki (2018). "Building Long-Term Empathy: A Large-Scale Comparison of Traditional and Virtual Reality Perspective-Taking". In: *PLOS ONE* 13.10, e0204494.

Horton, D. and R. R. Wohl (1956). "Mass Communication and Para-social Interaction: Observations on Intimacy at a Distance". In: *Psychiatry* 19.3, pp. 215–229.

Jin, S. V., A. Muqaddam, and E. Ryu (2019). "Instafamous and Social Media Influencer Marketing". In: *Marketing Intelligence & Planning* 37.5, pp. 567–579.

de-Juan-Ripoll, C., J. L. Soler-Domínguez, J. Guixeres, M. Contero, N. Álvarez Gutiérrez, and M. Alcañiz (2018). "Virtual Reality as a New Approach for Risk Taking Assessment". In: *Frontiers in Psychology* 9, p. 2532.

Kamenica, E. and M. Gentzkow (2011). "Bayesian Persuasion". In: *American Economic Review* 101.6, pp. 2590–2615.

Knoll, J. and J. Matthes (2017). "The Effectiveness of Celebrity Endorsements: A Meta-analysis". In: *Journal of the Academy of Marketing Science* 45.1, pp. 55–75.

Lacle-Melendez, J., S. Silva-Medina, and J. Bacca-Acosta (2025). "Virtual and Augmented Reality to Develop Empathy: A Systematic Literature Review". In: *Multimedia Tools and Applications* 84.11, pp. 8893–8927.

Lee, D. S. (2009). "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects". In: *The Review of Economic Studies* 76.3, pp. 1071–1102.

Lupia, A. and M. D. McCubbins (1998). *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge University Press.

Mawdsley, S. E. (2016). "Salk Hops: Teen Health Activism and the Fight against Polio, 1955 – 1960". In: *Cultural and Social History* 13.2, pp. 249–265.

Milkman, K. L., M. S. Patel, L. Gandhi, H. N. Graci, D. M. Gromet, H. Ho, J. S. Kay, T. W. Lee, M. Akinola, J. Beshears, et al. (2021). "A Megastudy of Text-Based Nudges Encouraging Patients to Get Vaccinated at an Upcoming Doctor's Appointment". In: *Proceedings of the National Academy of Sciences* 118.20, e2101165118.

Mullainathan, S., J. Schwartzstein, and A. Shleifer (2008). "Coarse Thinking and Persuasion". In: *The Quarterly Journal of Economics* 123.2, pp. 577–619.

Murphy, K. M. and R. H. Topel (2002). "Estimation and Inference in Two-step Econometric Models". In: *Journal of Business & Economic Statistics* 20.1, pp. 88–97.

O'Brien, G., R. Ganjigunta, and P. S. Dhillon (2024). "Wellness Influencer Responses to COVID-19 Vaccines on Social Media: a Longitudinal Observational Study". In: *Journal of Medical Internet Research* 26, e56651.

OECD (2024). *OECD Survey on Drivers of Trust in Public Institutions 2024 Results: Building Trust in a Complex Policy Environment*. Tech. rep. Paris: OECD Publishing. DOI: 10.1787/9a20554b-en.

— (2025). *Health at a Glance 2025: OECD Indicators*. Paris: OECD Publishing. DOI: 10.1787/8f9e3f98-en.

Pagan, A. (1984). "Econometric Issues in the Analysis of Regressions with Generated Regressors". In: *International Economic Review* 25.1, pp. 221–247.

Serra-Garcia, M. and N. Szech (2023). "Incentives and Defaults Can Increase COVID-19 Vaccine Intentions and Test Demand". In: *Management Science* 69.2, pp. 1037–1049.

World Health Organization (2025). *Vaccines and Immunization*. https://www.who.int/health-topics/vaccines-and-immunization/.

World Health Organization and UNICEF (2024). *Global Immunization Efforts Have Saved at Least 154 Million Lives Over the Past 50 Years*. https://www.who.int/news/item/24-04-2024-global-immunization-efforts-have-saved-at-least-154-million-lives-over-the-past-50-years.

Figure 1: Variation of the Outcome Variables

Panel A: Trust in Vaccines, Pre-Treatment Distribution

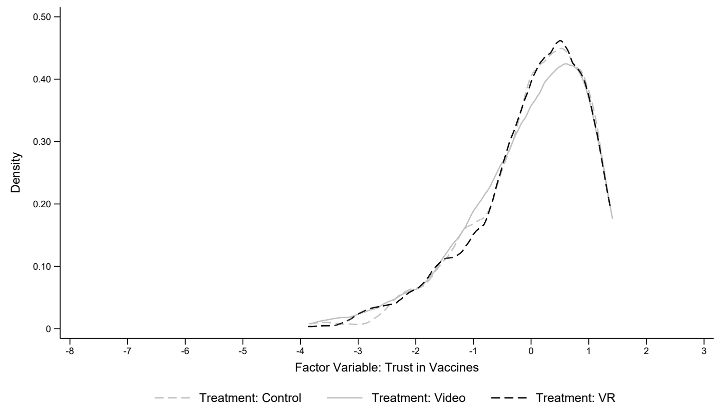

Panel B: Trust in Vaccines, Post-Treatment Distribution

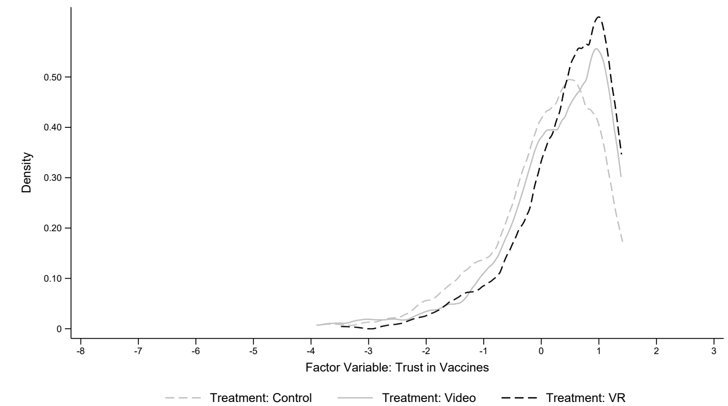

Panel C: Trust in the Healthcare System,
Pre-Treatment Distribution

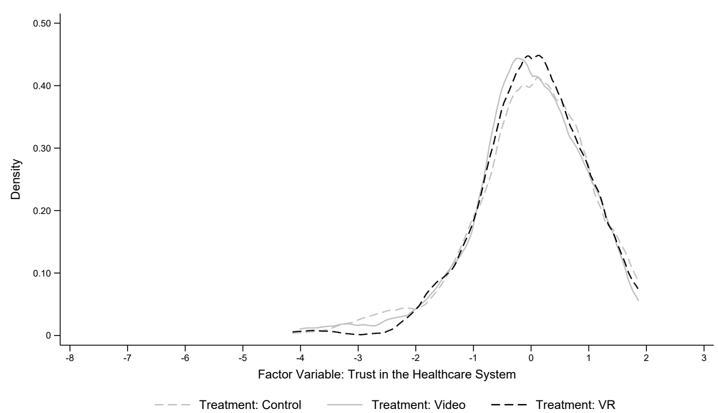

Panel D: Trust in the Healthcare System,
Post-Treatment Distribution

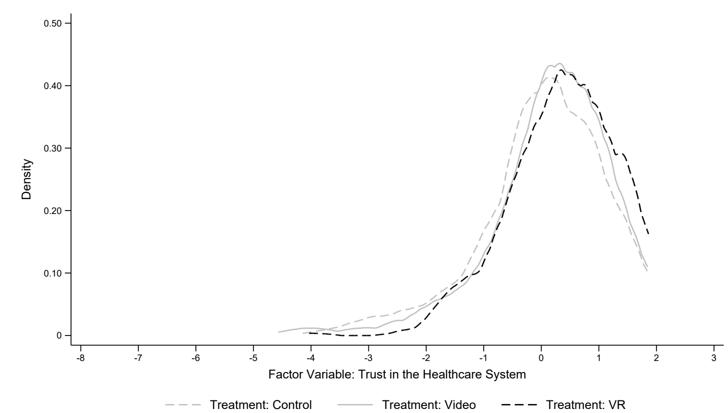

Notes: The figure shows the distribution of the two outcome indices by treatment group before and after the intervention. Panel A reports the pre-treatment distribution of the trust-in-vaccines factor by treatment group. Panel B reports the corresponding post-treatment distribution. Panel C reports the pre-treatment distribution of the healthcare-system trust factor by treatment group. Panel D reports the corresponding post-treatment distribution. The reported variables are derived using Principal Component Analysis, as described in Section 3.2.

Figure 2: Heterogeneous Effects: Information Treatment vs. Control

Panel A: Trust in Vaccines

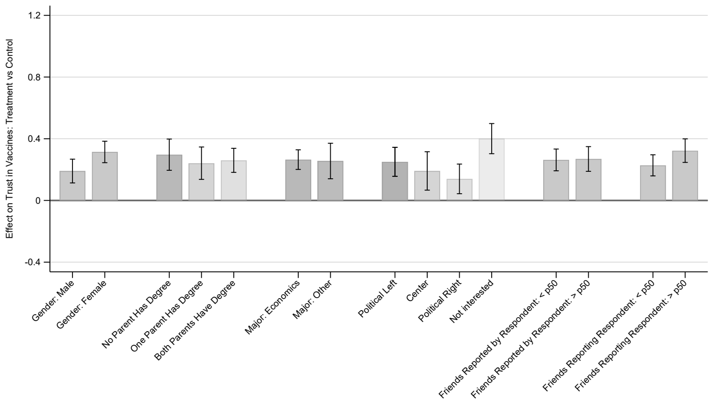

Panel B: Trust in Healthcare System

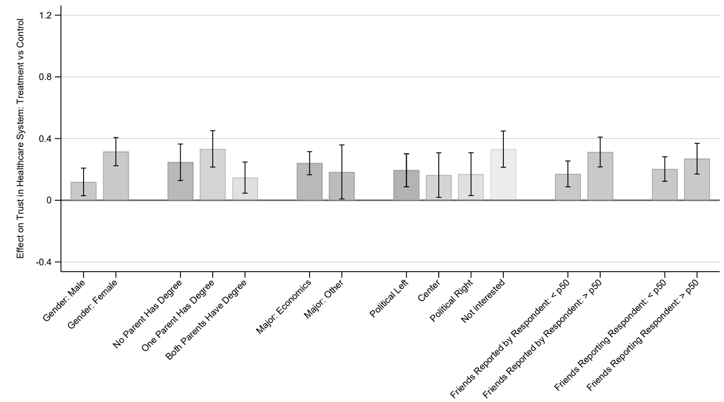

Panel C: Trust in Vaccines

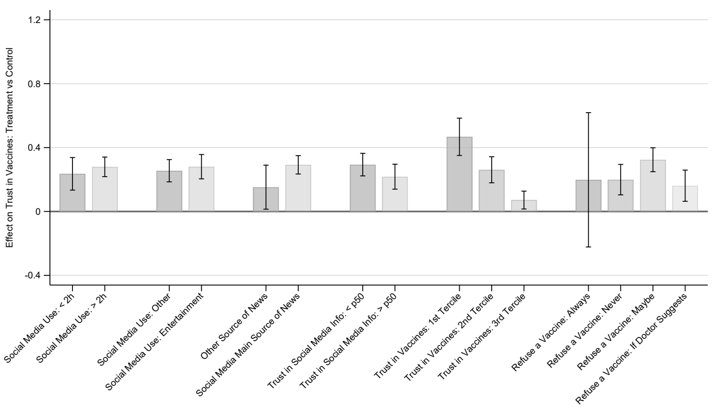

Panel D: Trust in Healthcare System

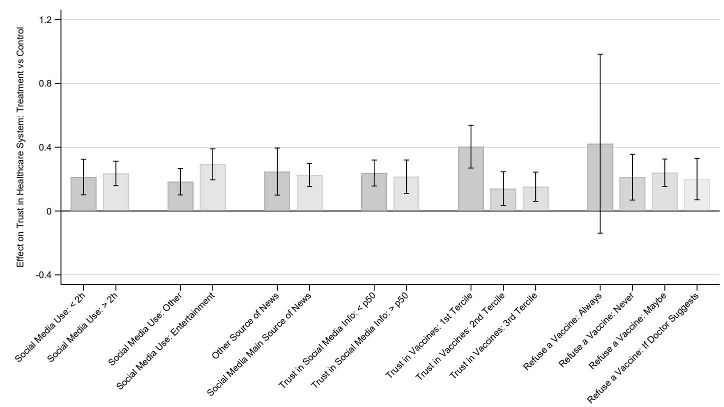

Notes: The figure plots subgroup-specific treatment-effect estimates and 95 percent confidence intervals for the comparison between treated participants and the control group. Panels A and B split the sample by demographic characteristics; Panels C and D split it by information-environment variables and baseline beliefs. Panels A and C use the trust-in-vaccines factor as outcome; Panels B and D use the healthcare-system trust factor.

Figure 3: Predicted Trust in Vaccines by Influencer Familiarity and Reputation

Panel A: Trust in Vaccines by Familiarity with the Influencer

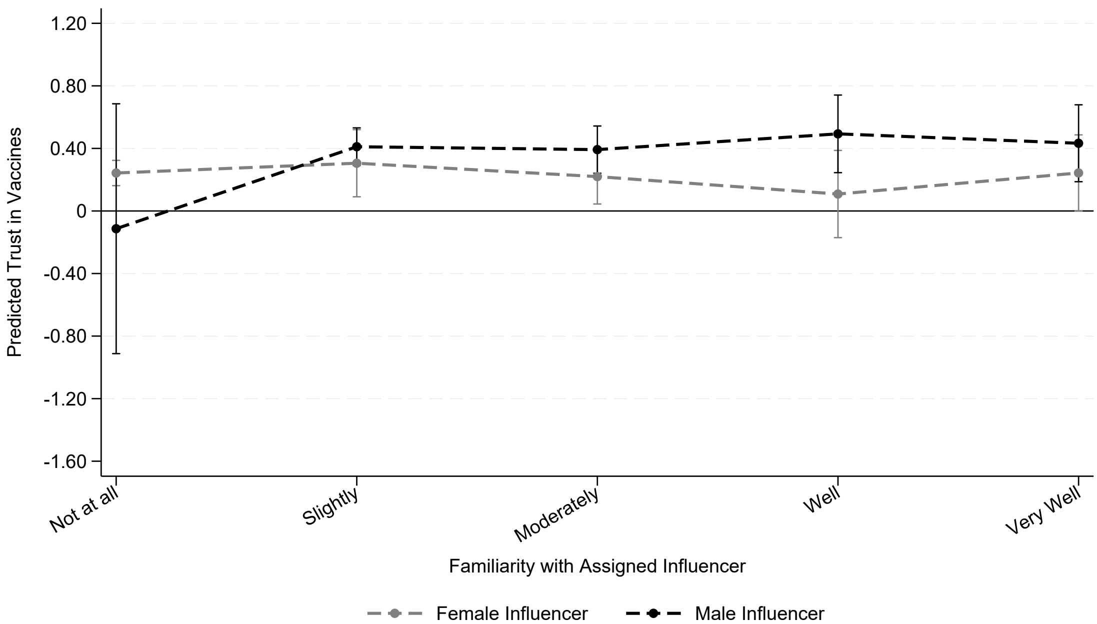

Panel B: Trust in Vaccines by Reputation of the Influencer

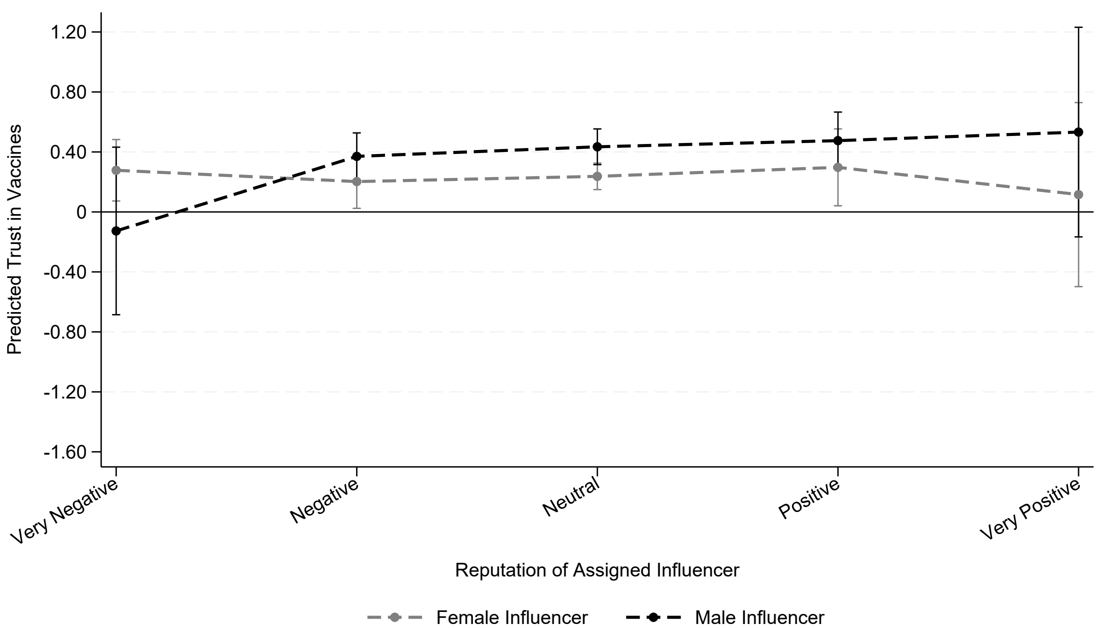

Notes: The figure plots regression-adjusted predicted levels of trust in vaccines (predictive margins) for participants assigned to the female and male influencers, evaluated across different levels of baseline familiarity and reputation. Panel A reports predicted levels of trust in vaccines separately for the female and male influencers across levels of prior familiarity. Panel B reports the corresponding predicted levels across levels of prior reputation. All estimates control for the baseline value of the outcome, individual characteristics, and prefecture, session-date, and session-number fixed effects. The 95 percent confidence intervals are constructed using standard errors clustered at the session level.

Table 1: Descriptive Statistics of the Baseline Sample

	Mean	SD	Min	Max	N
Share of Complete Surveys	0.98	0.15	0	1	1,245
Age (years)	20.41	3.78	18	59	1,218
Smokes More Than Once a Week	0.33	0.47	0	1	1,218
Exercises More Than Once a Week	0.93	0.25	0	1	1,218
Friends Reported	3.49	2.81	0	10	1,218
Friends Reporting Respondent	1.41	1.46	0	10	1,218

	Percentage	Frequency	N
Gender			1,218
Female	51	619	
Male	49	592	
Undisclosed/Other	1	7	
Study Level			1,218
Undergraduate Degree	95	1,156	
Graduate Degree	4	51	
Postgraduate or Other	1	11	
Department			1,218
Missing	27	325	
International and European Economic Studies	5	63	
Economics	19	232	
Administrative Science and Technology	8	94	
Business Organization and Administration	11	129	
Accounting and Finance	6	76	
Marketing and Communication	9	107	
Informatics	10	116	
Statistics	6	70	
Other	0	6	
Household Income (in Euro)			1,218
0–9,999	6	77	
10,000–19,999	16	190	
20,000–29,999	18	215	
30,000–39,999	15	186	
40,000–49,999	11	129	
More than 50,000	13	161	
I prefer not to answer	21	260	
Father's Education			1,218
High School or Lower	31	380	
Professional Training	19	229	
Bachelor's degree	29	348	
Master's degree	17	211	
Doctorate	4	50	
Mother's Education			1,218
High School or Lower	24	288	
Professional Training	15	183	
Bachelor's degree	39	470	
Master's degree	20	243	
Doctorate	3	34	
Political Ideology			1,218
Left/Center-Left	28	338	
Center	16	193	
Right/Center-Right	24	297	
Not Interested	32	390	
Treatment Arm Assignment			1,218
Control Group	34	420	
Video Group	34	420	
Virtual Reality Group	31	378	

Notes: The table reports descriptive statistics for the baseline demographic variables. Of the 1,245 participants who began the experiment, 1,218 (97.8 percent) completed the survey; the statistics are based on completers. The upper panel summarizes continuous variables; the lower panel reports percentages and frequencies for categorical variables.

Table 2: Summary of Balance Tests Across Treatment Groups

	Group Assignment			Treatment Assignment		
	Control	Video	VR	Treatment vs. Control	VR vs. Video	Scientist vs. Influencer
	(1)	(2)	(3)	(4)	(5)	(6)
Number of Performed Tests	42	42	42	42	42	42
# Significant at 1% Level	0	0	0	0	0	1
# Significant only at 5% Level	0	1	4	0	5	2
# Significant only at 10% Level	4	3	1	4	0	1
Share Significant at 1% Level	0.00	0.00	0.00	0.00	0.00	0.02
Share Significant only at 5% Level	0.00	0.02	0.10	0.00	0.12	0.05
Share Significant only at 10% Level	0.10	0.07	0.02	0.10	0.00	0.02

Notes: The table summarizes the share of statistically significant coefficients obtained from the balancing regressions across treatment groups. Counts of significant coefficients are mutually exclusive across thresholds (e.g., the 5 percent row excludes variables that are already significant at the 1 percent level). The full balancing regressions corresponding to columns (1) to (3) are reported in Tables A6 and A7, while those corresponding to columns (4) to (6) are reported in Tables A8 and A9.

Table 3: Effects of Information, Medium, and Messenger on Trust in Vaccines and the Healthcare System

	Trust in					
	Vaccines		Healthcare System		Placebo: Response Time	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Information Treatment vs. Control					
Treatment: Information	0.267***	0.265***	0.230***	0.229***	0.086	0.093
	(0.027)	(0.025)	(0.028)	(0.029)	(0.062)	(0.071)
Reference Group Mean	0.08	0.08	0.04	0.04	-0.04	-0.04
Reference Group SD	0.94	0.94	1.09	1.09	0.52	0.52
Observations	1,218	1,218	1,218	1,218	1,218	1,218
	Panel B: VR vs. Video					
Treatment: VR	0.127***	0.105***	0.120**	0.105**	0.199	0.189
	(0.039)	(0.038)	(0.057)	(0.053)	(0.127)	(0.141)
Reference Group Mean	0.26	0.26	0.17	0.17	-0.06	-0.06
Reference Group SD	0.94	0.94	1.07	1.07	0.50	0.50
Observations	798	798	798	798	798	798
	Panel C: Scientist vs. Influencer					
Treatment: Scientist	0.037	0.041	0.012	0.003	0.034	0.049
	(0.037)	(0.036)	(0.043)	(0.046)	(0.076)	(0.091)
Reference Group Mean	0.31	0.31	0.29	0.29	0.00	0.00
Reference Group SD	0.90	0.90	0.94	0.94	0.57	0.57
Observations	798	798	798	798	798	798
Session Day	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓

Notes: The table reports treatment effects on trust in vaccines, trust in the healthcare system, and the standardized time required to complete the corresponding survey section across three comparisons: (i) exposure to the information treatment (video or VR) vs. the control condition (Panel A, using equation (1)); (ii) VR vs. standard video among treated participants (Panel B, using equation (2)); and (iii) scientist vs. influencer speakers among treated participants (Panel C, using equation (3)). Columns (1), (3), and (5) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2), (4), and (6) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Effects of Messenger Gender and Type on Trust

	Trust in					
	Vaccines		Healthcare System		Placebo: Response Time	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment: Influencer – Male	0.137***	0.153***	0.017	0.046	-0.093	-0.077
	(0.047)	(0.049)	(0.058)	(0.057)	(0.064)	(0.066)
Treatment: Scientist – Female	0.096*	0.116**	0.034	0.048	0.030	0.062
	(0.050)	(0.051)	(0.061)	(0.062)	(0.148)	(0.179)
Treatment: Scientist – Male	0.114**	0.122**	0.006	0.007	-0.074	-0.058
	(0.048)	(0.049)	(0.052)	(0.055)	(0.083)	(0.079)
Reference Group Mean	0.17	0.17	0.20	0.20	0.02	0.02
Reference Group SD	0.99	0.99	0.95	0.95	0.55	0.55
Observations	798	798	798	798	798	798
Session Day	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓

Notes: The table reports treatment effects on trust in vaccines, trust in the healthcare system, and the standardized time required to complete the corresponding survey section across four treatment arms: female scientist, male scientist, female influencer, and male influencer, using equation (2). The omitted category is the female influencer treatment arm. Columns (1), (3), and (5) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2), (4), and (6) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Persistence of the Information Treatment

	Trust in			
	Vaccines (1)	Healthcare System (2)	Willing to Vaccinate (3)	Did Vaccinate (4)
	Panel A: First Follow-up (2 Months)			
Treatment: Information	0.240***	0.036	0.048	-0.011
	(0.048)	(0.053)	(0.037)	(0.016)
	[0.213, 0.264]	[0.016, 0.070]	[0.051, 0.069]	[-0.015, 0.012]
Reference Group Mean	-0.01	0.14	0.37	0.05
Reference Group SD	0.97	0.92	0.48	0.21
Observations	770	770	770	770
	Panel B: Second Follow-up (6 Months)			
Treatment: Information	0.184**	0.174*	0.113	-0.038
	(0.074)	(0.104)	(0.075)	(0.036)
	[0.121, 0.214]	[0.127, 0.234]	[0.094, 0.144]	[-0.039, 0.005]
Reference Group Mean	0.08	0.10	0.35	0.07
Reference Group SD	1.00	0.95	0.48	0.26
Observations	314	314	314	314
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓

Notes: The table reports the effect of exposure to the information treatment (video or VR) on outcomes measured in the two follow-up surveys, together with Lee bounds correcting for differential attrition (D. S. Lee 2009). Panel A reports outcomes from the first follow-up, administered approximately two months after treatment for participants in both experimental waves. Panel B reports outcomes from the second follow-up, administered approximately six months after treatment for participants in the second wave and eight to nine months after treatment for participants in the first wave. The outcomes are trust in vaccines (Column 1), trust in the healthcare system (Column 2), whether the participant reports being more willing to vaccinate (Column 3), and whether the participant received a vaccination between the baseline and the respective follow-up survey (Column 4). The omitted category is the control group. All specifications control for prefecture, session-day, session-order fixed effects, and the full set of participant characteristics described in Section 5. For trust in vaccines and trust in the healthcare system, we control for the corresponding pre-treatment measure. For willingness to vaccinate and vaccination uptake, we control for the pre-treatment measures of both trust in vaccines and trust in the healthcare system. Standard errors, clustered at the session level, are reported in parentheses. Lee bounds are reported in square brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Effects of Messenger Type on Impressions, Visual Attention, and Perceived Quality

	Impression of the							
	Speaker		Message		Visual Attention		Perceived Speaker Quality	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Scientist vs. Influencer							
Treatment: Scientist	0.386***	0.388***	0.076	0.059	-0.068	-0.042	0.391***	0.384***
	(0.074)	(0.074)	(0.072)	(0.071)	(0.045)	(0.051)	(0.073)	(0.073)
Reference Group Mean	-0.22	-0.22	-0.05	-0.05	0.40	0.40	-0.21	-0.21
Reference Group SD	1.13	1.13	1.06	1.06	0.34	0.34	1.10	1.10
Observations	798	798	798	798	378	378	798	798
	Panel B: Breakdown by Speaker Identity							
Treatment: Influencer – Male	-0.103	-0.081	0.045	0.059	0.082	0.080	-0.446***	-0.428***
	(0.108)	(0.112)	(0.108)	(0.114)	(0.062)	(0.060)	(0.107)	(0.111)
Treatment: Scientist – Female	0.289***	0.297***	0.019	0.014	-0.011	0.021	0.143	0.148
	(0.106)	(0.108)	(0.098)	(0.103)	(0.051)	(0.061)	(0.104)	(0.101)
Treatment: Scientist – Male	0.394***	0.411***	0.186*	0.174	-0.052	-0.035	0.212**	0.208**
	(0.103)	(0.102)	(0.107)	(0.108)	(0.057)	(0.061)	(0.087)	(0.093)
Reference Group Mean	-0.19	-0.19	-0.12	-0.12	0.38	0.38	-0.04	-0.04
Reference Group SD	1.17	1.17	1.09	1.09	0.34	0.34	1.12	1.12
Observations	798	798	798	798	378	378	798	798
Session Day	✓	✓	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓

Notes: The table reports treatment effects on mechanism outcomes among treated participants. The outcomes are impression of the speaker and impression of the message (both measured on a scale from -5, very negative, to 5, very positive), visual attention (measured as the share of seconds spent looking at the speaker, available only for participants assigned to the VR treatment), and a principal-component index of perceived speaker quality. Panel A reports estimates from equation (3), comparing participants exposed to a scientist vs. an influencer, with all influencer speakers (male and female) serving as the omitted reference group. Panel B reports estimates from equation (2), providing a full breakdown by speaker identity (female influencer, male influencer, female scientist, and male scientist), with the female influencer serving as the omitted reference category. All outcome variables are standardized using the mean and standard deviation of the overall sample (except for visual attention, which was collected only for participants in the VR treatment and is measured as the share of seconds the participant's eyes were on the speaker out of the total seconds the speaker was on screen). All specifications control for baseline trust in vaccines and baseline trust in the healthcare system, as well as prefecture, session-day, and session-order fixed effects. Columns (2), (4), (6), and (8) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 7: Effects of Messenger Type on Perceived Speaker Quality

	Perceived Speaker Quality: The Speaker Was						
	Charismatic	Persuasive	Reliable	Knowledgeable	Sympathetic	Attractive	Passionate
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: Scientist vs. Influencer						
Treatment: Scientist	0.054	0.212***	0.651***	0.978***	0.245***	-0.151**	0.113
	(0.073)	(0.072)	(0.072)	(0.064)	(0.076)	(0.074)	(0.076)
Reference Group Mean	-0.04	-0.13	-0.35	-0.53	-0.13	0.10	-0.05
Reference Group SD	1.09	1.05	1.04	1.02	1.13	1.02	1.05
Observations	798	798	798	798	798	798	798
	Panel B: Breakdown by Speaker Identity						
Treatment: Influencer – Male	-0.394***	-0.094	-0.111	-0.314***	-0.193*	-0.617***	-0.758***
	(0.115)	(0.106)	(0.107)	(0.112)	(0.115)	(0.096)	(0.109)
Treatment: Scientist – Female	-0.164	0.161	0.590***	0.823***	0.085	-0.453***	-0.306***
	(0.102)	(0.102)	(0.098)	(0.092)	(0.108)	(0.105)	(0.089)
Treatment: Scientist – Male	-0.109	0.184*	0.613***	0.824***	0.222**	-0.449***	-0.215**
	(0.092)	(0.093)	(0.093)	(0.100)	(0.100)	(0.100)	(0.094)
Reference Group Mean	0.10	-0.12	-0.33	-0.40	-0.06	0.37	0.29
Reference Group SD	1.12	1.05	1.08	1.03	1.21	1.02	0.93
Observations	798	798	798	798	798	798	798
Session Day	✓	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓

Notes: The table reports treatment effects on perceived speaker quality among treated participants. The outcomes are whether participants perceive the speaker as "Charismatic," "Persuasive," "Reliable," "Knowledgeable," "Sympathetic," "Attractive," and "Passionate," on a scale from 0 ("not at all") to 10 ("extremely"). Panel A reports estimates from equation (3), comparing participants exposed to a scientist vs. an influencer. Panel B reports estimates from equation (2), providing a full breakdown by speaker identity (female influencer, male influencer, female scientist, and male scientist), with the female influencer serving as the omitted reference category. All outcome variables are standardized using the mean and standard deviation of the overall sample. All specifications control for baseline trust in vaccines and baseline trust in the healthcare system, as well as prefecture, session-day, and session-order fixed effects. We also control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: Model Predictions and Corresponding Evidence

Model Prediction	Empirical Counterpart	Result
Information shifts beliefs toward the signal (Proposition 1)	Information treatment vs. control	+0.27 SD in vaccine trust (Table 3, Panel A)
Engagement-enhancing media amplify belief revision, holding content fixed (Proposition 2)	VR vs. video among treated	+0.13 SD (Table 3, Panel B)
Familiarity can offset the expertise advantage of experts (Proposition 3)	Scientist vs. influencer, average effects	No detectable difference; equivalence bounds rule out differences above 0.11 SD (Table 3, Panel C)
The familiarity gradient is steep only under positive prior reputation (Corollary 1)	Post-treatment trust by familiarity, by influencer	Steep for the male influencer, flat for the female influencer (Figure 3; Table A29)
Revision is decreasing in baseline beliefs (Proposition 6(iv))	Baseline-trust terciles	Effects decline across terciles (Figure 2)
Belief revision partially persists (Proposition 5)	Two- and Six-month Follow-up	90 and 62 percent of the vaccine-trust effect remains (Table 5)

Notes: The table summarizes the model's main predictions and the corresponding experimental evidence. Effect sizes are from the specifications without additional controls, reported in the odd-numbered columns of the cited tables; SD denotes standard deviations of the corresponding outcome index.

Online Appendix

A Additional Figures and Tables

Figure A1: Vaccination Coverage Trends

Panel A: Measles Vaccination Coverage (1st Dose)

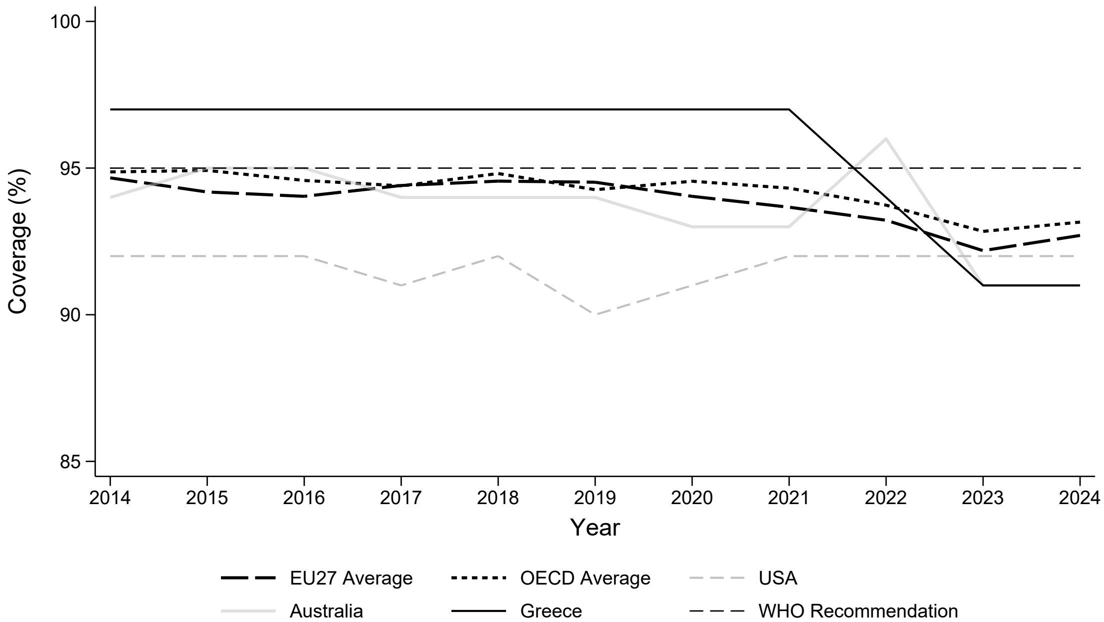

Panel B: HPV Vaccination Coverage for Girls by Age 15

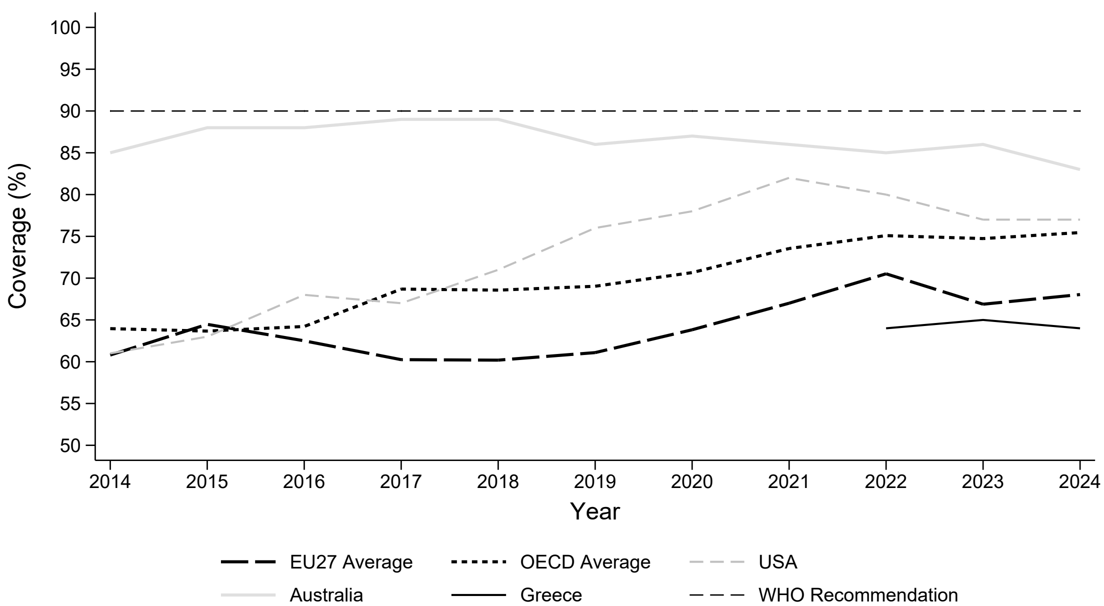

Note: For Greece we report estimates from Doulou et al. (2026) based on the doses dispensed using data from the electronic prescription database.

Notes: The figure reports vaccination coverage trends for selected countries and aggregates. Panel A reports first-dose measles vaccination coverage over time using World Health Organization data for Greece, the United States, Australia, the OECD average, and the EU27 average. The horizontal line marks the 95 percent threshold recommended by WHO, which is generally associated with herd immunity against measles. Panel B reports human papillomavirus (HPV) vaccination coverage among girls by age 15 for the same countries and aggregates. The horizontal line marks the 90 percent threshold recommended by the World Health Organization. HPV coverage data for all countries except Greece are obtained from the World Health Organization. For Greece, we report the estimate from Doulou et al. (2026), who use dispensed-dose data from the Greek electronic prescription database.

Figure A2: Public-Health Message Used in the Experiment

This message is part of a joint public-health initiative by the Athens University of Economics and Business and Monash University.

Vaccines are among the safest and most effective tools in medicine. Over the past 50 years, immunization has saved an estimated 154 million lives worldwide, and each year routine vaccination prevents around 3.5–5 million deaths. When enough people are vaccinated, we also protect those who cannot be—such as infants, the elderly, and individuals with weakened immune systems.

Greece performs well overall, but there are warning signs. Measles requires at least 95% coverage, yet current coverage in Greece is about 83%, leaving room for outbreaks. In 2024, Greece recorded 440 pertussis cases, the highest number since 2004, with infants under six months most at risk.

For students, the HPV vaccine protects against cancers of the cervix and throat. Greece introduced gender-neutral HPV vaccination in 2023, free of charge in 7th grade. Coverage remains below the WHO 90% target. If you missed your HPV doses during high school, now is the time to complete the series.

Stay up to date on your MMR (two doses), Tdap booster, and recommended meningococcal vaccines. These evidence-based steps protect campuses, teammates, and communities.

By closing immunity gaps, we protect our families and our future. Thank you for doing your part.

Notes: The figure reports the script that was read by the scientists and influencers in the vaccination-information treatment. The same text was used across all information-treatment arms; only the communication medium (VR vs. video) and messenger identity varied.

Figure A3: Insights from the VR: Median Exposure and Visual Attention by Speaker

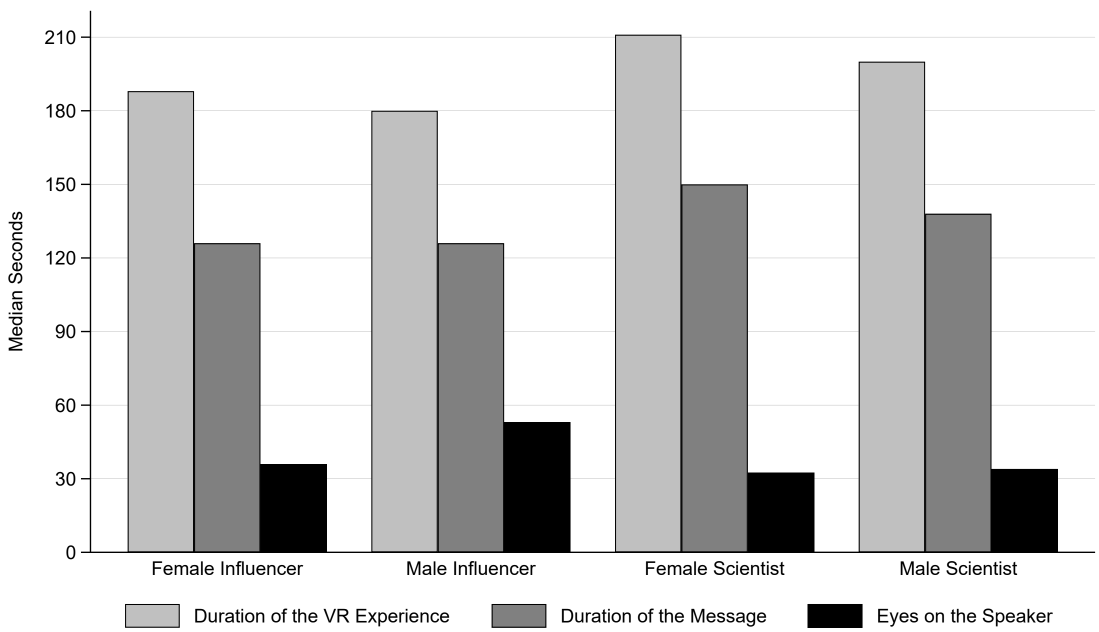

Notes: The figure reports median exposure time and visual attention by speaker within the VR treatment arm. The figure is descriptive: it compares the duration of each message with the number of seconds participants' gaze remained on the speaker during its delivery.

Figure A4: Participation Days and Session Order

Panel A: Participation by Day

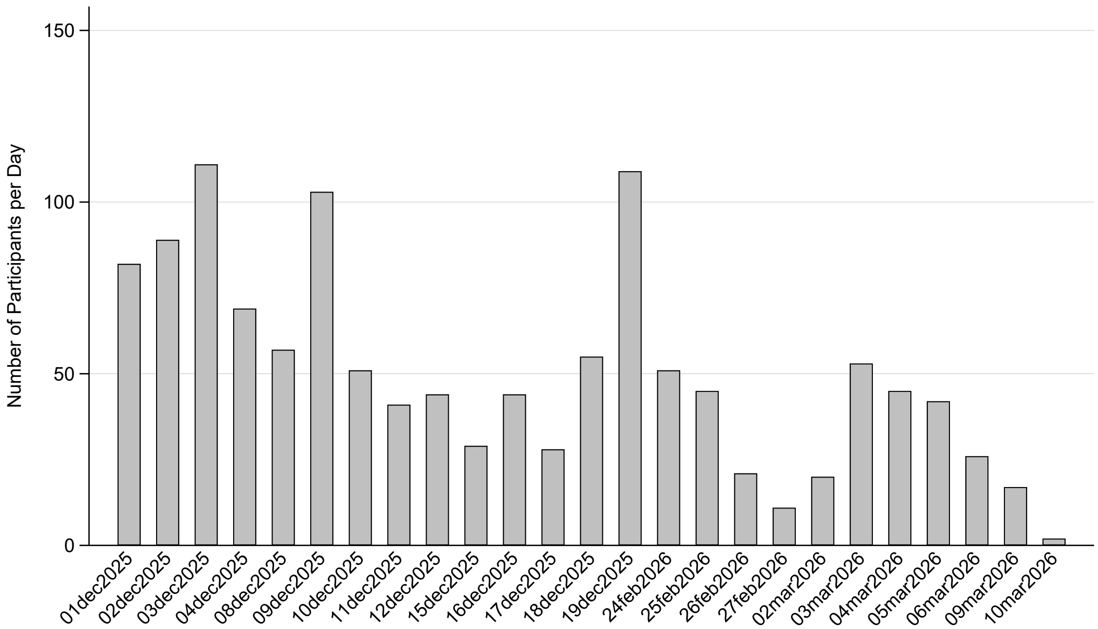

Panel B: Median Participation by Session

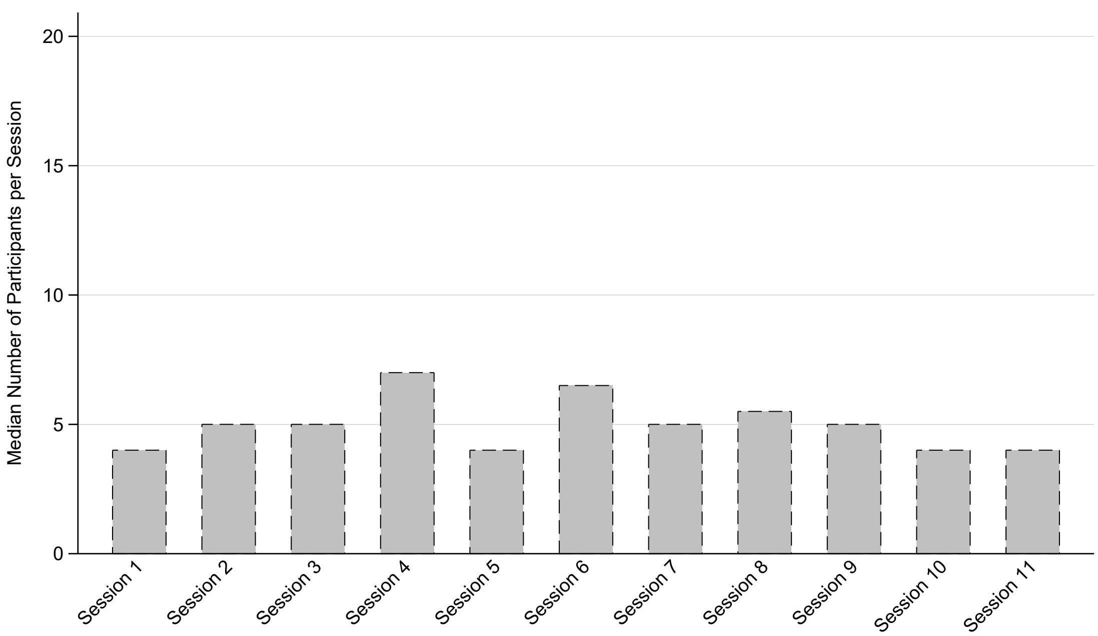

Notes: The figure reports participation over the data-collection period. Panel A reports the total number of participants by day of data collection during the experimental sessions conducted at the Athens University of Economics and Business. Panel B reports the median number of participants across sessions within each day. Sessions were conducted multiple times per day, and participants chose among available sessions corresponding to their randomized treatment assignment.

Figure A5: Heterogeneous Effects of VR Relative to Video on Trust

Panel A: Trust in Vaccines

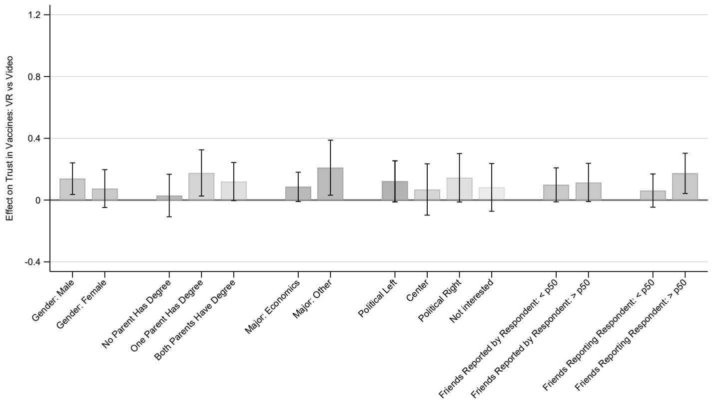

Panel B: Trust in the Healthcare System

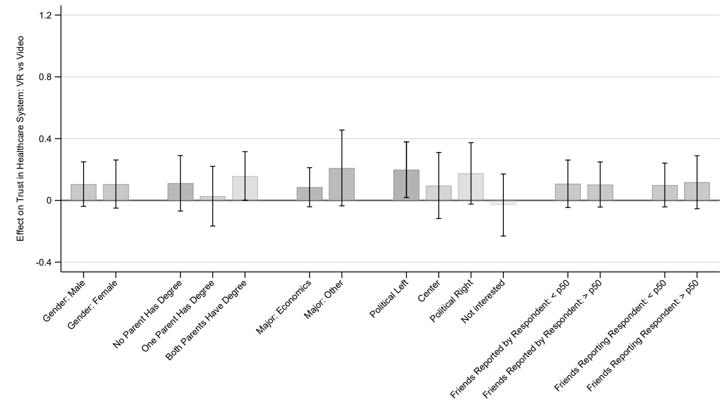

Panel C: Trust in Vaccines

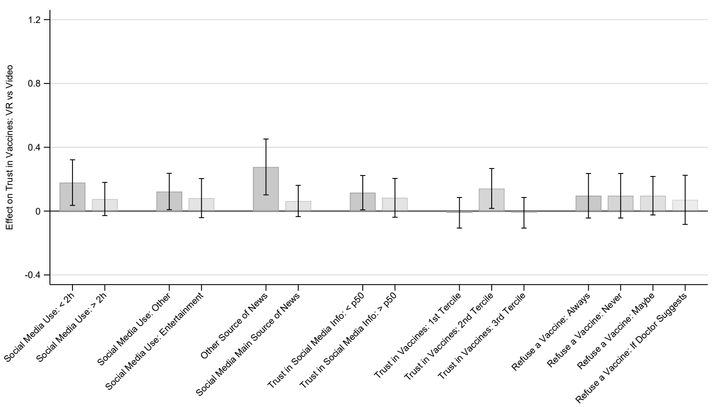

Panel D: Trust in the Healthcare System

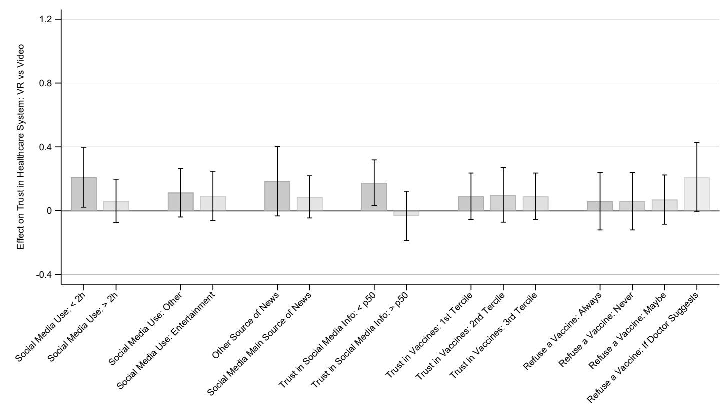

Notes: The figure plots subgroup-specific treatment-effect estimates and 95 percent confidence intervals for the comparison between treated participants who receive the message through virtual reality and treated participants who receive it through a standard two-dimensional video. Panels A and B split the sample by demographic characteristics; Panels C and D split it by information-environment variables and baseline beliefs. Panels A and C use the trust-in-vaccines factor as outcome; Panels B and D use the healthcare-system trust factor.

Figure A6: Heterogeneous Effects of Scientist Relative to Influencer

Panel A: Trust in Vaccines

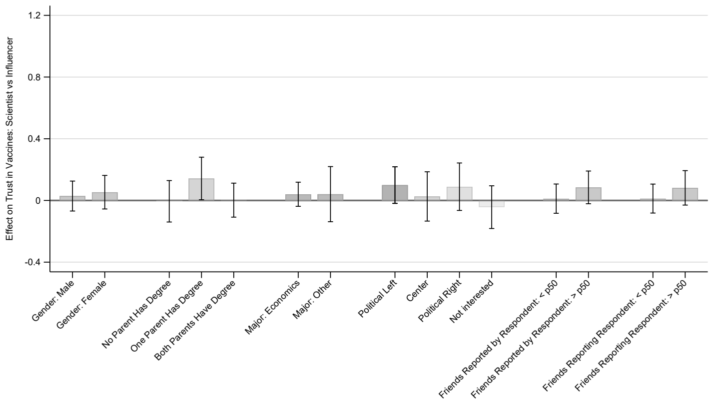

Panel B: Trust in the Healthcare System

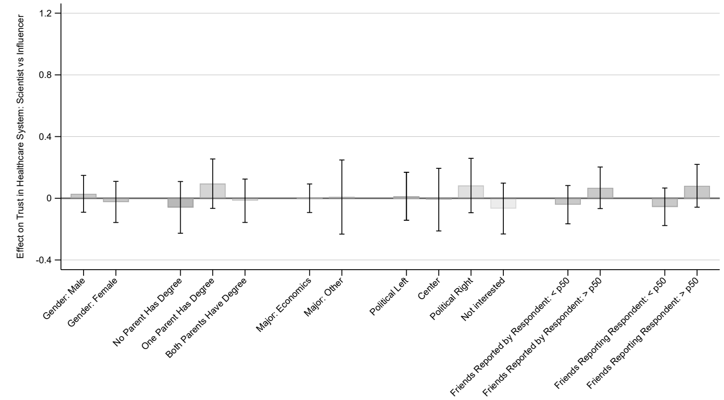

Panel C: Trust in Vaccines

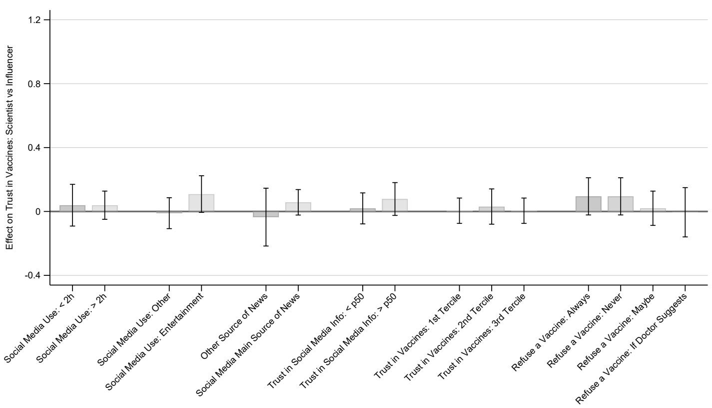

Panel D: Trust in the Healthcare System

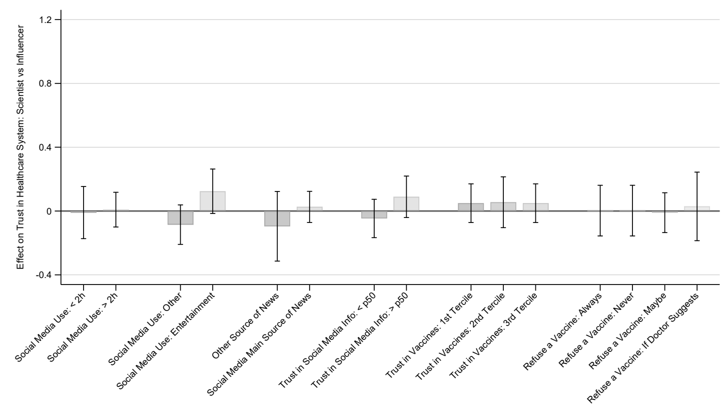

Notes: The figure plots subgroup-specific treatment-effect estimates and 95 percent confidence intervals for the comparison between treated participants exposed to a scientist and treated participants exposed to an influencer. Panels A and B split the sample by demographic characteristics; Panels C and D split it by information-environment variables and baseline beliefs. Panels A and C use the trust-in-vaccines factor as outcome; Panels B and D use the healthcare-system trust factor.

Table A1: Construction of the Trust Indices: Principal Component Analysis

	Component Weights	
	Trust in Vaccines	Trust in Healthcare System
Vaccination Efficacy	0.47	0.02
Vaccination Safety	0.50	0.01
Vaccination Important (Self)	0.51	-0.02
Vaccination Important (Others)	0.51	-0.02
Trust in Medicine	0.09	0.51
Trust in Doctors	-0.01	0.62
Trust in Health Authorities	-0.06	0.60
Eigenvalue	3.43	2.01
Variance Explained	0.49	0.29
Cronbach's α (Pre-treatment)	0.94	0.75
Cronbach's α (Post-treatment)	0.94	0.80
Cronbach's α (Follow-up #1)	0.94	0.77
Cronbach's α (Follow-up #2)	0.95	0.81

Notes: The table reports rotated component weights from a principal component analysis estimated using pre-treatment responses. Two orthogonal components are extracted and rotated using the varimax criterion. The resulting component weights and the pre-treatment item normalization are used to construct the indices at pre-treatment, post-treatment, and follow-up, ensuring a common definition of the outcomes across survey waves. The first component is interpreted as *Trust in Vaccines*, while the second is interpreted as *Trust in Healthcare System*. Variance explained represents the proportion of total variance accounted for by each rotated component. Cronbach's α is reported separately for each construct and survey wave.

Table A2: Construction of the Perceived Speaker Quality Index: Principal Component Analysis

	Perceived Speaker Quality (1)	Unexplained (2)
Charismatic	0.40	0.31
Persuasive	0.40	0.30
Reliable	0.40	0.31
Knowledgeable	0.36	0.42
Sympathetic	0.40	0.28
Attractive	0.34	0.49
Passionate	0.34	0.49
Variance Explained	0.63	
Cronbach's α	0.90	

Notes: The table reports the component weights on the first principal component from a principal component analysis of seven perceived speaker characteristics: charismatic, persuasive, reliable, knowledgeable, sympathetic, attractive, and passionate. The first principal component is used to construct the Perceived Speaker Quality index. All items receive positive and similar weights on the first component. Cronbach's α is reported for the seven-item scale.

Table A3: Participants' Information Environment and Social Media Use

	Mean	SD	Min	Max	N
Level of Trust in Information on Social Media (0 to 10)	4.81	1.61	0	10	1,218
Female Influencer					1,218
Do They Know Her? (1 to 5)	1.84	1.10	1	5	
Do They Have a Good Opinion of Her? (-2 to 2)	-0.41	0.92	-2	2	
Male Influencer					1,218
Do They Know Him? (1 to 5)	3.24	1.10	1	5	
Do They Have a Good Opinion of Him? (-2 to 2)	-0.07	1.01	-2	2	

	Percentage	Frequency	N
Daily Time Spent on Social Media			1,218
Under 30 min	2	28	
30-59 min	6	78	
1–2 hours	20	247	
2–3 hours	38	460	
3–5 hours	26	317	
More than 5 hours	7	88	
Main Purpose of Social Media Use			1,218
Keep in Touch with Friends	47	571	
Follow News/Events	8	102	
Express Opinions	0	3	
Entertainment	42	516	
Professional Networking	2	20	
Don't Use Social Media	0	6	
Primary Source of Information			1,218
Newspapers	10	125	
Television	5	55	
Social Media Platforms	81	990	
Don't Watch News	4	48	

Notes: The table reports descriptive statistics for baseline beliefs about information sources, media use, and related background characteristics. The upper panel summarizes continuous variables; the lower panel reports percentages and frequencies for categorical variables.

Table A4: Participants' Beliefs about Vaccines and the Healthcare System

	Mean	SD	Min	Max	N
To What Extent Do You Think Vaccines Are (0 to 10)					1,218
Effective	7.38	2.04	0	10	
Safe	7.07	2.10	0	10	
Important for Own Health	7.43	2.30	0	10	
Important for Public Health	7.99	2.08	0	10	
To What Extent Do You Trust (0 to 10)					1,218
Modern Medicine	7.84	1.72	0	10	
Doctors	7.72	1.64	0	10	
Public Health Authorities	5.87	2.20	0	10	
Factor Scores: Trust in					1,218
Vaccines	0.00	1.00	-4	1	
Healthcare System	0.00	1.00	-4	2	
Response Time (seconds)	44.63	24.08	12	251	1,218

	Percentage	Frequency	N
Would You Accept a Vaccine Recommended by Authorities			1,218
No	3	37	
Yes	19	227	
It Depends on the Vaccine	58	710	
It Depends on my Doctor's Advice	20	244	

Notes: The table reports descriptive statistics for the pre-treatment measures related to trust in vaccines and the healthcare system. The upper panel summarizes continuous variables; the lower panel reports percentages and frequencies for categorical variables.

Table A5: Descriptive Statistics of the Eye-Tracking Data

	Mean	SD	Min	Max	N
Number of Seconds:					378
In the VR Experience	226.75	87.89	134	737	
Watching the Speaker	136.81	15.22	124	280	
With Eyes on the Speaker	50.63	47.09	0	210	
With Eyes off the Speaker	86.17	45.60	0	150	
	Panel A: Female Influencer				
Number of Seconds:					97
In the VR Experience	219.61	84.59	135	591	
Watching the Speaker	128.75	15.55	124	252	
With Eyes on the Speaker	49.44	45.24	0	156	
With Eyes off the Speaker	79.31	42.08	0	127	
	Panel B: Male Influencer				
Number of Seconds:					73
In the VR Experience	207.07	68.74	134	562	
Watching the Speaker	127.85	18.06	124	280	
With Eyes on the Speaker	56.16	45.64	0	210	
With Eyes off the Speaker	71.68	41.82	0	126	
	Panel C: Female Scientist				
Number of Seconds:					94
In the VR Experience	250.38	104.32	165	737	
Watching the Speaker	151.07	9.94	149	234	
With Eyes on the Speaker	55.85	55.20	0	187	
With Eyes off the Speaker	95.22	52.66	0	150	
	Panel D: Male Scientist				
Number of Seconds:					114
In the VR Experience	225.94	83.46	152	631	
Watching the Speaker	137.64	0.55	136	138	
With Eyes on the Speaker	43.81	41.63	0	136	
With Eyes off the Speaker	93.83	41.63	2	138	

Notes: The table reports descriptive statistics for the VR eye-tracking measures. The first block summarizes the VR sample, and the subsequent panels break the same measures down by speaker identity. The eye-tracking sample includes participants whose individual identifiers could be reliably matched to their survey records. Where identifiers were entered incorrectly during the VR experience or in the survey, we attempted to trace the observations back to the corresponding participants. Overall, we were able to match eye-tracking and survey records for 378 participants.

Table A6: Balancing Exercise #1

	Control		Video		VR	
	(1)	(2)	(3)	(4)	(5)	(6)
Age (years)	0.009**	0.006*	-0.002	-0.001	-0.007**	-0.006**
	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)
Friends Reported	0.002	0.002	0.004	0.007	-0.006	-0.009*
	(0.008)	(0.006)	(0.007)	(0.005)	(0.007)	(0.005)
Friends Reporting Respondent	-0.016	-0.017	-0.004	-0.003	0.020*	0.019**
	(0.012)	(0.010)	(0.012)	(0.011)	(0.012)	(0.010)
Smoking: More Than Once a Week	-0.038	-0.046*	0.051	0.046*	-0.013	0.000
	(0.032)	(0.027)	(0.032)	(0.026)	(0.031)	(0.027)
Exercising: More Than Once a Week	-0.113**	-0.090*	0.061	0.038	0.051	0.052
	(0.049)	(0.048)	(0.043)	(0.042)	(0.045)	(0.040)
Trust in Social Media	0.013	0.009	-0.015*	-0.007	0.002	-0.002
	(0.009)	(0.008)	(0.008)	(0.007)	(0.009)	(0.007)
Female	0.031	0.032	0.002	-0.003	-0.033	-0.030
	(0.030)	(0.025)	(0.029)	(0.022)	(0.029)	(0.025)
Income: Above Median	-0.038	-0.040	0.048	0.043*	-0.010	-0.004
	(0.032)	(0.032)	(0.030)	(0.026)	(0.032)	(0.028)
Political Ideology: Left/Center-Left	-0.002	0.002	0.002	0.004	0.000	-0.006
	(0.034)	(0.032)	(0.033)	(0.027)	(0.033)	(0.028)
Political Ideology: Center	-0.010	-0.017	0.052	0.019	-0.042	-0.002
	(0.034)	(0.033)	(0.037)	(0.030)	(0.036)	(0.032)
Political Ideology: Right/Center-Right	-0.033	-0.032	-0.011	0.011	0.044	0.020
	(0.036)	(0.032)	(0.034)	(0.028)	(0.034)	(0.030)
Enrolled in: Undergraduate Degree	-0.096	-0.118*	0.108	0.084	-0.013	0.034
	(0.078)	(0.069)	(0.068)	(0.065)	(0.070)	(0.065)
Enrolled in: Graduate Degree	0.090	0.112	-0.114	-0.062	0.024	-0.051
	(0.087)	(0.075)	(0.074)	(0.072)	(0.080)	(0.077)
Father's Education: Above BSc	0.010	0.016	-0.003	0.006	-0.007	-0.022
	(0.027)	(0.024)	(0.026)	(0.021)	(0.025)	(0.019)
Mother's Education: Above BSc	-0.005	0.001	-0.002	-0.003	0.008	0.002
	(0.024)	(0.023)	(0.026)	(0.022)	(0.023)	(0.020)
Use Social Media: Friends	0.014	0.024	0.004	-0.018	-0.017	-0.006
	(0.027)	(0.025)	(0.030)	(0.025)	(0.029)	(0.024)
Use Social Media: News	0.041	0.012	0.009	0.019	-0.050	-0.031
	(0.051)	(0.045)	(0.049)	(0.038)	(0.051)	(0.043)
Use Social Media: Entertainment	-0.027	-0.026	-0.006	0.011	0.033	0.015
	(0.026)	(0.023)	(0.027)	(0.023)	(0.025)	(0.021)
Use Social Media: Professional	0.005	-0.031	0.005	0.032	-0.011	-0.001
	(0.107)	(0.095)	(0.114)	(0.115)	(0.103)	(0.104)
Use Social Media: No Use	-0.012	-0.052	-0.012	0.004	0.023	0.048
	(0.188)	(0.184)	(0.189)	(0.136)	(0.188)	(0.158)
More than 2h a Day on Social Media	0.023	0.031	-0.005	-0.016	-0.018	-0.015
	(0.030)	(0.026)	(0.030)	(0.023)	(0.030)	(0.024)
Session Day		✓		✓		✓
Session Order		✓		✓		✓
Prefecture FE		✓		✓		✓

Notes: The table reports balance checks across treatment groups for baseline demographic characteristics. Each coefficient is obtained from a separate regression of the treatment indicator on the corresponding baseline characteristic. Odd-numbered columns report specifications without fixed effects, while even-numbered columns include prefecture, session-day, and session-order fixed effects. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A7: Balancing Exercise #2

	Control		Video		VR	
	(1)	(2)	(3)	(4)	(5)	(6)
Familiarity with Female Influencer	-0.006	-0.005	0.015	0.017	-0.010	-0.012
	(0.014)	(0.011)	(0.013)	(0.012)	(0.012)	(0.010)
Familiarity with Male Influencer	0.005	0.014	-0.026	-0.027**	0.021	0.012
	(0.016)	(0.014)	(0.016)	(0.013)	(0.016)	(0.012)
Opinion of Female Influencer	0.009	0.019	0.013	0.005	-0.023	-0.024**
	(0.015)	(0.014)	(0.014)	(0.012)	(0.015)	(0.012)
Opinion of Male Influencer	0.000	0.010	0.004	-0.002	-0.003	-0.008
	(0.014)	(0.013)	(0.016)	(0.011)	(0.015)	(0.011)
Response Time (seconds)	0.001	0.001	-0.001	-0.001	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Vaccines: Effective	0.003	0.001	-0.006	-0.004	0.003	0.003
	(0.006)	(0.005)	(0.006)	(0.005)	(0.006)	(0.005)
Vaccines: Safe	0.004	0.002	-0.007	-0.004	0.002	0.002
	(0.006)	(0.005)	(0.007)	(0.006)	(0.006)	(0.005)
Vaccines: Important for Own Health	0.004	-0.001	-0.007	-0.001	0.003	0.002
	(0.006)	(0.005)	(0.006)	(0.005)	(0.006)	(0.004)
Vaccines: Important for Public Health	0.002	-0.003	-0.004	-0.001	0.002	0.003
	(0.006)	(0.005)	(0.007)	(0.005)	(0.006)	(0.005)
Trust: Modern Medicine	-0.002	-0.005	-0.009	-0.006	0.011	0.011**
	(0.008)	(0.007)	(0.008)	(0.006)	(0.007)	(0.006)
Trust: Doctors	-0.001	-0.003	-0.007	-0.005	0.009	0.008
	(0.008)	(0.007)	(0.008)	(0.007)	(0.007)	(0.007)
Trust: Public Health Authorities	0.003	0.005	-0.006	-0.009*	0.003	0.004
	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)
Trust in Vaccines	0.007	-0.001	-0.014	-0.006	0.007	0.007
	(0.012)	(0.011)	(0.014)	(0.011)	(0.013)	(0.010)
Trust in Healthcare System	0.001	-0.001	-0.016	-0.016	0.015	0.016
	(0.012)	(0.011)	(0.012)	(0.010)	(0.012)	(0.011)
Accept a Vaccine: Never	-0.018	-0.025	0.009	-0.004	0.008	0.029
	(0.037)	(0.032)	(0.033)	(0.029)	(0.035)	(0.029)
Accept a Vaccine: Depends on Vaccine	0.021	0.027	-0.050	-0.026	0.029	-0.001
	(0.031)	(0.028)	(0.033)	(0.028)	(0.030)	(0.028)
Accept a Vaccine: If Doctor Suggests	0.010	0.004	0.051	0.031	-0.060*	-0.035
	(0.038)	(0.034)	(0.039)	(0.032)	(0.033)	(0.032)
Information Source: Newspapers	-0.045	-0.061	0.035	0.037	0.011	0.023
	(0.053)	(0.050)	(0.052)	(0.042)	(0.049)	(0.041)
Information Source: Television	0.001	0.023	-0.056	-0.076	0.056	0.053
	(0.067)	(0.067)	(0.060)	(0.053)	(0.070)	(0.057)
Information Source: Social Media	-0.013	0.003	0.009	0.008	0.004	-0.011
	(0.040)	(0.038)	(0.038)	(0.032)	(0.040)	(0.034)
Information Source: Don't Watch	0.162**	0.112	-0.055	-0.038	-0.106*	-0.074
	(0.076)	(0.072)	(0.065)	(0.063)	(0.062)	(0.063)
Session Day		✓		✓		✓
Session Order		✓		✓		✓
Prefecture FE		✓		✓		✓

Notes: The table reports balance checks across treatment groups for baseline beliefs, trust, social-media use, and vaccine-related attitudes measured before treatment exposure. Each coefficient is obtained from a separate regression of the treatment indicator on the corresponding baseline characteristic. Odd-numbered columns report specifications without fixed effects, while even-numbered columns include prefecture, session-day, and session-order fixed effects. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A8: Balancing Exercise #3

	Treatment vs. Control		VR vs. Video		Scientist vs. Influencer	
	(1)	(2)	(3)	(4)	(5)	(6)
Age (years)	-0.009**	-0.006*	-0.004	-0.003	-0.001	-0.002
	(0.004)	(0.004)	(0.005)	(0.003)	(0.005)	(0.005)
Friends Reported	-0.002	-0.002	-0.008	-0.005	-0.023***	-0.020***
	(0.008)	(0.006)	(0.009)	(0.005)	(0.006)	(0.007)
Friends Reporting Respondent	0.016	0.017	0.019	0.021**	-0.012	-0.010
	(0.012)	(0.010)	(0.016)	(0.011)	(0.011)	(0.013)
Smoking: More Than Once a Week	0.038	0.046*	-0.047	-0.001	0.030	0.019
	(0.032)	(0.027)	(0.041)	(0.034)	(0.034)	(0.038)
Exercising: More Than Once a Week	0.113**	0.090*	-0.004	0.013	0.072	0.082
	(0.049)	(0.048)	(0.065)	(0.042)	(0.083)	(0.075)
Trust in Social Media	-0.013	-0.009	0.013	0.011	0.029**	0.027**
	(0.009)	(0.008)	(0.012)	(0.007)	(0.012)	(0.012)
Female	-0.031	-0.032	-0.028	-0.014	0.010	0.007
	(0.030)	(0.025)	(0.038)	(0.028)	(0.038)	(0.038)
Income: Above Median	0.038	0.040	-0.045	-0.014	-0.021	-0.042
	(0.032)	(0.032)	(0.042)	(0.030)	(0.048)	(0.051)
Political Ideology: Left/Center-Left	0.002	-0.002	-0.001	-0.018	0.012	0.006
	(0.034)	(0.032)	(0.043)	(0.029)	(0.043)	(0.044)
Political Ideology: Center	0.010	0.017	-0.071	-0.027	-0.008	0.027
	(0.034)	(0.033)	(0.049)	(0.029)	(0.043)	(0.044)
Political Ideology: Right/Center-Right	0.033	0.032	0.042	0.033	0.029	0.024
	(0.036)	(0.032)	(0.043)	(0.026)	(0.039)	(0.042)
Enrolled in: Undergraduate Degree	0.096	0.118*	-0.102	-0.063	-0.051	-0.046
	(0.078)	(0.069)	(0.100)	(0.074)	(0.083)	(0.094)
Enrolled in: Graduate Degree	-0.090	-0.112	0.117	0.043	0.066	0.072
	(0.087)	(0.075)	(0.110)	(0.089)	(0.087)	(0.103)
Father's Education: Above BSc	-0.010	-0.016	-0.003	-0.008	0.000	0.003
	(0.027)	(0.024)	(0.033)	(0.020)	(0.037)	(0.038)
Mother's Education: Above BSc	0.005	-0.001	0.008	0.022	0.034	0.052
	(0.024)	(0.023)	(0.032)	(0.022)	(0.041)	(0.045)
Use Social Media: Friends	-0.014	-0.024	-0.016	0.010	0.028	0.044
	(0.027)	(0.025)	(0.040)	(0.028)	(0.039)	(0.041)
Use Social Media: News	-0.041	-0.012	-0.049	-0.033	-0.033	-0.032
	(0.051)	(0.045)	(0.069)	(0.043)	(0.068)	(0.074)
Use Social Media: Entertainment	0.027	0.026	0.031	0.005	-0.025	-0.036
	(0.026)	(0.023)	(0.034)	(0.025)	(0.040)	(0.043)
Use Social Media: Professional	-0.005	0.031	-0.012	-0.135	0.016	-0.046
	(0.107)	(0.095)	(0.145)	(0.120)	(0.130)	(0.148)
Use Social Media: No Use	0.012	0.052	0.026	0.063	-0.023	-0.065
	(0.188)	(0.184)	(0.246)	(0.122)	(0.253)	(0.293)
More than 2h a Day on Social Media	-0.023	-0.031	-0.010	-0.004	0.029	0.017
	(0.030)	(0.026)	(0.039)	(0.027)	(0.039)	(0.044)
Session Day		✓		✓		✓
Session Order		✓		✓		✓
Prefecture FE		✓		✓		✓

Notes: The table reports balancing tests across experimental treatment groups for baseline demographic and socioeconomic characteristics. Each coefficient is obtained from a separate regression of the treatment indicator on the corresponding baseline characteristic. Columns (1)–(2) compare participants assigned to any treatment condition with those assigned to the control group. Columns (3)–(4) compare participants assigned to the VR treatment with those assigned to the standard video treatment. Columns (5)–(6) compare participants exposed to a scientist with those exposed to a social media influencer. Odd-numbered columns report specifications without fixed effects, while even-numbered columns include prefecture, session-day, and session-order fixed effects. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A9: Balancing Exercise #4

	Treatment vs. Control		VR vs. Video		Scientist vs. Influencer	
	(1)	(2)	(3)	(4)	(5)	(6)
Know Female Influencer	0.006	0.005	-0.018	-0.027**	-0.003	-0.002
	(0.014)	(0.011)	(0.015)	(0.011)	(0.016)	(0.016)
Know Male Influencer	-0.005	-0.014	0.035*	0.029**	-0.008	-0.004
	(0.016)	(0.014)	(0.020)	(0.013)	(0.018)	(0.021)
Trust Female Influencer	-0.009	-0.019	-0.028	-0.016	0.002	0.007
	(0.015)	(0.014)	(0.019)	(0.013)	(0.017)	(0.017)
Trust Male Influencer	0.000	-0.010	-0.005	-0.012	-0.003	-0.001
	(0.014)	(0.013)	(0.021)	(0.011)	(0.017)	(0.019)
Response Time (seconds)	-0.001	-0.001	0.001	0.001	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Vaccines: Effective	-0.003	-0.001	0.007	0.008	-0.007	-0.007
	(0.006)	(0.005)	(0.009)	(0.006)	(0.008)	(0.009)
Vaccines: Safe	-0.004	-0.002	0.007	0.007	0.002	0.006
	(0.006)	(0.005)	(0.009)	(0.006)	(0.008)	(0.009)
Vaccines: Important for Own Health	-0.004	0.001	0.008	0.004	0.006	0.009
	(0.006)	(0.005)	(0.008)	(0.005)	(0.007)	(0.008)
Vaccines: Important for Public Health	-0.002	0.003	0.004	0.005	0.000	0.004
	(0.006)	(0.005)	(0.009)	(0.006)	(0.008)	(0.009)
Trust: Modern Medicine	0.002	0.005	0.016*	0.015**	-0.009	-0.007
	(0.008)	(0.007)	(0.010)	(0.006)	(0.010)	(0.011)
Trust: Doctors	0.001	0.003	0.013	0.007	-0.014	-0.012
	(0.008)	(0.007)	(0.010)	(0.008)	(0.012)	(0.013)
Trust: Public Health Authorities	-0.003	-0.005	0.007	0.005	-0.005	-0.004
	(0.006)	(0.005)	(0.008)	(0.006)	(0.009)	(0.010)
Trust in Vaccines	-0.007	0.001	0.015	0.014	0.000	0.007
	(0.012)	(0.011)	(0.018)	(0.012)	(0.017)	(0.018)
Trust in Healthcare System	-0.001	0.001	0.025	0.019	-0.020	-0.018
	(0.012)	(0.011)	(0.016)	(0.012)	(0.019)	(0.021)
Refuse a Vaccine: Never	0.018	0.025	0.000	0.000	0.078	0.095*
	(0.037)	(0.032)	(0.043)	(0.032)	(0.051)	(0.053)
Refuse a Vaccine: Maybe	-0.021	-0.027	0.059	0.016	-0.081**	-0.094**
	(0.031)	(0.028)	(0.042)	(0.032)	(0.038)	(0.040)
Refuse a Vaccine: If Doctor Suggests	-0.010	-0.004	-0.086*	-0.022	0.019	0.020
	(0.038)	(0.034)	(0.047)	(0.038)	(0.044)	(0.049)
Information Source: Newspapers	0.045	0.061	-0.016	-0.002	-0.019	0.010
	(0.053)	(0.050)	(0.062)	(0.039)	(0.059)	(0.067)
Information Source: Television	-0.001	-0.023	0.086	0.096**	-0.140	-0.119
	(0.067)	(0.067)	(0.084)	(0.048)	(0.086)	(0.101)
Information Source: Social Media	0.013	-0.003	-0.003	0.000	0.048	0.021
	(0.040)	(0.038)	(0.052)	(0.032)	(0.042)	(0.045)
Information Source: Don't Watch	-0.162**	-0.112	-0.059	-0.134	0.020	0.032
	(0.076)	(0.072)	(0.099)	(0.093)	(0.104)	(0.118)
Session Day		✓		✓		✓
Session Order		✓		✓		✓
Prefecture FE		✓		✓		✓

Notes: The table reports balancing tests across experimental treatment groups for baseline beliefs, trust, social-media use, and vaccine-related attitudes measured before treatment exposure. Each coefficient is obtained from a separate regression of the treatment indicator on the corresponding baseline characteristic. Columns (1)–(2) compare participants assigned to any treatment condition with those assigned to the control group. Columns (3)–(4) compare participants assigned to the VR treatment with those assigned to the standard video treatment. Columns (5)–(6) compare participants exposed to a scientist with those exposed to a social media influencer. Odd-numbered columns report specifications without fixed effects, while even-numbered columns include prefecture, session-day, and session-order fixed effects. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A10: Treatment Effect on Trust in Vaccines and Healthcare System: Treatment vs. Control

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Treatment Effect on Beliefs on Vaccines							
	Belief that Vaccinations Are				Belief that Vaccinations Are Important for			
	Safe		Effective		Personal Health		Public Health	
Treatment	0.636***	0.640***	0.538***	0.524***	0.504***	0.501***	0.470***	0.465***
	(0.074)	(0.071)	(0.070)	(0.067)	(0.070)	(0.067)	(0.065)	(0.063)
Reference Group Mean	7.60	7.60	7.31	7.31	7.60	7.60	8.04	8.04
Reference Group SD	1.95	1.95	1.96	1.96	2.13	2.13	2.01	2.01
Observations	1,218	1,218	1,218	1,218	1,218	1,218	1,218	1,218
	Panel B: Treatment Effect on Beliefs on the Healthcare System							
	Trust in						Placebo	
	Modern Medicine		Medical Doctors		Health Authorities		Response Time	
Treatment	0.177***	0.189***	0.312***	0.298***	0.583***	0.584***	0.086	0.093
	(0.057)	(0.058)	(0.065)	(0.061)	(0.078)	(0.076)	(0.065)	(0.075)
Reference Group Mean	7.84	7.84	7.69	7.69	6.10	6.10	-0.04	-0.04
Reference Group SD	1.84	1.84	1.73	1.73	2.36	2.36	0.52	0.52
Observations	1,218	1,218	1,218	1,218	1,218	1,218	1,218	1,218
Session Day	✓	✓	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓

Notes: The table reports the effect of exposure to the information treatment (video or VR) relative to the control condition on the individual survey items underlying the trust-in-vaccines and healthcare-system trust indices. Panel A reports outcomes measuring beliefs that vaccines are safe and effective and beliefs that vaccines are important for personal and public health. Panel B reports trust in modern medicine, medical doctors, and health authorities, as well as the standardized time required to complete the corresponding survey section as a placebo outcome. The omitted category is the control group. Odd-numbered columns control for the baseline value of the corresponding item, as well as prefecture, session-day, and session-order fixed effects. Even-numbered columns additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A11: Treatment Effect on Trust in Vaccines and Healthcare System: Video vs. VR vs. Control

	Trust in			
	Vaccines		Healthcare System	
	(1)	(2)	(3)	(4)
Treatment: Video	0.204***	0.209***	0.171***	0.175***
	(0.029)	(0.027)	(0.036)	(0.035)
Treatment: VR	0.332***	0.323***	0.290***	0.286***
	(0.033)	(0.031)	(0.036)	(0.035)
Reference Group Mean	0.08	0.08	0.04	0.04
Reference Group SD	0.94	0.94	1.09	1.09
Observations	1,218	1,218	1,218	1,218
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls		✓		✓

Notes: The table reports treatment effects on trust in vaccines and trust in the healthcare system, separating the information treatment by communication medium. The table reports coefficients for the standard video and VR treatment arms relative to the control condition. Columns (1) and (3) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2) and (4) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A12: Treatment Effect on Trust in Vaccines and Healthcare System: VR vs. Video

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Treatment Effect on Beliefs on Vaccines							
	Belief that Vaccinations Are				Belief that Vaccinations Are Important for			
	Safe		Effective		Personal Health		Public Health	
Treatment: VR	0.214*	0.158	0.362***	0.331***	0.262**	0.220**	0.178*	0.143
	(0.121)	(0.119)	(0.122)	(0.121)	(0.107)	(0.108)	(0.101)	(0.103)
Reference Group Mean	8.03	8.03	7.64	7.64	7.91	7.91	8.42	8.42
Reference Group SD	1.85	1.85	1.96	1.96	2.07	2.07	1.94	1.94
Observations	798	798	798	798	798	798	798	798
	Panel B: Treatment Effect on Beliefs on the Healthcare System							
	Trust in						Placebo	
	Modern Medicine		Medical Doctors		Health Authorities		Response Time	
Treatment: VR	0.231**	0.232**	0.196*	0.154	0.146	0.121	0.199	0.189
	(0.104)	(0.102)	(0.103)	(0.101)	(0.135)	(0.139)	(0.126)	(0.140)
Reference Group Mean	7.92	7.92	7.87	7.87	6.47	6.47	-0.06	-0.06
Reference Group SD	1.77	1.77	1.72	1.72	2.15	2.15	0.50	0.50
Observations	798	798	798	798	798	798	798	798
Session Day	✓	✓	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓

Notes: The table reports the effect of VR delivery relative to standard video delivery among treated participants on the individual survey items underlying the trust-in-vaccines and healthcare-system trust indices. Panel A reports outcomes measuring beliefs that vaccines are safe and effective and beliefs that vaccines are important for personal and public health. Panel B reports trust in modern medicine, medical doctors, and health authorities, as well as the standardized time required to complete the corresponding survey section as a placebo outcome. The omitted category is the standard video treatment. Odd-numbered columns control for the baseline value of the corresponding item, as well as prefecture, session-day, and session-order fixed effects. Even-numbered columns additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A13: Treatment Effect on Trust in Vaccines and Healthcare System: Scientist vs. Influencer

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Treatment Effect on Beliefs on Vaccines							
	Belief that Vaccinations Are				Belief that Vaccinations Are Important for			
	Safe		Effective		Personal Health		Public Health	
Treatment: Scientist	0.087	0.079	0.036	0.027	0.094	0.118	0.112	0.107
	(0.073)	(0.074)	(0.091)	(0.088)	(0.095)	(0.096)	(0.087)	(0.088)
Reference Group Mean	8.17	8.17	7.77	7.77	7.95	7.95	8.46	8.46
Reference Group SD	1.78	1.78	1.90	1.90	2.07	2.07	1.81	1.81
Observations	798	798	798	798	798	798	798	798
	Panel B: Treatment Effect on Beliefs on the Healthcare System							
	Trust in						Placebo	
	Modern Medicine		Medical Doctors		Health Authorities		Response Time	
Treatment: Scientist	-0.093	-0.112	0.067	0.055	0.045	0.036	0.034	0.049
	(0.074)	(0.076)	(0.078)	(0.080)	(0.103)	(0.104)	(0.076)	(0.091)
Reference Group Mean	8.13	8.13	8.02	8.02	6.62	6.62	0.00	0.00
Reference Group SD	1.58	1.58	1.53	1.53	2.04	2.04	0.57	0.57
Observations	798	798	798	798	798	798	798	798
Session Day	✓	✓	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓

Notes: The table shows the effect of exposure to a scientist speaker relative to an influencer speaker among treated participants on the individual survey items underlying the trust-in-vaccines and healthcare-system trust indices. Panel A reports outcomes measuring beliefs that vaccines are safe and effective and beliefs that vaccines are important for personal and public health. Panel B reports trust in modern medicine, medical doctors, and health authorities, as well as the standardized time required to complete the corresponding survey section as a placebo outcome. The omitted category is the influencer treatment arm. Odd-numbered columns control for the baseline value of the corresponding item, as well as prefecture, session-day, and session-order fixed effects. Even-numbered columns additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A14: Treatment Effect on Trust in Vaccines and Healthcare System: Scientist vs. Influencer vs. Control

	Trust in			
	Vaccines		Healthcare System	
	(1)	(2)	(3)	(4)
Treatment: Influencer	0.243***	0.242***	0.218***	0.221***
	(0.032)	(0.031)	(0.037)	(0.037)
Treatment: Scientist	0.289***	0.288***	0.241***	0.237***
	(0.033)	(0.030)	(0.033)	(0.034)
Reference Group Mean	0.08	0.08	0.04	0.04
Reference Group SD	0.94	0.94	1.09	1.09
Observations	1,218	1,218	1,218	1,218
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls		✓		✓

Notes: The table reports treatment effects on trust in vaccines and trust in the healthcare system, separating the information treatment by speaker type. The table reports coefficients for scientist and influencer speakers relative to the control condition. Columns (1) and (3) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2) and (4) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A15: Treatment Effect on Trust in Vaccines and Healthcare System: Treatment vs. Control, Breakdown by Speaker Identity

	Trust in			
	Vaccines		Healthcare System	
	(1)	(2)	(3)	(4)
Treatment: Influencer Female	0.196***	0.190***	0.214***	0.207***
	(0.039)	(0.040)	(0.043)	(0.043)
Treatment: Influencer Male	0.294***	0.300***	0.222***	0.237***
	(0.041)	(0.040)	(0.050)	(0.050)
Treatment: Scientist Female	0.285***	0.287***	0.270***	0.266***
	(0.040)	(0.037)	(0.046)	(0.046)
Treatment: Scientist Male	0.293***	0.288***	0.215***	0.212***
	(0.040)	(0.038)	(0.042)	(0.044)
Reference Group Mean	0.08	0.08	0.04	0.04
Reference Group SD	0.94	0.94	1.09	1.09
Observations	1,218	1,218	1,218	1,218
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls		✓		✓

Notes: The table reports treatment effects on trust in vaccines and trust in the healthcare system, separating the information treatment into the four speaker-identity arms: female scientist, male scientist, female influencer, and male influencer. The omitted category is the control condition. Columns (1) and (3) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2) and (4) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A16: Treatment Effect on Trust in Vaccines and Healthcare System—Common Factor Analysis

	Trust in			
	Vaccines		Healthcare System	
	(1)	(2)	(3)	(4)
	Panel A: Treatment vs. Control			
Treatment: Information	0.295***	0.292***	0.205***	0.204***
	(0.030)	(0.028)	(0.028)	(0.028)
Reference Group Mean	-0.18	-0.18	-0.14	-0.14
Reference Group SD	1.03	1.03	1.04	1.04
Observations	1,218	1,218	1,218	1,218
	Panel B: VR vs. Video			
Treatment: VR	0.145***	0.123***	0.127**	0.112**
	(0.042)	(0.041)	(0.054)	(0.049)
Reference Group Mean	0.02	0.02	-0.02	-0.02
Reference Group SD	1.03	1.03	1.03	1.03
Observations	798	798	798	798
	Panel C: Scientist vs. Influencer			
Treatment: Scientist	0.040	0.044	0.006	-0.003
	(0.041)	(0.040)	(0.041)	(0.044)
Reference Group Mean	0.06	0.06	0.09	0.09
Reference Group SD	0.99	0.99	0.91	0.91
Observations	798	798	798	798
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls		✓		✓

Notes: The table re-estimates the main treatment effects using common-factor measures of trust in vaccines and trust in the healthcare system instead of the principal-component indices used in the main tables. Panel A compares exposure to the information treatment (video or VR) vs. the control condition, Panel B compares VR vs. standard video among treated participants, and Panel C compares scientist vs. influencer speakers among treated participants. The omitted category varies accordingly across panels. Columns (1) and (3) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2) and (4) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A17: Treatment Effect on Trust in Vaccines and Healthcare System—Bootstrapped Standard Errors

	Trust in Vaccines		Trust in Healthcare System	
	(1)	(2)	(3)	(4)
Treatment: Treatment vs Control	0.267***	0.265***	0.230***	0.229***
	(0.032)	(0.029)	(0.034)	(0.034)
N	1,218	1,218	1,218	1,218
Treatment: VR vs Video	0.127**	0.105*	0.120	0.105
	(0.051)	(0.055)	(0.076)	(0.075)
N	798	798	798	798
Treatment: Scientist vs Influencer	0.037	0.041	0.012	0.003
	(0.037)	(0.035)	(0.043)	(0.046)
N	798	798	798	798
Treatment: Female Scientist vs Female Influencer	0.096*	0.116**	0.034	0.048
	(0.050)	(0.051)	(0.062)	(0.067)
N	798	798	798	798
Treatment: Male Influencer vs Female Influencer	0.137***	0.153***	0.017	0.046
	(0.048)	(0.052)	(0.060)	(0.058)
N	798	798	798	798
Treatment: Male Scientist vs Female Influencer	0.114**	0.122**	0.006	0.007
	(0.050)	(0.051)	(0.051)	(0.057)
N	798	798	798	798
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls		✓		✓

Notes: The table re-estimates the main treatment effects on trust in vaccines and trust in the healthcare system using bootstrapped standard errors. The reported comparisons are exposure to the information treatment (video or VR) vs. the control condition, VR vs. standard video among treated participants, and scientist vs. influencer speakers among treated participants. Outcomes are principal-component indices constructed from the underlying survey items. To account for the sampling uncertainty introduced by index construction, standard errors are obtained using a session-clustered bootstrap procedure that re-estimates the principal components within each bootstrap replication. Columns (1) and (3) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2) and (4) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Bootstrapped standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A18: Effect of Treatment Assignment on Follow-Up Participation

	Follow-up Participation	
	(1)	(2)
Panel A: First Follow-up (2 Months)		
Treatment: Video	-0.028	-0.018
	(0.035)	(0.035)
Treatment: VR	-0.031	-0.034
	(0.038)	(0.037)
Reference Group Mean	0.63	0.63
Reference Group SD	0.48	0.48
Observations	1,218	1,218
Panel B: Second Follow-up (6 Months)		
Treatment: Video	-0.025	-0.018
	(0.028)	(0.029)
Treatment: VR	-0.036	-0.033
	(0.031)	(0.031)
Reference Group Mean	0.26	0.26
Reference Group SD	0.44	0.44
Observations	1,218	1,218
Session Day	✓	✓
Session Order	✓	✓
Prefecture FE	✓	✓
Controls		✓

Notes: The table reports whether assignment to the information treatment predicts participation in the two follow-up surveys. The first follow-up was administered approximately two months after treatment for participants in both experimental waves. The second follow-up was administered approximately six months after treatment for participants in the second wave and eight to nine months after treatment for participants in the first wave. The omitted category is the control group. Column (1) includes prefecture, session-day, and session-order fixed effects. Column (2) additionally controls for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A19: Persistence of the Information Treatment by Propensity to Vaccinate, First Follow-up (2 Months)

	Trust in			
	Vaccines (1)	Healthcare System (2)	Willing to Vaccinate (3)	Did Vaccinate (4)
	Panel A: Participants with a Low Propensity to Vaccinate			
Treatment: Information	0.344***	0.091	0.092*	-0.018
	(0.084)	(0.083)	(0.051)	(0.018)
	[0.296, 0.388]	[0.053, 0.140]	[0.093, 0.126]	[-0.023, 0.015]
Reference Group Mean	-0.28	-0.08	0.32	0.03
Reference Group SD	1.02	0.88	0.47	0.17
Observations	457	457	457	457
	Panel B: Participants with a High Propensity to Vaccinate			
Treatment: Information	0.134**	0.008	-0.040	-0.038
	(0.059)	(0.085)	(0.065)	(0.040)
	[0.138, 0.164]	[-0.007, 0.022]	[-0.030, -0.021]	[-0.040, -0.018]
Reference Group Mean	0.40	0.48	0.45	0.07
Reference Group SD	0.72	0.87	0.50	0.26
Observations	313	313	313	313
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓

Notes: The table reports the effect of exposure to the information treatment (video or VR) on outcomes measured in the first follow-up survey, together with Lee bounds correcting for differential attrition (D. S. Lee 2009). The follow-up was administered approximately two months after treatment for participants in both experimental waves. Panel A restricts the sample to participants with a low pre-treatment propensity to vaccinate, defined as those who report that they would receive a vaccine recommended by a public health authority "depending on the vaccine" or "would never" receive such a vaccine. Panel B restricts the sample to participants with a high pre-treatment propensity to vaccinate, defined as those who report that they would "always" receive a recommended vaccine or would receive it "only if recommended by their doctor." The outcomes are trust in vaccines (Column 1), trust in the healthcare system (Column 2), whether the participant reports being more willing to vaccinate (Column 3), and whether the participant received a vaccination between the baseline and follow-up survey (Column 4). The omitted category is the control group. All specifications control for prefecture, session-day, session-order fixed effects, and the full set of participant characteristics described in Section 5. For trust in vaccines and trust in the healthcare system, we control for the corresponding pre-treatment measure. For willingness to vaccinate and vaccination uptake, we control for the pre-treatment measures of both trust in vaccines and trust in the healthcare system. Standard errors, clustered at the session level, are reported in parentheses. Lee bounds are reported in square brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A20: Persistence of the Treatment Effects, First Follow-up (2 Months)

	Trust in			
	Vaccines (1)	Healthcare System (2)	Willing to Vaccinate (3)	Did Vaccinate (4)
	Panel A: VR vs. Video			
Treatment: VR	-0.083	-0.019	-0.062	0.052***
	(0.067)	(0.070)	(0.061)	(0.017)
	[-0.157, 0.029]	[-0.133, 0.085]	[-0.161, 0.001]	[-0.015, 0.057]
Reference Group Mean	0.21	0.10	0.45	0.03
Reference Group SD	0.85	1.01	0.50	0.18
Observations	500	500	500	500
	Panel B: Scientist vs. Influencer			
Treatment: Scientist	0.003	0.047	-0.017	0.012
	(0.053)	(0.077)	(0.048)	(0.020)
	[-0.049, 0.048]	[-0.028, 0.129]	[-0.054, 0.011]	[0.012, 0.056]
Reference Group Mean	0.19	0.15	0.44	0.03
Reference Group SD	0.83	0.89	0.50	0.18
Observations	500	500	500	500
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓

Notes: The table reports the effects of the communication medium and messenger type on outcomes measured in the first follow-up survey, together with Lee bounds correcting for differential attrition (D. S. Lee 2009). The follow-up was administered approximately two months after treatment for participants in both experimental waves. Panel A estimates the effect of VR delivery relative to standard video delivery among participants assigned to the information treatment, with video as the omitted category. Panel B estimates the effect of a scientist messenger relative to a social media influencer among participants assigned to the information treatment, with the influencer as the omitted category. The outcomes are trust in vaccines (Column 1), trust in the healthcare system (Column 2), whether the participant reports being more willing to vaccinate (Column 3), and whether the participant received a vaccination between the baseline and follow-up survey (Column 4). All specifications control for prefecture, session-day, session-order fixed effects, and the full set of participant characteristics described in Section 5. For trust in vaccines and trust in the healthcare system, we control for the corresponding pre-treatment measure. For willingness to vaccinate and vaccination uptake, we control for the pre-treatment measures of both trust in vaccines and trust in the healthcare system. Standard errors, clustered at the session level, are reported in parentheses. Lee bounds are reported in square brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A21: Multiple-Hypothesis Testing: Information Treatment vs. Control

Heterogeneity Dimension	Raw p-value	FDR-adjusted q-value
Trust in Vaccines		
Gender	0.022**	0.255
Parental Education	0.760	0.972
Household Income	0.846	0.972
Major in Economics	0.874	0.972
Political Ideology	0.001***	0.050*
Friends Reported by Respondent	0.857	0.972
Friends Reporting Respondent	0.055*	0.360
Time Spent on Social Media	0.439	0.815
Social Media Used for Entertainment	0.594	0.927
Social Media as Main News Source	0.072*	0.401
Trust in Social Media Information	0.172	0.585
Baseline Trust in Vaccines	<0.001***	<0.001***
Baseline Propensity to Vaccinate	0.042**	0.297
Trust in Healthcare System		
Gender	0.003***	0.076*
Parental Education	0.065*	0.388
Household Income	0.799	0.972
Major in Economics	0.616	0.927
Political Ideology	0.161	0.585
Friends Reported by Respondent	0.034**	0.294
Friends Reporting Respondent	0.307	0.775
Time Spent on Social Media	0.787	0.972
Social Media Used for Entertainment	0.089*	0.465
Social Media as Main News Source	0.775	0.972
Trust in Social Media Information	0.756	0.972
Baseline Trust in Healthcare System	0.004***	0.076*
Baseline Propensity to Vaccinate	0.833	0.972

Notes: The table reports omnibus tests of whether the effect of receiving information, relative to the control condition, is constant across the categories of each heterogeneity dimension. Raw p-values are obtained from joint Wald tests of the treatment-by-characteristic interaction terms. The adjusted q-values apply the Benjamini–Hochberg false-discovery-rate correction across all 78 heterogeneity tests reported in the three multiple-testing tables. *** < 0.01, ** < 0.05, * < 0.10.

Table A22: Multiple-Hypothesis Testing: VR vs. Video

Heterogeneity Dimension	Raw p-value	FDR-adjusted q-value
Trust in Vaccines		
Gender	0.331	0.775
Parental Education	0.301	0.775
Household Income	0.704	0.972
Major in Economics	0.168	0.585
Political Ideology	0.901	0.976
Friends Reported by Respondent	0.820	0.972
Friends Reporting Respondent	0.152	0.585
Time Spent on Social Media	0.217	0.662
Social Media Used for Entertainment	0.570	0.927
Social Media as Main News Source	0.029**	0.279
Trust in Social Media Information	0.638	0.939
Baseline Trust in Vaccines	0.021**	0.255
Baseline Propensity to Vaccinate	0.618	0.927
Trust in Healthcare System		
Gender	0.961	0.996
Parental Education	0.487	0.883
Household Income	0.588	0.927
Major in Economics	0.372	0.806
Political Ideology	0.221	0.662
Friends Reported by Respondent	0.884	0.972
Friends Reporting Respondent	0.845	0.972
Time Spent on Social Media	0.168	0.585
Social Media Used for Entertainment	0.869	0.972
Social Media as Main News Source	0.401	0.810
Trust in Social Media Information	0.023**	0.255
Baseline Trust in Healthcare System	0.941	0.996
Baseline Propensity to Vaccinate	0.570	0.927

Notes: The table reports omnibus tests of whether the effect of receiving information via virtual reality, relative to receiving it via 2D video, is constant across the categories of each heterogeneity dimension. Raw p-values are obtained from joint Wald tests of the treatment-by-characteristic interaction terms. The adjusted q-values apply the Benjamini–Hochberg false-discovery-rate correction across all 78 heterogeneity tests reported in the three multiple-testing tables. *** < 0.01, ** < 0.05, * < 0.10.

Table A23: Multiple-Hypothesis Testing: Scientist vs. Influencer

Heterogeneity Dimension	Raw p-value	FDR-adjusted q-value
Trust in Vaccines		
Gender	0.782	0.972
Parental Education	0.212	0.662
Household Income	0.528	0.922
Major in Economics	0.996	0.996
Political Ideology	0.433	0.815
Friends Reported by Respondent	0.299	0.775
Friends Reporting Respondent	0.348	0.775
Time Spent on Social Media	0.985	0.996
Social Media Used for Entertainment	0.122	0.585
Social Media as Main News Source	0.335	0.775
Trust in Social Media Information	0.405	0.810
Baseline Trust in Vaccines	0.710	0.972
Baseline Propensity to Vaccinate	0.532	0.922
Trust in Healthcare System		
Gender	0.598	0.927
Parental Education	0.422	0.815
Household Income	0.395	0.810
Major in Economics	0.961	0.996
Political Ideology	0.688	0.972
Friends Reported by Respondent	0.260	0.751
Friends Reporting Respondent	0.132	0.585
Time Spent on Social Media	0.829	0.972
Social Media Used for Entertainment	0.038**	0.297
Social Media as Main News Source	0.335	0.775
Trust in Social Media Information	0.138	0.585
Baseline Trust in Healthcare System	0.342	0.775
Baseline Propensity to Vaccinate	0.988	0.996

Notes: The table reports omnibus tests of whether the effect of receiving information from the scientist, relative to the influencer speaker, is constant across the categories of each heterogeneity dimension. Raw p-values are obtained from joint Wald tests of the treatment-by-characteristic interaction terms. The adjusted q-values apply the Benjamini–Hochberg false-discovery-rate correction across all 78 heterogeneity tests reported in the three multiple-testing tables. *** < 0.01, ** < 0.05, * < 0.10.

Table A24: Treatment Effects by Speaker Type and Gender Match

	Trust in			
	Vaccines		Healthcare System	
	(1)	(2)	(3)	(4)
Same-Gender Speaker: Influencers	-0.067	-0.093*	-0.039	-0.058
	(0.050)	(0.048)	(0.068)	(0.064)
Same-Gender Speaker: Scientists	-0.109**	-0.086*	-0.056	-0.006
	(0.045)	(0.047)	(0.067)	(0.069)
Difference: Scientists − Influencers	-0.042	0.007	-0.017	0.051
	(0.065)	(0.067)	(0.099)	(0.098)
Reference Group Mean	0.37	0.37	0.24	0.24
Reference Group SD	0.86	0.86	0.98	0.98
Observations	795	795	795	795
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls		✓		✓

Notes: The table reports estimates of the effect of being assigned to a speaker of the same gender relative to a speaker of the opposite gender, separately for influencers and scientists, using equation (2). The omitted category in the underlying regression is participants assigned to an opposite-gender influencer. The row Same-Gender Speaker: influencers reports the within-influencer contrast between same- and opposite-gender speaker assignments. The row Same-Gender Speaker: scientistsreports the corresponding within-scientist contrast. The row Difference: Scientists − Influencers reports the difference between the two homophily effects. Columns (1) and (3) control for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. Columns (2) and (4) additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A25: Effects on Willingness to Discuss Vaccines, Healthcare, and Other Topics

	Vaccines		Healthcare Authorities		Social Media		Artificial Intelligence		Politics	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Information Treatment vs. Control										
Treatment: Information	0.434***	0.425***	0.387***	0.385***	0.039	0.041	0.215***	0.226***	0.212***	0.171***
	(0.063)	(0.062)	(0.065)	(0.062)	(0.062)	(0.065)	(0.064)	(0.067)	(0.060)	(0.052)
Reference Group Mean	-0.28	-0.28	-0.25	-0.25	-0.06	-0.06	-0.17	-0.17	-0.13	-0.13
Reference Group SD	1.04	1.04	1.06	1.06	1.04	1.04	1.10	1.10	1.02	1.02
Observations	1,218	1,218	1,218	1,218	1,218	1,218	1,218	1,218	1,218	1,218
Panel B: VR vs. Video										
Treatment: VR	0.064	0.046	0.100	0.080	0.186**	0.168*	0.227***	0.268***	-0.013	-0.018
	(0.085)	(0.086)	(0.072)	(0.077)	(0.082)	(0.086)	(0.073)	(0.072)	(0.068)	(0.070)
Reference Group Mean	0.11	0.11	0.08	0.08	-0.03	-0.03	-0.01	-0.01	0.04	0.04
Reference Group SD	0.95	0.95	0.96	0.96	1.01	1.01	0.97	0.97	1.00	1.00
Observations	798	798	798	798	798	798	798	798	798	798
Panel C: Scientist vs. Influencer										
Treatment: Scientist	0.074	0.050	-0.082	-0.100	-0.119	-0.101	-0.093	-0.072	0.005	-0.023
	(0.077)	(0.076)	(0.078)	(0.078)	(0.088)	(0.087)	(0.076)	(0.074)	(0.075)	(0.070)
Reference Group Mean	0.11	0.11	0.17	0.17	0.08	0.08	0.13	0.13	0.07	0.07
Reference Group SD	0.93	0.93	0.90	0.90	0.95	0.95	0.91	0.91	0.97	0.97
Observations	798	798	798	798	798	798	798	798	798	798
Session Day	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓		✓

Notes: The table reports treatment effects on willingness to discuss different topics (standardized). Panel A compares exposure to the information treatment (video or VR) vs. the control condition, Panel B compares VR vs. standard video among treated participants, and Panel C compares scientist vs. influencer speakers among treated participants. The omitted category varies accordingly across panels. All regressions include prefecture, session-day, and session-order fixed effects, as well as the pre-treatment (baseline) measures of trust in vaccines and trust in the healthcare system. Controls include the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A26: Willingness to Discuss: Breakdown by Speaker Identity

	Vaccines		Healthcare Authorities		Social Media		Artificial Intelligence		Politics	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment: Influencer – Male	0.002	0.036	-0.008	0.010	0.015	0.082	0.072	0.124	-0.029	-0.020
	(0.094)	(0.092)	(0.087)	(0.082)	(0.096)	(0.094)	(0.096)	(0.101)	(0.109)	(0.097)
Treatment: Scientist – Female	0.066	0.043	-0.106	-0.139	-0.193	-0.155	-0.036	-0.010	-0.036	-0.097
	(0.107)	(0.103)	(0.095)	(0.093)	(0.118)	(0.116)	(0.102)	(0.097)	(0.112)	(0.100)
Treatment: Scientist – Male	0.080	0.088	-0.066	-0.050	-0.035	0.033	-0.082	-0.018	0.020	0.028
	(0.104)	(0.106)	(0.098)	(0.101)	(0.105)	(0.110)	(0.110)	(0.112)	(0.117)	(0.104)
Reference Group Mean	0.08	0.08	0.14	0.14	0.06	0.06	0.10	0.10	0.04	0.04
Reference Group SD	0.93	0.93	0.91	0.91	0.93	0.93	0.93	0.93	1.01	1.01
Observations	798	798	798	798	798	798	798	798	798	798
Session Day	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓		✓		✓

Notes: The table reports treatment effects on willingness to discuss different topics (standardized), separating the treated sample into the four speaker-identity arms: female scientist, male scientist, female influencer, and male influencer. The omitted category is the female influencer treatment arm. All specifications include prefecture, session-day, and session-order fixed effects, as well as the pre-treatment (baseline) measures of trust in vaccines and trust in the healthcare system. Controls include the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A27: Mediation Analysis: Trust in Vaccines

	(1)	(2)	(3)	(4)	(5)	(6)
	Charismatic	Reliable	Knowledgeable	Sympathetic	Attractive	Passionate
Male Influencer vs. Female Influencer						
Total Effect	0.153***	0.153***	0.153***	0.153***	0.153***	0.153***
	(0.049)	(0.049)	(0.049)	(0.049)	(0.049)	(0.049)
Indirect Effect	-0.050***	-0.018	-0.051***	-0.028	-0.063***	-0.092***
	(0.018)	(0.019)	(0.019)	(0.017)	(0.014)	(0.023)
Direct Effect	0.204***	0.172***	0.204***	0.181***	0.217***	0.246***
	(0.049)	(0.049)	(0.048)	(0.051)	(0.050)	(0.052)
Female Scientist vs. Female Influencer						
Total Effect	0.116**	0.116**	0.116**	0.116**	0.116**	0.116**
	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)
Indirect Effect	-0.022	0.104***	0.134***	0.012	-0.048***	-0.039***
	(0.014)	(0.022)	(0.025)	(0.017)	(0.014)	(0.014)
Direct Effect	0.139***	0.013	-0.018	0.104**	0.164***	0.155***
	(0.050)	(0.051)	(0.055)	(0.050)	(0.050)	(0.052)
Male Scientist vs. Female Influencer						
Total Effect	0.122**	0.122**	0.122**	0.122**	0.122**	0.122**
	(0.049)	(0.049)	(0.049)	(0.049)	(0.049)	(0.049)
Indirect Effect	-0.014	0.109***	0.135***	0.033**	-0.046***	-0.026*
	(0.013)	(0.023)	(0.027)	(0.016)	(0.014)	(0.014)
Direct Effect	0.135***	0.012	-0.013	0.088*	0.168***	0.148***
	(0.047)	(0.049)	(0.056)	(0.046)	(0.047)	(0.049)
Session Day	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

Notes: The table reports mediation estimates of the speaker-specific treatment effects on Trust in Vaccines through participants' post-treatment evaluations of the speaker. The analysis is restricted to participants assigned to an information treatment and compares the male influencer, female scientist, and male scientist treatment arms with the omitted female influencer treatment arm. Each column considers one speaker characteristic as a potential mediator. The total effect reports the overall difference in post-treatment Trust in Vaccines between each treatment arm and the female influencer, before conditioning on the mediator. The indirect effect captures the portion of this difference associated with differences in the corresponding speaker evaluation. It is calculated as the product of (i) the effect of assignment to each speaker, relative to the female influencer, on the mediator and (ii) the association between the mediator and post-treatment Trust in Vaccines, conditional on speaker assignment. The direct effect is the remaining difference in post-treatment Trust in Vaccines between each speaker treatment arm and the female influencer after conditioning on the corresponding mediator. In this linear setting, the total effect is equal to the sum of the indirect and direct effects. Each specification controls for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. We additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors for the indirect effects are obtained from 1,000 bootstrap replications, resampling at the session level to preserve the clustered structure of the experiment. Standard errors for the total and direct effects are clustered at the session level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A28: Mediation Analysis: Trust in the Healthcare System

	(1)	(2)	(3)	(4)	(5)	(6)
	Charismatic	Reliable	Knowledgeable	Sympathetic	Attractive	Passionate
Male Influencer vs. Female Influencer						
Total Effect	0.046	0.046	0.046	0.046	0.046	0.046
	(0.057)	(0.057)	(0.057)	(0.057)	(0.057)	(0.057)
Indirect Effect	-0.050***	-0.014	-0.048**	-0.023	-0.076***	-0.069***
	(0.018)	(0.019)	(0.021)	(0.016)	(0.018)	(0.022)
Direct Effect	0.096	0.059	0.094	0.069	0.122**	0.115*
	(0.058)	(0.058)	(0.061)	(0.057)	(0.060)	(0.059)
Female Scientist vs. Female Influencer						
Total Effect	0.048	0.048	0.048	0.048	0.048	0.048
	(0.062)	(0.062)	(0.062)	(0.062)	(0.062)	(0.062)
Indirect Effect	-0.019	0.106***	0.137***	0.014	-0.056***	-0.027**
	(0.014)	(0.025)	(0.030)	(0.016)	(0.015)	(0.013)
Direct Effect	0.067	-0.058	-0.089	0.034	0.104*	0.075
	(0.060)	(0.062)	(0.069)	(0.062)	(0.062)	(0.063)
Male Scientist vs. Female Influencer						
Total Effect	0.007	0.007	0.007	0.007	0.007	0.007
	(0.055)	(0.055)	(0.055)	(0.055)	(0.055)	(0.055)
Indirect Effect	-0.014	0.104***	0.134***	0.029*	-0.057***	-0.020*
	(0.013)	(0.026)	(0.031)	(0.015)	(0.016)	(0.012)
Direct Effect	0.022	-0.097*	-0.127**	-0.022	0.064	0.027
	(0.054)	(0.056)	(0.061)	(0.055)	(0.056)	(0.056)
Session Day	✓	✓	✓	✓	✓	✓
Session Order	✓	✓	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓

Notes: The table reports mediation estimates of the speaker-specific treatment effects on Trust in the Healthcare System through participants' post-treatment evaluations of the speaker. The analysis is restricted to participants assigned to an information treatment and compares the male influencer, female scientist, and male scientist treatment arms with the omitted female influencer treatment arm. Each column considers one speaker characteristic as a potential mediator. The total effect reports the overall difference in post-treatment Trust in the Healthcare System between each treatment arm and the female influencer, before conditioning on the mediator. The indirect effect captures the portion of this difference associated with differences in the corresponding speaker evaluation. It is calculated as the product of (i) the effect of assignment to each speaker, relative to the female influencer, on the mediator and (ii) the association between the mediator and post-treatment Trust in the Healthcare System, conditional on speaker assignment. The direct effect is the remaining difference in post-treatment Trust in the Healthcare System between each speaker treatment arm and the female influencer after conditioning on the corresponding mediator. In this linear setting, the total effect is equal to the sum of the indirect and direct effects. Each specification controls for the baseline value of the outcome, as well as prefecture, session-day, and session-order fixed effects. We additionally control for the full set of participant characteristics and session-size controls described in Section 5. Standard errors for the indirect effects are obtained from 1,000 bootstrap replications, resampling at the session level to preserve the clustered structure of the experiment. Standard errors for the total and direct effects are clustered at the session level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A29: Stability Exercise: Influencer Effects Controlling for Prior Familiarity and Opinion

	Trust in Vaccines			
	(1)	(2)	(3)	(4)
Male Influencer	0.150***	0.104	0.129**	0.097
	(0.051)	(0.075)	(0.052)	(0.076)
Reference Group Mean	0.17	0.17	0.17	0.17
Reference Group SD	0.99	0.99	0.99	0.99
Observations	381	381	381	381
Session Day	✓	✓	✓	✓
Session Order	✓	✓	✓	✓
Prefecture FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Knowledge Control		✓		✓
Opinion Control			✓	✓

Notes: The table reports the differential effect of assignment to the male influencer relative to the female influencer on post-treatment trust in vaccines. The sample is restricted to participants assigned to one of the two influencer treatments and for whom both prior familiarity with and prior opinion of the assigned influencer are observed. Column (1) reports the baseline specification. Column (2) additionally controls for prior familiarity with the assigned influencer, Column (3) controls instead for prior opinion of the assigned influencer, and Column (4) includes both controls. Prior familiarity and prior opinion were measured before participants were exposed to the experimental message. All specifications control for the corresponding pre-treatment measure of trust in vaccines, treatment group, prefecture, session-day, and session-order fixed effects, as well as the full set of participant characteristics and session-size controls described in Section 5. Standard errors, clustered at the session level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Experimental Protocol and Questionnaires

B.1 Pre-Registration

The experiment was registered in the AEA RCT Registry as trial AEARCTR-0016954 under the title "Beyond the Facts: Who Changes Minds—Professors or Influencers?" (Dioikitopoulos, Megalokonomou, Sartori, and Zenou 2025). The registration was submitted on October 7, 2025 and published on October 13, 2025, seven weeks before data collection began. Data collection started on December 1, 2025, approximately four weeks after the date anticipated in the registration, as additional time was required to complete piloting of the VR environment and equipment setup across headsets.

Table B1 maps each element of the registration to its implementation in the paper. The registration has not been amended since its publication.

Three aspects of implementation merit explicit disclosure. First, the registration planned 65 sessions of approximately 30 students each, for a target of 2,000 observations. Session attendance was lower than anticipated: the daily median number of participants per session ranged between four and seven, rather than the planned thirty. To reach the largest feasible sample, we expanded the number of sessions from the planned 65 to 197 across two waves, yielding a final sample of 1,245 participants. Because the realized sample falls short of the registered target, we complement the scientist–influencer comparison with equivalence bounds (Section 6.1), so that null results are interpreted relative to the precision the realized sample affords. Second, the registration specifies random assignment of treatment arms at the session level, with all participants in a session receiving the same treatment; this describes the assignment of sessions to the control, video, and VR arms. Within information-treatment sessions, the identity of the speaker was additionally randomized at the individual level, a finer randomization than the registration describes. Third, the registration lists the primary outcome families—participants' beliefs about vaccinations, the healthcare system, and health professionals—but no analysis plan, secondary outcomes, or minimum detectable effect sizes were filed. Accordingly, the construction of the outcome indices by principal component analysis, the heterogeneity and mechanism analyses, the mediation exercises, and the follow-up analyses were not pre-specified; we treat them as exploratory and, for the heterogeneity results, report false-discovery-rate (FDR) corrections (Tables A21–A23).

Table B1: Pre-Registration and Implementation

Registry Element	As Registered (October 2025)	As Implemented
Interventions	Identical informational content delivered by a scientific expert or a popular non-expert communicator; immersive versus conventional communication channels	Identical script delivered by scientists or influencers (male and female), through VR or standard video, against a neutral-content control (Section 2.2)
Primary outcomes	Beliefs about vaccinations, the healthcare system, and health professionals	Trust in Vaccines and Trust in the Healthcare System indices; trust in health professionals is a component of the latter (Section 3.2)
Secondary outcomes	None registered	Speaker evaluations, visual attention, willingness to discuss, follow-up outcomes: exploratory
Randomization unit	Experimental sessions; clustered treatment	Sessions assigned to control/video/VR; speaker identity additionally randomized within information sessions (Section 2.2)
Inference	Clustered assignment acknowledged	Standard errors clustered at the session level throughout
Planned sample	65 sessions $\times$ $\sim$30 students; 2,000 observations	197 scheduled sessions across two waves; 1,245 participants
Analysis plan	None filed	Index construction, heterogeneity, mechanisms, mediation, and persistence analyses labeled exploratory; FDR corrections for heterogeneity

Notes: The table maps the elements of AEA RCT Registry trial AEARCTR-0016954 to their implementation in the paper. The registration was published on October 13, 2025; data collection began on December 1, 2025. Registered fields are quoted or paraphrased from the public registration (Dioikitopoulos, Megalokonomou, Sartori, and Zenou 2025).

B.2 Cost-Effectiveness Calculation

The experimental estimates also permit a back-of-envelope assessment of the cost-effectiveness of each delivery choice. Consider first the medium. Delivering the message as a standard online video costs on the order of €0.02–€0.05 per completed view through paid social media placement; combined with our same-medium estimate of 0.22 standard deviations (Table A11), information provision through conventional video costs roughly two to three euro cents per 0.1 standard deviations of improved vaccine trust per person. Immersive delivery is substantially more expensive: a consumer VR headset currently retails for approximately €300 and, amortized over a campaign of two thousand exposures with staffed administration, implies a cost of roughly €0.80–€1.00 per exposure. Set against the estimated VR premium of 0.11 standard deviations, the immersive medium purchases additional belief change at approximately €0.70–€0.90 per 0.1 standard deviations—about thirty times the unit cost of conventional video. Two observations qualify this multiple. First, it is falling rapidly: the headsets used in this study cost approximately €1,000 at the time of data collection, while devices of comparable capability now retail for less than a third of that amount. Second, the calculation implies that immersive delivery clears a cost-benefit test precisely where our heterogeneity results locate the largest returns: among initially hesitant individuals, for whom treatment effects are roughly three times larger and for whom we observe a nine-percentage-point increase in reported willingness to vaccinate, and in settings such as schools, universities, and clinics where audiences are already assembled and equipment costs can be amortized intensively.

The messenger margin admits a similar calculation. A mid-tier social media influencer with an audience comparable to our speakers' commands a fee in the low thousands of euros for sponsored content, whereas a physician's recording time costs an order of magnitude less. The influencer's fee, however, embeds distribution: it delivers the message organically to an audience of hundreds of thousands of followers—approximately one to two euro cents per exposure—and, crucially, to the audience whose prior familiarity our results indicate makes the message most effective. The physician's video, by contrast, must purchase its audience. Because persuasion per exposure is statistically comparable across the two messenger types, the choice between them reduces to reach and audience familiarity rather than credibility—the same conclusion the model delivers analytically.

B.3 Experimental Protocol

The study was implemented in Greece in two waves, in December 2025 and in February–March 2026, and involved 1,245 participants. Recruitment was conducted through centralized university communication channels and flyers disseminated around the university campus (see Figure B1 for the flyer we used for the recruitment in the first wave of the study). Upon registering, participants were randomly assigned to a treatment group (either control, VR, or video group) and were then offered a set of available laboratory sessions corresponding to that assignment (a breakdown of the schedule is shown in Figure B2), so that they could choose among the sessions available within their assigned group. Upon arrival, participants completed the registration process and received a unique individual code, which was used to access the survey and, when applicable, the virtual reality session.

All experimental sessions were supervised by the research team. Every laboratory session was staffed by three to four research assistants from a team of 17 trained for the study. Research assistants checked participants in, verified identifiers, launched the survey, and, in VR sessions, fitted the headsets and assisted with the acclimatization sequence. Staffing levels and the session protocol were the same across the control, video, and VR arms. For the VR treatment, headsets and accessories were sanitized before each use, and participants received on-site assistance to ensure correct interaction with the environment and to avoid technical disruptions. Participants were compensated for their time with a €5 voucher (see Figure B3) and entered into a raffle for several prizes, including tablets and Apple devices.

B.4 Questionnaires

B.4.1 Experimental Questionnaire—Part 1 (Survey Administered During the Laboratory Session)

Block 0: Consent and Session Information

1. Participation consent.
   - Accept
   - Refuse
2. Please select today's date: *(Dropdown selection: e.g., 01/12/2025)*
3. Please select the session number that applies to you: *(Dropdown selection: e.g., Session 1)*
4. Please enter the code written on the ticket provided at the beginning of the experiment: *(Text Entry: e.g., 1057)*

Block 1: Demographic Characteristics

5. In what year were you born? *(Dropdown selection: e.g., 1992)*
6. What is your gender?
   - Male
   - Female
   - Other
   - Prefer not to say
7. What year are you currently in at university?
   - 1st Year BSc
   - 2nd Year BSc
   - 3rd Year BSc
   - 4th Year BSc
   - 1st Year MSc
   - 2nd Year MSc
   - PhD
   - Other
8. What course are you studying at the Athens University of Economics and Business?
   - International and European Economic Studies – International Economics and Finance

- International and European Economic Studies – International and European Political Economy
- Economics – Economic Theory and Policy
- Economics – Business Economics and Finance
- Economics – International and European Economics
- Management Science and Technology – Operations Research and Business Analytics
- Management Science and Technology – Operations and Supply Chain Management
- Management Science and Technology – Software and Data Analysis Technologies
- Management Science and Technology – Information Systems and Electronic Business
- Management Science and Technology – Strategy, Entrepreneurship and Human Resources
- Business Administration – Business Administration
- Business Administration – Information Systems Management
- Business Administration – Accounting and Financial Management
- Business Administration – Marketing
- Accounting and Finance – Accounting
- Accounting and Finance – Finance
- Marketing and Communication – International Management, Innovation and Entrepreneurship
- Marketing and Communication – Human Resources Management
- Marketing and Communication – Business Analytics
- Marketing and Communication – Digital Marketing
- Informatics – Theoretical Computer Science
- Informatics – Computer Systems and Networks
- Informatics – Information Systems and Information Security
- Informatics – Databases and Knowledge Management
- Informatics – Operational Research and Economics of Information Technology
- Informatics – Computational Mathematics and Scientific Calculations
- Statistics – Statistics
- Other

9. What is your annual family income?

- 0–9,999
- 10,000–19,999
- 20,000–29,999
- 30,000–39,999

- 40,000–49,999
- 50,000–59,999
- 60,000–69,999
- More than 70,000
- Prefer not to say

10. What is the highest educational qualification your father has achieved?
    - High School or Lower
    - Professional Training
    - BSc
    - MSc
    - PhD or Higher
11. What is the highest educational qualification your mother has achieved?
    - High School or Lower
    - Professional Training
    - BSc
    - MSc
    - PhD or Higher
12. Which of these categories describes best the occupation of your father?
    - Managers
    - Professionals
    - Technicians and Associate Professionals
    - Clerical Support Workers
    - Service and Sales Workers
    - Skilled Agricultural, Forestry and Fishery Workers
    - Craft and Related Trades Workers
    - Plant and Machine Operators, and Assemblers
    - Other Occupation
    - Armed Forces Occupations
13. Which of these categories describes best the occupation of your mother?
    - Managers
    - Professionals

- Technicians and Associate Professionals
- Clerical Support Workers
- Service and Sales Workers
- Skilled Agricultural, Forestry and Fishery Workers
- Craft and Related Trades Workers
- Plant and Machine Operators, and Assemblers
- Other Occupation
- Armed Forces Occupations

14. Which political ideology attracts you or represents you the most?
    - Left
    - Center-left
    - Center
    - Center-right
    - Right
    - I am not interested in politics

15. During the past 30 days, on how many days did you smoke cigarettes?
    - Never
    - Rarely – no more than once a week
    - Often – more than once a week
    - Every day or almost every day

16. During the past 30 days, on how many days did you do physical exercise?
    - Never
    - Rarely – no more than once a week
    - Often – more than once a week
    - Every day or almost every day

Block 2: Trust in Science, Social Media, and Content Creators

17. On a typical day, how much time do you spend on social media platforms?
    - Less than 30 minutes
    - 30–59 minutes
    - 1–2 hours

- 2–3 hours
- 3–5 hours
- More than 5 hours

18. For which of the following purposes do you primarily use social media? *Select one option.*
    - To keep in touch with family/friends
    - To follow news/current events
    - To share my opinions
    - For entertainment (music, videos, memes)
    - For professional networking
    - I don't use social media

19. Where do you usually get most of your news about what is happening in the world?
    - Newspapers
    - Television
    - Social Media
    - I don't follow the news

20. How much do you trust the information you see on social media platforms? *Slider: Scale from 0 (Not at All) to 10 (Completely).*

21. Below you will see a list of Greek content creators. How familiar are you with them? *Slider: Scale from 1 (I do not know them at all) to 5 (Extremely familiar).*
    - Influencer Name #1 – Instagram Handle #1
    - Influencer Name #2 – Instagram Handle #2
    - Influencer Name #3 – Instagram Handle #3
    - Influencer Name #4 – Instagram Handle #4
    - Influencer Name #5 – Instagram Handle #5
    - Influencer Name #6 – Instagram Handle #6

22. Thinking about the content produced by each of the following creators, what is your overall opinion of them? *Slider: Scale from -2 (Very Negative) to 2 (Very Positive).*
    - Influencer Name #1 – Instagram Handle #1
    - Influencer Name #2 – Instagram Handle #2
    - Influencer Name #3 – Instagram Handle #3
    - Influencer Name #4 – Instagram Handle #4
    - Influencer Name #5 – Instagram Handle #5
    - Influencer Name #6 – Instagram Handle #6

Block 3: Pre-treatment Beliefs on Vaccines and the Healthcare System

23. How much do you think vaccines are: *Slider: Scale from 0 (Not at All) to 10 (Completely).*
    - Effective at Preventing Pandemics and Severe Diseases
    - Safe for the Health of the Subject
    - Important for your Own Health
    - Important for Public Health
24. At this moment, how much do you trust the following? *Slider: Scale from 0 (Not at All) to 10 (Completely).*
    - Conventional Medicine and Surgery
    - Healthcare Professionals
    - Public Health Authorities
25. Would you accept to take a vaccine that is recommended to you by public health authorities?
    - No
    - Yes
    - It depends on the vaccine
    - It depends on the doctor's advice

Block 4: Friendship Network

26. We will now ask you about your friendships during your time at the Athens University of Economics and Business. Please list up to 10 friends who also attend the university.

    It is important for the study that you answer these questions carefully and truthfully. Your chances of winning one of the prizes (Apple products, such as phones, tablets, and wireless earphones) will increase based on the accuracy of your responses, which will be automatically cross-checked with the responses of your fellow students. This verification process is carried out exclusively by computer: the individuals you mention will never see your responses, and you will never know who has mentioned you.

    Once the data have been collected, your first and last names will be replaced with a code that does not allow identification. Therefore, the researchers analyzing the data will never have access to either your identity or the identity of your friends.

    If you do not have any friends to list, or if you have fewer than 10, simply leave the remaining fields empty.
27. Please indicate name of person #1: *(Text Entry: e.g., Yannis)*
28. Please indicate surname of person #1: *(Text Entry: e.g., Papadopoulos)*
29. In a typical university week, about how many hours do you spend with person #1 outside of class?

- 0 hours (We don't spend time together in a typical week)
- Less than an hour
- 1–2 hours
- 3–5 hours
- 6–10 hours
- 11–20 hours
- More than 20 hours

30. What best describes your relationship with person #1?
    - Mostly social/personal (friends, leisure activities)
    - Mostly academic (study together, coursework, projects)
    - Both, but mainly social
    - Both, but mainly academic

31. How close do you feel to person #1? *Slider: Scale from 0 (Not at All) to 10 (Extremely).*

    Questions 27–31 were then repeated for up to 10 friends.

Block 5: Treatment Block

The exact wording and format of the treatment block varied depending on experimental assignment. Participants were randomly assigned to one of the following conditions:

- **VR Treatment Condition**

  "Please now wear the VR headset. Once you complete the experience, you will be able to continue with the questionnaire."

- **Video Treatment Condition**

  "You will now watch a short video, which is part of a public health campaign by the Athens University of Economics and Business and Monash University. Please watch it carefully."

  The video was directly embedded within the survey platform.

- **Control Condition**

  "You will now watch a short video about London. Please watch it carefully."

  The video was directly embedded within the survey platform.

Block 6: Post-treatment Beliefs on Vaccines and Public Health

32. At this moment, to what extent do you believe vaccines are: *Slider: Scale from 0 (Not at All) to 10 (Completely).*

    - Effective at Preventing Pandemics and Severe Diseases
    - Safe for the Health of the Subject
    - Important for your Own Health
    - Important for Public Health

33. At this moment, how much do you trust the following? *Slider: Scale from 0 (Not at All) to 10 (Completely).*

    - Conventional Medicine and Surgery
    - Healthcare Professionals
    - Public Health Authorities

34. After completing this questionnaire, how likely are you to start discussions about the following topics? *Slider: Scale from 0 (Not at All Likely) to 10 (Very Likely).*

    - Vaccines
    - Public Health
    - Social Media
    - Artificial Intelligence (AI)
    - Politics

Block 7: Mechanisms

The following questions were asked only to participants assigned to the VR and Video Treatment conditions.

35. Overall, how positive or negative was your impression of the speaker? *Slider: Scale from -5 (Very Negative) to 5 (Very Positive).*

    - Speaker
    - Message

36. We would now like to know how you perceived the speaker. For each characteristic below, please evaluate how well it describes the speaker. *Slider: Scale from 0 (Not at All) to 10 (Extremely).*

    - Charismatic
    - Persuasive
    - Reliable

- Knowledgeable
- Sympathetic
- Attractive
- Passionate

Block 8: Contact Information and Prize Draw

37. Please provide your first name. *(Text Entry: e.g., Yannis)*

38. Please provide your surname. *(Text Entry: e.g., Papadopoulos)*

39. Please provide your email address. *(Text Entry: e.g., Yannis.Papadopoulos@gmail.com)*

40. Please provide your phone number. *(Text Entry: e.g., +30 1234567890)*

B.4.2 Questionnaire—Part 2 (Online Follow-Up Survey)

Block 0: Consent

1. Participation consent.
    - Accept
    - Refuse

Block 1: Follow-up Beliefs and Vaccination Behavior

2. How much do you think vaccines are: *Slider: Scale from 0 (Not at All) to 10 (Completely).*
    - Effective at Preventing Pandemics and Severe Diseases
    - Safe for the Health of the Subject
    - Important for your Own Health
    - Important for Public Health
3. At this moment, how much do you trust the following? *Slider: Scale from 0 (Not at All) to 10 (Completely).*
    - Conventional Medicine and Surgery
    - Healthcare Professionals
    - Public Health Authorities
4. Are you more willing to get vaccinated since you participated in the first wave of the survey?
    - No
    - Yes
5. Have you received any additional vaccines since you participated in the first wave of the survey?
    - No
    - Yes

Block 2: Friendship Network

6. We will now ask you about your friendships during your time at the Athens University of Economics and Business. Please list up to 10 friends who also attend the university.

    It is important for the study that you answer these questions carefully and truthfully. Your chances of winning one of the prizes (Apple products, such as phones, tablets, and wireless earphones) will increase based on the accuracy of your responses, which will be automatically cross-checked with the responses of your fellow students. This verification process is carried out exclusively by computer: the individuals you mention will never see your responses, and you will never know who has mentioned you.

Once the data have been collected, your first and last names will be replaced with a code that does not allow identification. Therefore, the researchers analyzing the data will never have access to either your identity or the identity of your friends.

If you do not have any friends to list, or if you have fewer than 10, simply leave the remaining fields empty.

7. Please indicate name of person #1: *(Text Entry: e.g., Yannis)*

8. Please indicate surname of person #1: *(Text Entry: e.g., Papadopoulos)*

9. In a typical university week, about how many hours do you spend with person #1 outside of class?

   - 0 hours (We don't spend time together in a typical week)
   - Less than an hour
   - 1–2 hours
   - 3–5 hours
   - 6–10 hours
   - 11–20 hours
   - More than 20 hours

10. What best describes your relationship with person #1?

    - Mostly social/personal (friends, leisure activities)
    - Mostly academic (study together, coursework, projects)
    - Both, but mainly social
    - Both, but mainly academic

11. How close do you feel to person #1? *Slider: Scale from 0 (Not at All) to 10 (Extremely).*

    Questions 7–11 were then repeated for up to 10 friends.

Block 3: Contact Information and Prize Draw

12. Please provide your first name. *(Text Entry: e.g., Yannis)*

13. Please provide your surname. *(Text Entry: e.g., Papadopoulos)*

14. Please provide your email address. *(Text Entry: e.g., Yannis.Papadopoulos@gmail.com)*

15. Please provide your phone number. *(Text Entry: e.g., +30 1234567890)*

B.4.3 Questionnaire—Part 3 (Online Follow-Up Survey)

Block 0: Consent

1. Participation consent.

    - Accept
    - Refuse

Block 1: Follow-up Beliefs and Vaccination Behavior

2. How much do you think vaccines are: *Slider: Scale from 0 (Not at All) to 10 (Completely).*

    - Effective at Preventing Pandemics and Severe Diseases
    - Safe for the Health of the Subject
    - Important for your Own Health
    - Important for Public Health

3. At this moment, how much do you trust the following? *Slider: Scale from 0 (Not at All) to 10 (Completely).*

    - Conventional Medicine and Surgery
    - Healthcare Professionals
    - Public Health Authorities

4. Are you more willing to get vaccinated since you participated in the first wave of the survey?

    - No
    - Yes

5. Have you received any additional vaccines since you participated in the first wave of the survey?

    - No
    - Yes

6. Which vaccine(s) did you receive (you can select more than one answer)?

    - Human Papillomavirus Vaccine (HPV)
    - Measles Vaccine (MMR)
    - Diphtheria-tetanus-pertussis Vaccine (DTP/Tdap)
    - Other (Please Specify)

Figure B1: Recruitment Flyer Distributed to Students

Notes: The figure reports the flyer that was circulated around campus and distributed through university mailing lists to recruit participants into the study for the first wave of data collection.

Figure B2: Example of Survey Session Schedule During Week 2

08/12/2025				09/12/2025				10/12/2025				11/12/2025				12/12/2025				Weekend
1	11:05	11:25	**CONTROL**	**1**	11:00	11:40	**VR**	**1**	11:10	11:40	**VIDEO**	**1**	11:00	11:40	**VR**	**1**	11:05	11:25	**CONTROL**	
	11:25	11:35	break		11:40	11:50	break		11:40	11:50	break		11:40	11:50	break		11:25	11:35	break	
2	11:35	12:05	**VIDEO**	**2**	11:50	12:20	**VIDEO**	**2**	11:50	12:10	**CONTROL**	**2**	11:50	12:20	**VIDEO**	**2**	11:35	12:05	**VIDEO**	
	12:05	12:15	break		12:20	12:30	break		12:10	12:20	break		12:20	12:30	break		12:05	12:15	break	
3	12:15	12:55	**VR**	**3**	12:30	12:50	**CONTROL**	**3**	12:20	13:00	**VR**	**3**	12:30	12:50	**CONTROL**	**3**	12:15	12:55	**VR**	
	12:55	13:05	break		12:50	13:00	break		13:00	13:10	break		12:50	13:00	break		12:55	13:05	break	
4	13:05	13:25	**CONTROL**	**4**	13:00	13:40	**VR**	**4**	13:10	13:40	**VIDEO**	**4**	13:00	13:40	**VR**	**4**	13:05	13:25	**CONTROL**	
	13:25	13:35	break		13:40	13:50	break		13:40	13:50	break		13:40	13:50	break		13:25	13:35	break	
5	13:35	14:15	**VR**	**5**	13:50	14:10	**CONTROL**	**5**	13:50	14:30	**VR**	**5**	13:50	14:10	**CONTROL**	**5**	13:35	14:15	**VR**	
	14:15	14:25	break		14:10	14:20	break		14:30	14:40	break		14:10	14:20	break		14:15	14:25	break	
6	14:25	14:55	**VIDEO**	**6**	14:20	14:50	**VIDEO**	**6**	14:40	15:00	**CONTROL**	**6**	14:20	14:50	**VIDEO**	**6**	14:25	14:55	**VIDEO**	
	14:55	15:05	break		14:50	15:00	break		15:00	15:10	break		14:50	15:00	break		14:55	15:05	break	
7	15:05	15:25	**CONTROL**	**7**	15:00	15:40	**VR**	**7**	15:10	15:40	**VIDEO**	**7**	15:00	15:40	**VR**	**7**	15:05	15:25	**CONTROL**	
	15:25	15:35	break		15:40	15:50	break		15:40	15:50	break		15:40	15:50	break		15:25	15:35	break	
8	15:35	16:05	**VIDEO**	**8**	15:50	16:20	**VIDEO**	**8**	15:50	16:30	**VR**	**8**	15:50	16:20	**VIDEO**	**8**	15:35	16:05	**VIDEO**	
	16:05	16:15	break		16:20	16:30	break		16:30	16:40	break		16:20	16:30	break		16:05	16:15	break	
9	16:15	16:55	**VR**	**9**	16:30	16:50	**CONTROL**	**9**	16:40	17:00	**CONTROL**	**9**	16:30	16:50	**CONTROL**	**9**	16:15	16:55	**VR**	
	16:55	17:05	break		16:50	17:00	break		17:00	17:10	break		16:50	17:00	break		16:55	17:05	break	
10	17:05	17:25	**CONTROL**	**10**	17:00	17:40	**VR**	**10**	17:10	17:40	**VIDEO**	**10**	17:00	17:20	**CONTROL**	**10**	17:05	17:25	**CONTROL**	
	17:25	17:35	break		17:40	17:50	break		17:40	17:50	break		17:20	17:30	break		17:25	17:35	break	
11	17:35	18:05	**VIDEO**	**11**	17:50	18:10	**CONTROL**	**11**	17:50	18:30	**VR**	**11**	17:30	18:10	**VR**	**11**	17:35	18:05	**VIDEO**	

Notes: The figure reports an example of the survey schedule during the second week of data collection. Session types were randomized across days and time slots throughout the experimental period.

Figure B3: Participation Voucher Provided to Students

Notes: The figure reports the €5 voucher that students who participated in the study received as compensation for completing the survey.

C Proofs of Propositions

Proof of Proposition 1: From equation (18), since $\gamma_m^i > 0$, $\lambda^{(m)}(a) > 0$, $a > 0$, and $\theta - \mu_0 > 0$ by assumption, the product is strictly positive. For the second part, note that $\lambda^{(m)}(a) = \sigma_0^2/(\sigma_0^2 + \tilde{\sigma}_m^2(a))$, which is strictly increasing in σ_0^2 and approaches 0 as $\sigma_0^2 \to 0$.

Proof of Proposition 2: Taking expectations of (7),

$$\mathbb{E}[\tilde{x}_m(a) - \mu_0 \mid \theta] \;=\; a(\theta - \mu_0),$$

so the expected belief revision is

$$\mathbb{E}\Big[\Delta\mu^{(m,i)}(a)\Big] = \gamma_m^i \, \lambda^{(m)}(a) \, a \, (\theta - \mu_0), \qquad \lambda^{(m)}(a) = \frac{\sigma_0^2}{\sigma_0^2 + a^2\sigma_m^2 + \sigma_\eta^2}, \tag{C1}$$

where the Bayesian weight $\lambda^{(m)}(a)$ is *decreasing* in a (more absorbed content also means more absorbed message noise). The engagement multiplier $F(a)$ is defined in (8). Differentiating,

$$F'(a) = \frac{\sigma_0^2 \left(\sigma_0^2 + \sigma_\eta^2 - a^2\sigma_m^2\right)}{\left(\sigma_0^2 + \sigma_\eta^2 + a^2\sigma_m^2\right)^2},$$

which is strictly positive for all $a \in (0, 1]$ whenever $\sigma_m^2 \le \sigma_0^2 + \sigma_\eta^2$ (implied by the strict inequality in Assumption 2), since then $a^2\sigma_m^2 \le \sigma_m^2 \le \sigma_0^2 + \sigma_\eta^2$ with strict inequality for $a < 1$ or $\sigma_m^2 < \sigma_0^2 + \sigma_\eta^2$. Let $a = a_\ell + \alpha_m$ (equation (10)). By Assumption 3, $a_{VR} > a_{Video}$, hence $F(a_{VR} + \alpha_m) > F(a_{Video} + \alpha_m)$ and the expected revision (C1) is strictly larger under VR.

Proof of Proposition 3: Equating the two expected revisions from (18) immediately yields (19). Substituting (15) and rearranging gives (20).

Proof of Proposition 4:

Part (a). From Proposition 2, F is increasing in a. Since $a_I = a_\ell + \alpha_I > a_\ell + \alpha_S = a_S$, the claim follows.

Part (b). Set $r_I^i f_I^i = r_S^i f_S^i = 0$, so $\gamma_I^i = \bar{\gamma}_I$, $\gamma_S^i = \bar{\gamma}_S$; parity requires $F^{(I)}(a_I)/F^{(S)}(a_S) = \bar{\gamma}_S/\bar{\gamma}_I$. By the proof of Proposition 2, $F^{(I)}$ is strictly increasing on $[0, 1]$, hence invertible on its range, with $F^{(I)}(1) = \sigma_0^2/(\sigma_0^2 + \sigma_\eta^2 + \sigma_I^2)$ its supremum on $a_I \in [0, 1]$. The displayed threshold $\bar{\alpha}$ solves $F^{(I)}(a_\ell + \alpha_S + \bar{\alpha}) = (\bar{\gamma}_S/\bar{\gamma}_I)F^{(S)}(a_S)$ whenever the right-hand side lies in the range of $F^{(I)}$ on $[0, 1]$; monotonicity of $F^{(I)}$ then gives parity, or a strict advantage to the influencer, for every $\alpha_I - \alpha_S \ge \bar{\alpha}$. If the target value exceeds $F^{(I)}(1)$, boundedness of $F^{(I)}$ on $[0, 1]$ rules out an engagement-only solution, and (i) or (ii) becomes necessary.

Part (c). Identity (21) follows by adding and subtracting $\gamma_S^i F_I$. Decompose $F_I - F_S$ into a pure-engagement term and a pure-noise term:

$$F_I - F_S = \underbrace{\left[F^{(I)}(a_I) - F^{(I)}(a_S)\right]}_{\text{engagement, } \ge 0} + \underbrace{\left[F^{(I)}(a_S) - F^{(S)}(a_S)\right]}_{\text{noise, } \le 0}.$$

The bound on the engagement channel follows from the mean value theorem,

$$F^{(I)}(a_I) - F^{(I)}(a_S) = F^{(I)\prime}(\xi)(\alpha_I - \alpha_S) \leq \frac{\sigma_0^2}{\sigma_0^2 + \sigma_\eta^2}(\alpha_I - \alpha_S).$$

Since the noise bracket is ≤ 0, the displayed bound on $\gamma_S^i(F_I - F_S)$ follows a fortiori. Condition (22) then compares the two channels directly.

Proof of Proposition 5: Part (a): taking expectations of (25) immediately yields the expression; $\Pi_\ell = \kappa_\ell + (1-\rho)(1-\kappa_\ell) \in (\kappa_\ell, 1)$ for $\rho \in (0,1)$ and $\kappa_\ell \in (0,1)$. Part (b): it is straightforward since Π_ℓ depends on ℓ. Part (c): if $\kappa_{\ell,V} > \kappa_{\ell,H}$ (higher consolidation for vaccine trust), then $\Pi_{\ell,V} > \Pi_{\ell,H}$, generating greater persistence on the vaccine dimension.