Preferences, Beliefs, and Follow-Through: Understanding the Demand for Flu Vaccines*

Daniel W. Sacks† Justin Sydnor‡

March 20, 2025

VERY PRELIMINARY
PLEASE DO NOT CITE OR CIRCULATE

Abstract

Vaccine take-up is lower than targets set by public health agencies. While contemporary accounts of vaccine hesitancy focus on misinformation and scientifically unwarranted beliefs, economic models of vaccination emphasize individuals trading off perceived costs and benefits of vaccination. As little is known about how people perceive these costs and benefits, we conduct an online experiment to measure preferences for and beliefs about characteristics of the flu vaccine. We have two main findings. First, nearly all the difference in vaccination rates across vaccine intention groups is explained by differences in preferences over and beliefs about vaccine effectiveness and short-run side effects. Relative to vaccine confident participants, vaccine hesitant participants value effectiveness lower, dislike side effects more, and believe vaccines are less effective and have worse side effects. Scientifically unwarranted beliefs appear to play a quantitatively minor role in flu vaccine hesitancy. Second, beliefs are biased in a way that favors vaccination, except that vaccine hesitant respondents overestimate the short-run side effects of vaccination. Correcting misperceptions of short-run side effects would substantially increase vaccination rates.

JEL codes: D83, I12, I18, L65
Key words: vaccines, belief elicitation, demand estimation

*Support for this research was provided by the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin–Madison with funding from the Wisconsin Alumni Research Foundation.

†Risk & Insurance, Wisconsin School of Business, University of Wisconsin - Madison (dan.sacks@wisc.edu).

‡Risk & Insurance, Wisconsin School of Business, University of Wisconsin - Madison and NBER. (justin.sydnor@wisc.edu.edu).

1 Introduction

Seasonal diseases such as influenza, COVID-19, and Respiratory Syncytial Virus inflict millions of Americans annually, and kill tens of thousands (Naquin, 2024; Ahmad, 2024). The worldwide burden of these diseases is greater still (Paget et al., 2019). Safe vaccines exist to reduce the incidence of these diseases, yet demand for these vaccines is low: despite a price of zero, and a CDC recommendation that all adults vaccinate against COVID and Flu, in typical only 18-35 percent of adults report doing so (Kriss, 2024). Low vaccination rates are concerning both because of the disease burden of flu and because vaccines against infectious disease can prevent the spread of that disease, a classic externality (Ward, 2014; White, 2021; Freedman et al., 2022)

This low vaccination rate may be explained by at least three hypotheses. One hypothesis centers on misinformation, i.e. scientifically unwarranted beliefs about extreme adverse consequences of vaccinations, including views that vaccines cause autism, cancer, death or severe illness, or the targeted disease itself.[1] An alternative hypothesis centers on the intention-action gap: many people would like to get vaccinated, but fail to follow through. This hypothesis implicitly motivates the many efforts to increase vaccination through nudges and other light-touch interventions (for example, Milkman et al. (2022); Patel et al. (2023)). Third, economic epidemiology assumes that vaccination decisions are privately rational–though possibly socially suboptimal–given individual's preferences over health and their beliefs about prevailing risks and vaccine consequences (Philipson, 2000; Goodkin-Gold et al., 2020).

While these potential reasons for vaccine hesitancy are not mutually exclusive, distinguishing the relative importance of these hypotheses is critical for multiple reasons. First, low vaccine demand is a challenge to standard models of demand for health care, which imply high take-up of low-cost, health-improving options. Second, many policy interventions aim to increase vaccination rates. The welfare consequences of these interventions depend on the reasons for low vaccination rates. If low vaccination rates reflect a normatively-relevant aversion to vaccination, for example because of costly side effects, then these interventions can reduce welfare, although externalities may make them desirable. If however

[1] For academic examinations of this hypothesis, see for example Horne et al. (2015) or Nyhan and Reifler (2015). For an example of a policymaker embracing this hypothesis, see Green (2025).

low vaccination rates reflect normatively-irrelevant factors, such as incorrect beliefs or failure to follow-through, then increasing vaccination rates is unambiguously positive.

In this paper, we provide experimental evidence on vaccine intentions, beliefs about vaccines, preferences for different potential characteristics of vaccines, and the accuracy of people's beliefs. We measure beliefs by asking people about their perceived flu risk and the consequences–in terms of flu rate and short-term side effects—of vaccination. We measure (stated) preferences using a conjoint design. A key advantage of this design is that we can vary characteristics which are otherwise fixed: vaccine effectiveness, side effects, delivery mode (needle or inhalation), convenience, and price (which we proxy with different gift card rewards). Effectiveness and side effects are natural measures for the primary consequences of vaccination. We additionally vary delivery mode because many sources indicate needle phobia is a common reason for vaccine hesitancy (e.g. Freeman et al. (2023)), and needle-free vaccines are increasingly available (U.S. Food and Drug Administration, 2024). We focus on the flu vaccine because the health burden of flu is substantial, vaccination rates are low, and the vaccine has existed for many decades, making it relatively straightforward to describe its characteristics and elicit beliefs about it.

We have two main findings. First, differences in preferences over vaccine characteristics, and differences in beliefs about those characteristics, are centrally important in explaining differences in vaccination intentions. Second, although beliefs are predictive of real-world risks and effects, they are generally biased in a way that is favorable to vaccination: participants believe flu is more common than the Centers for Disease Control and Prevention (2024a) estimates, and more dangerous. An important exception is that vaccine hesitant participants overestimate short-term side effects of vaccination. These findings have several implications. First, because most apparent vaccine hesitancy is explained by beliefs and preferences over effectiveness and short-term side effects, scientifically unwarranted beliefs play a quantitatively smaller role. Second, because beliefs are mostly excessively favorable to vaccination, correcting misperceptions in general is unlikely to increase vaccination rates. However, correcting misperceptions of short-run side effects can raise vaccination rates.

Our analysis begins by examining beliefs. To measure beliefs, we ask participants—in a percentage

point scale—the likelihood of them getting the flu if unvaccinated, as well as the consequences of flu if they get it, the reduction in flu risk from vaccination, and the likelihood of experiencing side effects from vaccination. We ask beliefs about average, national flu risk, and about personal risk. Our first result is that, perhaps surprisingly, vaccine confident and hesitant participants have similar beliefs about flu risk and vaccine effects for the average American. Thus stated belief differences do not reflect general vaccine animosity, or general desire to signal a particular social stance. Second, however, vaccine confident and hesitant participants have sharply different beliefs about flu risk and vaccination effects for themselves. Respondents intending to vaccinate think they are twice as likely to get the flu (absent vaccination) as do respondents not intending to vaccinate, with similar differences in hospitalization and death rates. Those intending to vaccinate also think the effectiveness of the flu vaccine is fifty percent higher. While respondents who intend to vaccinate or are already vaccinated perceive no negative side effects from vaccination, respondents who do not intend to vaccinate perceive a 15 percentage point chance that vaccination causes side effects severe enough to inhibit at least a day of activity. Despite their differences, vaccine confident and vaccine hesitant alike hold beliefs that are roughly inline with what the CDC reports. This is possible because the CDC reports a wide range of uncertainty for flu-related statistics. Vaccine hesitant beliefs are at the bottom of the range, and vaccine confident are at or above the upper end of the range.

Our second result allows us to quantify the relative importance of beliefs and preferences in explaining differences in vaccine intentions. To estimate preferences, we apply a hierarchical Bayes model of demand for vaccines to the conjoint preference data. We estimate strong preferences both for vaccine effectiveness and against side effects. Preferences against side effects reduce the willingness-to-pay for vaccines, and failing to account for side effects would result in substantial underestimates of the value of vaccine effectiveness.

We use the estimated preferences and group-specific beliefs to simulate vaccinate take-up under different scenarios. Predicted demand for the current flu vaccine differs sharply between groups, providing an initial validity check. People intending to vaccinate or already vaccinated have a predicted vaccination rate around 90 percent; for respondents not intending to vaccinate, the rate is 3 percent. To quantify the

role of beliefs and preferences, we re-simulate demand but changing the mean values for preference or beliefs for groups that have not yet vaccinated to match the means for the already vaccinated. We find that preferences alone are more important than beliefs alone, but the two are complementary, because vaccine confident respondents both value effectiveness more highly, and believe effectiveness is greater, than vaccine hesitant respondents. Preferences for effectiveness and side effects are especially important; convenience and needle aversion are quantitatively unimportant.

We model preferences for vaccine effectiveness and side-effects, but we do not explicitly model preferences for or beliefs about scientifically unwarranted possibilities, like vaccines causing cancer or having long-run, unspecified health consequences. Nonetheless, our estimates imply a small role for such scientifically unwarranted beliefs. Our model of vaccine preferences includes an individual-level "inside option" term which reflects all unmodelled aspects of overall vaccine preferences. If scientifically unwarranted beliefs were important, we would expect that differences in the value of the inside option term would explain differences in vaccination rates between vaccine confident and hesitant respondents. Instead we find a quantitatively small role for the value of the inside option, and hence extreme beliefs. This small role for extreme beliefs is consistent with basic patterns in the data: we find that a majority of vaccine hesitant respondents say they would vaccinate for a vaccine with sufficiently high effectiveness and low short term side effects effects. This choice would be hard to understand if many vaccine hesitant people believed vaccines had terrible long-term side effects.

These results rely on stated preference data, without which it would be difficult or impossible to identify preferences over vaccine characteristics. However stated preference data come with well-known concerns relative to revealed preferences (e.g. Diamond and Hausman (1994); Hausman (2012)). To address these concerns, we show that stated preference data are informative for predicting real choices in a number of contexts, and for predicting incentivized choices in a follow-up experiment condition. For example, our estimated demand system predicts vaccination rates close to, albeit higher than, the actual vaccination rate.

Our first two results show the importance of preferences and beliefs for explaining vaccine intentions.

The final part of our paper investigates the accuracy of these beliefs, and how respondents follow through on their intentions.

Our third result is that beliefs are inaccurate: people overstate their own flu risk and the short term side effects of vaccination. In our initial survey, respondents intending to vaccinate thought they would have a 30 percent chance of getting the flu; two months later, 15 percent of those respondents reported they probably got the flu. Also in the initial survey, marginally vaccine hesitant respondents on average believed getting vaccinated would increase their experience of severe side effects by five percentage points. However, when we induced this group to vaccinate with randomly assigned incentives, we found their experienced symptom rate actually decreased. For this marginally hesitant group, our demand model implies that inaccurate beliefs about side effects reduce vaccination rates by 20 percentage points. Thus beliefs are inaccurate and these inaccuracies likely contribute to vaccine decisions. However, because many participant' beliefs are biased towards vaccination, for the average participant, correcting beliefs would not necessarily raise average vaccination rates.

Our final result is that, while actual vaccinations fall short of intended or expected vaccination rates, a light-touch intervention increases vaccinations in the short run. After we assigned incentives to vaccinate, we asked whether they expected to vaccinate before the follow-up survey. While these expectations are predictive of vaccination rates, they severely overstate it, consistent with follow-through problems. The final part of our baseline survey was a randomly assigned "nudge-style" message, encouraging subjects to follow a link to sign up for a vaccination. This nudge increased short run vaccination rates by 7.5 percentage points, an effect concentrated entirely among people who expected to get vaccinated. At a two month follow-up, however, the effect had fallen to zero, and two-month vaccination rates remain well below expected vaccination rates. This result suggests that, in addition to preferences and beliefs, follow-through problems are an important contributor to vaccination rates.

This paper contributes to the literature on vaccine take-up, and, more generally, on health and risky behaviors. Much of this literature uses natural experiments to measure the impact of specific interventions on vaccine take-up, for example school mandates, government recommendation, sibling vaccina-

tions, or disease prevalence (for example, Abrevaya and Mulligan (2011); Lawler (2017); Carpenter and Lawler (2019); Churchill et al. (2024); Humlum et al. (2024)). Other work uses designed experiments to assess the impact of nudges, information provision, endorsements, advertising, or incentives on vaccinations (for example, Bronchetti et al. (2015); Milkman et al. (2022); Patel et al. (2023); Athey et al. (2023); Alsan and Eichmeyer (2024); Larsen et al. (2023); Schneider et al. (2023); Ho et al. (2023); Campos-Mercade et al. (2024)). A third strand of this literature uses stated preference experiments to find very low willingness to pay for vaccines (e.g. Carlin et al. (2022)).

Our primary contribution to this literature is to quantify the importance of beliefs and preferences for vaccine take-up decisions.[2] This quantification is important for several reasons. First, without estimates of preferences and beliefs, it is impossible to assess the welfare consequences of vaccine mandates, or even lighter touch interventions such as nudges (Allcott et al., 2022). Second, understanding beliefs is crucial for designing and assessing information provision, endorsements, and other persuasive efforts to increase vaccination. Such interventions are most helpful when they target mistaken beliefs that are relevant for vaccination decision. We find that the most relevant beliefs—about vaccine effectiveness and flu risk—are roughly accurate or too favorable to vaccination. Extreme, unwarranted beliefs are less important, and hence correcting them is less likely to increase vaccine intentions, consistent with the surprising experimental evidence in, for example, Nyhan and Reifler (2015).

2 Background

We study vaccines to protect against the influenza virus ("flu"), which annually infects roughly a billion people world wide , including millions in America, where flu season runs from September through March (Centers for Disease Control and Prevention, 2024a; World Health Organization, 2025). A small fraction of infections are severe enough to warrant hospitalizations. The CDC estimates that during the 2023-24 flu season, Americans experienced 40 million cases of symptomatic flu illness, requiring 470,000 hospi-

[2]Some papers, including Athey et al. (2023), measure the impact of their experiments on beliefs about vaccines and illness risk, but they use qualitative measures that cannot be benchmarked to objective risk or consequences, nor mapped to vaccine take-up.

talizations, of which 28,000 resulted in deaths (Center for Disease Control, 2025a). There is considerable uncertainty in these estimates, however, because they are derived not from population-wide surveillance but from models estimated on data from a network of reporting hospitals (Center for Disease Control, 2025b; Rolfes et al., 2018; Reed et al., 2015). Accounting for uncertainty in these estimates, the rate of symptomatic influenza in the US in 2023-24 was 10 to 18 percent, the rate of hospitalization conditional on a flu case was 0.6-2.3 percent, and the rate of death conditional on flu hospitalization was 2.8-14 percent. [3]

These risks are heterogeneous in the population. For example, the CDC estimates that the flu death unconditional on infection is 8.3 per 100,000 for all ages, but increases from 1.9 for 18-49 year-olds to 9.1 for 50-64 year-olds and 32.1 for people 65 and older (Center for Disease Control, 2025b). It is likely that aspects other than age influence flu risk. As flu is an infectious disease, people with more and closer interactions with others—especially likely flu carriers—are themselves more likely to contract the flu. Thus family and occupational characteristics influence flu risk, as do personal choices such as frequency and mode of socialization.

Although vaccination mitigates the harms of the flu, the flu vaccine is not perfectly effective. The CDC reports vaccine effectiveness rates by season, and they range from a low 19 percent in 2014-15 to a high of 60 percent, although most years are between 38 and 60 percent. Vaccine effectiveness refer to the percent reduction in flu illness among the vaccinated, relative to the control (Centers for Disease Control and Prevention, 2024b). [4] The CDC relies on multiple studies of effectiveness, and these studies consider effectiveness against a range of outcomes, from outpatient visits (Rolfes et al., 2019) to hospitalizations (**?**). The effectiveness of the flu vaccine depends on multiple factors. Ever year the vaccine is prepared to match the strains forecasted to circulate, but this forecast is not perfect. The immunity provided by the vaccine wanes over time. For both those reasons, the CDC recommends annual vaccinations. Different

[3] The CDC reports total counts rather than conditional rates. We convert to conditional rates because this format is easier to express to our survey participants. To make bounds on rates, we use extrema of uncertainty intervals for counts. For example the lower bound on the hospitalization rate is the lower bound on the hospitalization count uncertainty interval divided by the upper bound on the infection count uncertainty interval.

[4] The CDC uses the term vaccine "efficacy" to refer to estimates from clinical trials, and effectiveness to real-world studies. They regularly report effectiveness, not efficacy.

vaccination options are available, an injection or a nasal spray, and the injection is more effective than the nasal spray.

Actual vaccination rates are far below CDC's recommend and target levels. Appendix Figure A.1 shows the weekly vaccination rate for 2023-24 flu season. This starts below 20 percent in October and rises steady to about 45 percent by January, where it roughly remains. Although vaccination rates are increasing in age, even among 65 and older, the vaccination rate is below the CDC's target of 70 percent until the very end of the flu season. Low vaccination rates indicate vaccine hesitancy, which the World Health Organization's Strategic Advisory Group of Experts (WHO SAGE) defines as "the delay in acceptance or refusal of vaccination despite availability of vaccine services" (MacDonald et al., 2015).

A common view of vaccine hesitancy is that it derives from scientifically unwarranted views of vaccines. Indeed, researchers sometimes focus on correcting such misconceptions, like the belief that the flu vaccine can cause the flu (Alsan and Eichmeyer, 2024; Nyhan and Reifler, 2015). But under the WHO SAGE definition, vaccine hesitancy can reflect a misinformed decision based on, for example, incorrect views about vaccine safety or effectiveness. But it can also correct a fully informed tradeoff between moderate short-term side effects or other vaccine costs and an imperfectly effective vaccine against moderate health risks.

The WHO SAGE scale to measure hesitancy does not distinguish between these views (Larson et al., 2015). The scale includes 10 items, with Likert-scale questions about the health benefits of vaccines, their side side effects, and trust in vaccine programs and health care providers (Larson et al., 2015). [5] Thus the scale classifies people as hesitant for a range of circumstances: if they refuse to vaccinate at all, but also if they correctly view themselves at low risk of disease, or say that a vaccine is ineffective because it has 60% effectiveness should be called effective. Thus, quantitative measures of beliefs, rather than qualitative agree/disagree, are necessary to understand vaccine hesitancy, as are measure of people's preferences over vaccine characteristics. This paper paper evidence on these preferences and beliefs, and quantifies their importance for vaccination take-up.

[5] The Larson et al. (2015) scale was developed for pediatric vaccinations. However modifications of it for adults are also used, e.g. Akel et al. (2021).

3 Experiment Design

3.1 Experiment design

Our experiment design addresses four goals: to elicit beliefs about flu risk and flu vaccine characteristics, to measure preferences over flu vaccine characteristics, to validate stated preferences, and to assess the accuracy of individual's beliefs. Figure 1 shows the design of the experiment, which ran across three waves. Wave 1 contains the core of the experiment and waves 2 and 3 are simple follow ups to measure vaccination, experienced side effects, and experienced flu. We run the experiment on Prolific, a platform for recruiting experiment subjects. The study was approved by the University of Wisconsin IRB under protocol number 2024-1162.

We begin in wave 1 by asking a set of questions about demographics, health, trust in doctors and the government, and and vaccine intentions. The health questions include questions about flu risk and vaccine contraindications. Blocks 2 and 3 contain our belief elicitation and stated preference tasks, which we describe in more detail in Sections 4 and 5. We randomize the order of these sections, and we include an attention check at the start of each section Within the belief block, we ask first about objective quantities (for example, percent of Americans who got the flu in 2023-24), and then about subjective quantities (personal flu risk). After asking about objective quantities, we reveal the CDC's estimates of these quantities to a random subset of participants. In the stated preference task, we ask respondents to choose between vaccine A or vaccine B, and vary the characteristics of the vaccines. We then ask whether they would vaccinate at all, given the choice of A or B. We exclude subjects whose answers are inconsistent between these questions (e.g. they prefer A to B, but B to A or no vaccinate).

The final block of wave 1 consists of a set of incentivized choices, which we use both to validate the conjoint and to provide experimental variation in vaccination status. In each of five choices, we ask participants what reward they would prefer for participating in a follow up survey: an unconditional $3, or $2 plus a bonus that they get if they show proof of vaccination in the follow up survey. We set the bonus amount to $1.50, $3, $10, $20, and $50. We explain that we will pick one of their answers and actually

implement it. Answers to these questions therefore provide incentivized information on willingness to accept for vaccination. Because these questions only make sense for people who are able to vaccinate, we exclude from block 4 people who are already vaccinated and people who have a contraindication for vaccination (the latter group is exlcuded from our main sample).

After respondents make their choice, we pick one of their choices at random to implement. We tell them what incentive (if any) they face to vaccinate, and ask them if they expect to vaccinate before the next survey. Finally, for a subset[6] of respondents, we randomize assignment to a "nudge" treatment, where we encourage them to sign up for a vaccine appointment, and provide links to pharmacies to do so.

Wave 1 of our survey therefore contains the main measures and manipulations. We measure outcomes—vaccination and experienced side effects—in two follow-up waves. Wave 2 opened two weeks after wave 1. All subjects not already vaccinated and not contraindicated are eligible for wave 2. We sent multiple reminder messages to all subjects encouraging them to complete wave 2. We ask about experienced side effects (as we explain in Section 6) and vaccination. Then, for subjects offered a reward to vaccinate, we ask for a record upload. To earn their reward, subjects upload a pharmacy receipt or a screenshot of their vaccine record. In Wave 3 of the survey, which is open to the same set of subjects as wave 2, and ran two months after the initial survey, we again measure vaccination, and also ask about flu experiences. To avoid experimenter demand effects, we mask our intentions by opening with questions about pharmacy shopping, and include questions about non-vaccine medications, and non-flu illness.

3.2 Managing induced demand and expressive answering

Two natural concerns with surveys are that participants may respond to perceived experimenter or survey enumerator demand, and that they may answer questions, particularly belief questions, to express their identity rather than to fully communicate their actual beliefs. Our online, anonymous context is less likely to suffer from these concerns than in-person surveys, and our pilot testing and open response questions revealed that participants did not vaccinations to be an especially sensitive subject. Nonetheless

[6] We only randomize respondents who choose the $50 vaccine incentive, to avoid encouraging vaccination among those most hostile to it.

we address these concerns in several ways.

In the experiment design, we followed best practices to minimize induce demand (Stantcheva, 2023; Haaland et al., 2023). Our recruitment material did not mention the study purpose but emphasized that we are "non-partisan researchers."[7] When we define vaccine characteristics, we do so in a way that clearly indicates upsides as well as downsides (e.g. imperfect effectiveness, severe side effects), providing face-saving language. We framed all belief elicitation in a neutral and open-ended way. Our wave three survey uses unrelated questions to minimize the connection to our earlier survey, and indeed virtually no respondents thought the survey was about vaccine preferences or below when asked at the end.

Our results suggest that these efforts were successful. Our results for beliefs, described in Section 4 show that participants do not use the belief questions to express their identity. Specifically, under the hypothesis of expressive answering, we might expect that vaccine hesitant subjects would predict very low flu risk in the US as a whole, and low vaccine effectiveness. However we show below that beliefs about national averages are not correlated with vaccine intentions. Finally, we show that beliefs and preferences are individually predictive of actual vaccinations (Appendix Table A.2), and beliefs about personal flu risk are predictive of actual flu incidence 7, meaning that our belief and preference measures capture distinct information which corresponds to choices and experiences.

3.3 Analysis sample

We limit the subject pool to Americans aged 18 or older, taking surveys on a computer, and living in the 26 states that allow people to look up their vaccine history in a state-maintained database. This last limitation ensures that participants can verify their vaccination. Because the Prolific sample skews young, like other online samples, we imposed an age quota that 10 percent of our sample be 65 or older.[8]

Table 1 shows how the sample size changes as we impose our inclusion criteria. 3,406 participants started wave 1. Our main analysis sample excludes respondents with at least one failed attention check,

[7] Our consent page included the title of the study, "preferences, beliefs, and the demand for vaccine," because our IRB required it.

[8] Filling this quota was difficult because our elder subjects took longer to completed the survey and so saw a lower hourly reward. We ended up increasing the payment for this group.

at least one inconsistent stated preference,[9] multiple survey attempts, and non-missing key variables. We further exclude the 8 percent of participants who indicate a contraindication to the flu vaccine. Our main analysis sample consists of the remaining 2,357 subjects. Of these, 1,757 are not already vaccinated and hence eligible for the incentivized part of the survey, as well as the follow up surveys. The actual follow up rate is 96 percent at wave 2 (about two weeks later) and 84 percent at wave 3 (about two months later).

We summarize the demographic characteristics of the main analysis sample in Table 2. For comparison, we also report the characteristics of respondents in the 2023 BRFSS, a nationally representative survey of health behaviors. Like most online panels, our sample skews towards young, high educated, and high income. However, when we reweight the BRFSS to match our sample on age, sex, income, we find a flu vaccination rate of 41 percent, very similar to our sample's vaccination rate of 42 percent. Thus our sample is roughly representative on this key metric.[10] Other work measuring vaccine hesitancy has also used online panelists (e.g. Akel et al. (2021)), and we find similar hesitancy rates.[11]

We organize much of analysis around flu vaccine intentions. The sample is roughly evenly divided across four intention groups: already vaccinated, intend to vaccinate, may or may not vaccinate, and do not intend to vaccinate. The BRFSS does not ask an analogous version of this question; instead it asks about whether the respondent has gotten a flu shot in the last 12 months. Actual vaccination rates are similar in our sample and in the BRFSS.[12]

[9] Consistency here is a very weak requirement. In each preference task, we ask respondents if they prefer vaccine A to B, given a requirement they choose one. Then we ask them which option they prefer of A/B/no vaccine. We require that the answers be consistent between these questions.

[10] In calculating our sample vaccination rate, we include people with contraindications, because excluding them reduces the vaccination rate relative to BRFSS, and we do not have a measure of contraindication in the BRFSS.

[11] Specifically, we use two items from the commonly used vaccine hesitancy scales developed by Larson et al. (2015) and adapted by Akel et al. (2021) for flu vaccines in the US adults. The Larson et al. scale has 10-items, two of which we use here with appropriate modification for our context. The other items are qualitative qualitative questions about perceived risks and effectiveness of vaccines. We do not ask those questions because they are redundant with out belief questions. Akel et al. (2021) find mean scores of 2.29 and 1.99 on our common questions; we find mean scores of 1.99 and 1.91. We report the distribution of responses in Appendix Table A.1.

[12] Our sample excludes people with a self-reported vaccine contraindication. We show below that including these people brings our vaccination rate more in line with the BRFSS rate. To measure the in-sample vaccination rate, we set vaccination to 1 for people who say they are already vaccinated at baseline or who report vaccination in the wave 3 survey. We set it to missing for people not already vaccinated at baseline and not present in the wave 3 survey.

3.4 Pre-registration

This experiment was pre-registered in the AEA RCT registry.[13] The pre-registration describes our general goals (measure preferences and beliefs for vaccines, biases in perceived side effects) and our interventions. It also describes our sample size rule, which was dictated by our budget. We preregistered some models and statistical tests, as well as some validity tests for the stated preference data. Appendix F describes some omissions from our pre-registration (i.e. requried decisions we failed to anticipate), and explains how the analyses in the paper deviate from the preregistration (generally minimally, though we did not pre-register our focus on vaccine intention groups).

We emphasize here that survey wave 3 was not preregistered. Two factors motivated this additional wave. First, we realized we had an opportunity to measure whether flu experiences aligned with flu beliefs. Second, we intended to use *verified* vaccinations, with record uploads, to measure vaccination status for people facing incentives to vaccinate. However many people failed to upload their records, and qualitative comments indicated that it was difficult for them to do so. To measure vaccinations without giving an incentive to lie, we ran a third survey wave with a number of unrealted questions, as well as flu and vaccine questions. Qualitative questions at the end of wave three show that less than 1 percent of subjects saw a connection to the wave 1 study.

4 Beliefs about flu risk and vaccine characteristics

We describe our approach to measuring beliefs, show that vaccine intentions do not correlate with beliefs about national average flu risk or vaccine effects, but correlate closely with beliefs about personal risk and vaccine effects.

[13] Registration here: `https://www.socialscienceregistry.org/trials/14763`. We registered immediately before launching the live survey.

4.1 Measuring beliefs about flu risk and vaccine effects

Measurement target: Our goal is to measure people beliefs' about their personal flu risk and consequences of vaccination. By "flu risk", we mean the likelihood of adverse flu-related outcomes, for the participants themselves. This distribution is multidimensional because it encompasses both the likelihood of catching the flu and the severity of illness conditional on disease. We focus on flu risk conditional on not vaccinating, because this risk, along with the consequences of vaccination, determines the incentive to vaccinate.

We focus on beliefs about two consequences of vaccination: effectiveness and short-term side effects. Effectiveness is the percent reduction in flu risk from vaccination, relative to not being vaccinated, for the respondents themselves. Short-term side effects are the effect of vaccination on the likelihood of experiencing disruptive short-term side effects, such as fatigue, muscle soreness, or fever. We do not attempt to measure beliefs about long run consequences of vaccination, and therefore do not measure the most extreme beliefs, such as the view that vaccines cause cancer or autism, or weaken the immune system. We focus on effectiveness and short-term side effects for two reasons. First, we can compare beliefs to objective benchmarks for our experiment sample (i.e. we do not need to wait decades to collect outcome data), and second, they are important aspects of the vaccination decision where disagreement is possible and scientific uncertainty exists, because systematic data are not collected on the experience of these types of side effects.[14]

Measuring beliefs about flu risk: The exact text of all our elicitations is in B; here we summarize. To elicit beliefs about flu risks, we begin by saying "Some people choose to get the flu vaccine, and others choose not to. For the next few questions, we'd like to ask your beliefs about what will happen to you **if you choose not to get the flu vaccine.**" We provide face-saving language (Stantcheva, 2023) and emphasize that we are asking about risk conditional on not vaccinating. Then we ask about likelihood of

[14]Side effects are tracked through the vaccine adverse event reporting system (VAERS). The system relies on voluntary reports by patients and providers. In principle this database could capture moderate side effects such as a day of muscle ache. In practice it is extremely unlikely that VAERS captures moderate side effects, because it contains fewer than 10,000 entries for flu vaccines in 2023-24, despite more than 150 million flu vaccinations, implying a rate of less than 1 per 10,000. We find a subjective side effect rate of roughly 5 percent on average.

getting the flu in the next year, likelihood of severe case given flu, likelihood of hospitalization given flu, and likelihood of death given hospitalization. We ask about these conditional probabilities because the objective unconditional probabilities are very small (for example, the unconditional hospitalization rate is 0.14 percent), and we worried that elicited probabilities would be very noisy.

For each of our probability questions, we first ask a 7-point multiple choice/Likert scale response, and then we ask a quantitative version. The Likert responses range from "would definitely happen" to "would definitely not happen." We ask these multiple choice questions for two reasons. First, they can be helpful for correcting measurement error in the quantitative responses.[15] Second, we use them to set the scale for the quantitative elicitation, as we explain presently.

To measure flu risk beliefs quantitatively, we tell participants, "We just asked you about how likely certain flu-related events would be, if you don't get the flu vaccine. We're going to ask you the same questions, but this time we'd like your best guess of the numerical percent for each of them." We then re-ask our questions, but asking for the "percent change" of the risk. We let participants adjust a slider, and we set the range of the slider based on their multiple choice answers. If they say an outcome would not happen or is unlikely, the slider ranges from 0 to 25; otherwise it ranges from 0 to 100. Although sliders likely introduce trembling-hand measurement error (Stantcheva, 2023), we use sliders because they let us indicate the correct scale of the response (i.e. 0 to 100 for a percent, not 0 to 1). This is important for our beliefs, which in some cases could be less than 1 percent (e.g. hospitalization risk) and in others be above 10 percent (flu risk).

Beliefs about vaccine characteristics: We measure beliefs about vaccine effectiveness similar to how we measure beliefs about flu risk. We first ask "If you do get the flu vaccine, how much would you say that reduces your chances of getting the flu?" and give multiple choice answers (eliminates, substantially reduces, ..., increases). We then use a text entry box to get a number; we let participants enter a negative number (to indicate increasing risk), although none do so. Notice that vaccine effectiveness here measures reduction in the likelihood of getting the flu, rather than reductions in disease severity con-

[15] Giglio et al. (2021) also use multiple survey measures of the same belief to correct for measurement error.

ditional on flu. We focus on this single measure to avoid survey fatigue, and because the CDC reports aggregate effectiveness, not distinguishing among different dimensions of effectiveness (Centers for Disease Control and Prevention, 2024b).

Eliciting beliefs about side effect rates is more challenging because it is less intuitive to think of a side effect rate as a causal contrast (relative to vaccine effectiveness, where the causal contrast is natural). The problem is that the language "side effect rate" does not evoke a causal contrast, and may refer to the expected level of an outcome. [16] In our pilot studies, we found that asking about side effect rates of vaccination produced very high side effect rates.[17] We therefore measure vaccine side effects using a different, two-step approach.

First we ask participants their likelihood of experiencing side-effects symptoms over the next two weeks, assuming they do not get vaccinated. To do so, we instruct them:

> After getting vaccinated, sometimes people experience headaches, muscle aches, upset stomach, tiredness, joint pain, fever, or general discomfort. These side effects are **severe if they disrupt your daily activities.**
>
> People experience these symptoms for all kinds of reasons, not just from the flu vaccine.

Our instructions emphasize the time frame, the definition of severity, and the possibility of experiencing these side effects without being vaccinated. We emphasize the two week time period because our second wave occurs two weeks later, and asks about experienced side effects over the prior two weeks.

Next we ask elicit beliefs about these symptoms:

> For this question, imagine that **you do not get vaccinated** in the next two weeks.
>
> Thinking about yourself, what do you think are the chances you would experience a case of headache, muscle ache, upset stomach, tiredness, joint pain or fever, **severe enough to disrupt your daily activities** in the next two weeks?

[16] Smith et al. (2021) notes the importance of making a clear causal contrast when eliciting subjective expectations of treatment effects.

[17] We found a believed side effect rate of roughly 25 percent, but vaccine information sheets report symptom-specific side effect rates of 0-2 percent, for 3-5 symptoms.

We again give a multiple choice version of the question, and then an open response version, and then we repeat the process by instructing participants to assume they get vaccinated over the next two weeks. The quantitative questions give us believed rate of experience symptoms of side effects, without vaccination and with vaccination, say $rate_i(0)$ and $rate_i(1)$. We then construct a believed side effect rate of vaccination as $rate_i(1) - rate_i(0)$. ask the analogous question but

Objective benchmarks: While our main interest is on beliefs about personal risks and vaccine consequence ("subjective beliefs"), we also elicit beliefs about analogous objects for the average American ("objective beliefs"), because they provide useful benchmarks, letting us quickly spot gross errors, if present. To elicit objective beliefs, we ask participants what fraction of Americans got the flu, what fraction of those with the flu were hospitalized, and what fraction of those hospitalized died, for 2023-24. We also ask them what they think the CDC says is the effectiveness of the flu vaccine. We do not elicit objective beliefs for severe flu or side effects because there is no clear benchmark to compare them to.

Measurement error: It is likely that our beliefs are measured with error. These errors include trembling hand issues with sliders, but also more substantial errors owing from innumeracy or unfamiliarity with the underlying concepts. These errors mean that individual-level responses may not be reliable indicators of individual-level beliefs. Most of our analysis therefore focuses on group-level average responses, which we take to be close to group level mean beliefs. This interpretation is justified by the large body of research demonstrating that aggregate predictions often perform well even when individual predictions are poor, i.e. the "wisdom of the crowd."

Discussion: This approach to measuring beliefs is simple and direct. It allows us to compare reported beliefs to objective probabilities as well as actual experiences, as we elicit beliefs about the aspects of flu risk that the CDC reports, and in follow up surveys, our experience measures correspond to our belief measures. Because we elicit beliefs about measurable, reported vaccine characteristics, we can combine our belief estimates with our stated preference demand estimates to predict vaccination rates. The alternative approach to inferring beliefs—recovering them from choices—requires a rational expectations assumptions to separate beliefs from preferences (Manski, 2004). This assumption is especially strong in

the vaccine context. This is because people may reasonably believe that their personal flu risk and vaccination effects are different from the average person, but they have little personal data on which to base their beliefs, as people with strong views would find personal experimentation very costly.

Measuring beliefs in this way is subject to three concerns, however. First, we do not ask participants about scientifically unwarranted beliefs, for example that vaccines have long term negative health consequences. These beliefs would be difficult to elicit and we lack a way to manipulate this characteristics in the stated preference experiment (since we would not offer a vaccine with a long term health risk). This concern means that we do not directly measure extreme beliefs, but, as we explain in our demand model, we pick up these beliefs as a residual.

Second, a common objection to measuring quantitative beliefs is that some participants may lack the numeracy to give meaningful answers to this question. A third and related concern is that our belief measures are not incentivized, and subjects therefore may not provide adequate effort in answering. While some amount of measurement error is likely inevitable, we limit its impact by using multiple choice questions along with quantitative ones, at times using the multiple choices as instruments for the quantitative responses. We show below that objective beliefs are reasonably close to what the CDC reports, suggesting innumeracy is not a strong problem. Further evidence can be seen in the distribution of beliefs, Appendix Figure A.2, which shows that respondents use the whole response scale, beliefs cluster towards zero where such clustering is appropriate, and relatively few respondents choose 50% for beliefs, except for vaccine effectiveness, where such a belief is warranted.[18] We also show in Appendix Table A.2 that beliefs strongly predict vaccination rates, in the way we would expect, showing there is signal as well as noise in the belief measures. This table also shows that beliefs predict vaccination rates conditional on our preference estimates,[19] indicating that beliefs and preferences are distinct objects, each with their own explanatory power. We further show in Appendix C that beliefs are predicted by individual level risk factors (age and

[18] Respondents tend to round to 5 percentage point increments, although not always. These patterns are typical of quantitative beliefs in general (Manski, 2004).

[19] To do this, we use the individual-level posterior mean preferences estimates described in Section 5 to predict vaccination rates at objective vaccine characteristics (50% effectiveness, 1% side effect rates). In a single regression, beliefs and this predicted vaccination rate are strongly predictive of actual vaccination rates.

chronic conditions), although small sample sizes make these predictions noisy at time. Finally, we show that although beliefs are on average biased for realized risk, they are nonetheless predictive of it: people who believe they face a higher flu risk are indeed more likely to get the flu (Figure 7).

4.2 Objective and subjective beliefs

We begin in Figure 2 by showing objective beliefs. The figure shows the mean belief about national risks or effectiveness, by vaccine intention group. The horizontal lines show the 95 percent uncertainty interval reported by the CDC (Centers for Disease Control and Prevention, 2024a).[20] The figure shows two important facts. First, there is very little disagreement across intention groups about national averages. Second, although beliefs are not hugely different from the CDC's estimates, they are generally biased in a way that favors vaccination. People think flu is more common, hospitalization is more likely, and vaccines more effective than the CDC reports. Beliefs about death given hospitalization are consistent with the (wide) range that the CDC reports.

These results show that, on average, beliefs are roughly well-calibrated, and that beliefs are not entirely driven by vaccine preferences. However beliefs about average risk or effectiveness are not necessarily relevant for individual decision making, because for example individuals may believe they are less risky than average. We therefore report the mean beliefs about personal risk and vaccine effects in Figure 3. In addition to the personal analogs of the beliefs in Figure 2, the figure reports beliefs about personal risk of severe flu (but not hospitalization) and personal side effect consequences of vaccination.

Beliefs about personal risks and effects contrast with beliefs about national avearges in two important ways. First, people genearlly believe flu is less risky for themselves, and vaccines less effective, than they believe for the general population. This perhaps reflects the young age distribution of our sample. Second, in contrast to beliefs about national averages, beliefs about personal risks and effects vary sharply with vaccine intentions. Although vaccine confident and vaccine hesitant participants alike roughly agree about overall flu risk and vaccine effectiveness, they differ in how they perceive flu risk and vaccine effects for

[20] The CDC reports overall rates 100,000. We convert all rates to conditional probabilities (unconditional flu risk, hospitalization given flu, and death given hospitalization).

themselves. participants intending to vaccinate perceive greater risk of flu and worse consequences of getting the flu (conditional on not vaccinating) than do participants not intending to vaccinate. participants intending to vaccinate also perceive higher vaccine effectiveness, and much lower side effect consequences, than do people not intending to vaccinate. These differences can be large: the difference in believed flu likelihood is 100 percent. Differences in side effect beliefs are especially large. While people not intending to vaccinate believe that the flu vaccine will increase their experience of disruptive symptoms by 15 percentage points, people not intending to vaccinate believe vaccinating *reduces* their experience of these symptoms, presumably because of the benefit of averted flu.

Despite the substantial differences in beliefs about personal risks and vaccine effects, no group has average beliefs that are far outside the average values reported by the CDC. Across all intention groups, believed risk of flu, and risk of hospitalization given flu, is higher than what CDC estimates in recent years. Perceived vaccine effectiveness is slightly higher for the top intention groups than what the CDC estimates, and slightly lower for the "do not intend to vaccinate group. Of course there is individual-specific heterogeneity in risks and effectiveness, and the CDC estimates need not be the correct value for these groups. The comparison simply shows that no group's beliefs are terribly far from plausible national averages, and even for the groups not intending to vaccinate.

5 Preferences for vaccines

5.1 Conjoint design

We measure preferences for vaccine characteristics using a conjoint design, a standard tool in the marketing literature (e.g.,Allenby et al. (2019)), with many recent economics applications (e.g. Ben-Akiva et al. (2019); Kessler et al. (2019); Moshary et al. (2023); Chan (2022)). In a conjoint design, subjects face a series of choice sets, each containing multiple products, and, possibly, a no purchase option. By carefully varying the characteristics of available products across choice sets, researchers can learn preferences for particular characteristics. In the typical application, and in ours, choices are hypothetical and not incen-

tivized, so we view the results as reflecting stated preferences. Below we describe several tests to validate the stated preference data with revealed preferences.

We focus on five vaccine characteristics that we expect to be important for choice: price, effectiveness, side effects, delivery mechanism (needle or nasal spray), and convenience. Following standard advice for conjoint designs (Ben-Akiva et al., 2019), we try to make attributes and choice setting as natural as possible. Since vaccines rarely have a price—they are covered by insurance—we operationalize price by saying that some vaccines come with gift card rewards. Vaccine effectiveness is reported by CDC and in clinical trials, and our description of effectiveness mirrors this language. Needle delivery is straightforward to describe, as its alternative, a nasal spray.

The most challenging characteristics to operationalize are side effects and convenience. Although side effects are described on vaccine package inserts, the inserts describe not an overall side effect rate, but the rate of many individual side effects: fever, nausea, fatigue, muscle ache, and rash. It is not practical to assess preferences for these individual side effects. Instead we report and vary an overall "severe side effect rate," which we define as the fraction of participants in clinical trials who report a experiencing a side effect (fever, nausea, fatigue, muscle ache, or rash) severe enough to limit their activities for at least one day. This definition corresponds to a grade three side effect. We set the levels of side effect rate at 0, 1, or 5 percent. We choose these levels because they are within the range of reported grade 3 side effect rates.[21] Finally, to operationalize convenience, we ask participants to imagine that the vaccines are available either at their pharmacy or health care provider, or at a clinic that is open 24 hours a day and next to their home (the convenient option).

We give each participant the same 10 choice sets, in a random order. We generated choice sets to maximize D-efficiency. Appendix Table A.3 reports the options available in each choice set, and Appendix Figure A.3 shows how we present the choices to the participants. We randomize the order of the choices. Before beginning the conjoint, we explain the meaning of each characteristics, and we include reminder text when we show each choice. Finally, before the conjoint, we include a baseline choice that resembles

[21] Five percent is about the highest severe side effect rate consistent with package inserts, obtained under the assumption that all reported side effects are mutually exclusive.

the prevailing vaccine. The baseline choice asks participants to choose between a vaccine with no reward, 50 percent effectiveness, 1 percent side effect rate, needle delivery, and not convenient, or no vaccination at all.

5.2 Preference specification, estimation, and identification

Preference specification: To analyze data from the conjoint, we assume that participant i receives conditional indirect utility u from choice j in choice set t:

$$\begin{aligned} u_{ijt} &= -\alpha_i \text{Reward}_{ijt} + \beta_{VE,i}\text{VE}_{ijt} - \beta_{SE,i}\text{SE}_{ijt} \\ &\quad + \beta_{N,i}\text{Needle}_{ijt} + \beta_{C,i}\text{Convenient}_{ijt} + \beta_{0i} + \epsilon_{ijt}, \\ u_{i0t} &= \epsilon_{i0t}. \end{aligned}$$

Here choice $j = 0$ is the no-vaccine option and VE and SE refer to vaccine effectiveness and side effect.

We assume that the individual level parameters are distributed as:

$$(\ln \alpha, \ln \beta_{VE}, \ln \beta_{SE}, \beta_N, \beta_C, \beta_0) \sim N(\mu_i, \Sigma), \epsilon_{ijk} \sim T1EV. \tag{1}$$

This specification imposes a positive sensitivity to reward and to effectiveness, and a negative sensitivity to side effects. For notational convenience, we denote the vector of individual level parameters θ_i.

We consider two specifications for μ_i. Our primary specification lets μ_i vary flexibly by vaccine intention group. In Appendix Figure A.4 we report our core simulation results where μ varies with age, education, and score on a vaccine hesitancy scale (but not intention). This latter specification was prespecified, but parameterizing μ in terms of vaccine intentions makes it straightforward to decompose differences in vaccination rates into preferences and beliefs, as described below. In Appendix Figure A.5, we report simulations where preferences depend on both intentions and our prespecified covariates. Our core results are very similar across these specification.

Our preference specification does not explicitly capture scientifically unwarranted beliefs, such as the

view that vaccines cause autism or harm health in the long-run. We omit such beliefs because we do not have a way to manipulate such characteristics in the conjoint. Nonetheless, our specification implicitly capture the role of such beliefs, via β_{0i}, which is the utility a participant receives from a vaccine with no explicit characteristics, i.e. the value of the "inside option." A participant who holds extremely negative beliefs about a vaccine would have a very negative value of β_{0i}. If very negative beliefs play an important role in vaccine intentions, then we should see that differences in β_{0i} have a large effect on cross-group differences in vaccination rates.

Estimation: We estimate the model using standard Bayesian methods, using the Gibbs-in-Metropolis algorithm implemented by the BAYESM R package (Rossi et al., 2024). We assume conjugate priors for all parameters. We use 300,000 draws, discard the first 30,000, and keep every 300th draw. This simulation procedure produces $N_d(= 900)$ draws from the posterior distribution of individual specific preference parameters, θ_i. We use these extensively in our simulations. We report diagnostics on the resulting Markov chains in Appendix D; the chains quickly exhibit good mixing and stability.

Identification: Our conjoint experiment provides the variation in characteristics needed to identify the distribution of individual preferences. Speaking loosely, the mean of a preference parameter, say β_x is identified by how choice probabilities change when characteristic x changes, holding fixed other characteristics. As we manipulate x experimentally, we are not worried about confounding from unobserved product characteristics. The persistence in person-specific choice errors pins down the cross-person variability in preferences, i.e. Σ. For example, if a person does not choose a product with a high reward, that could indicate a low ϵ or a low α_i. If some people consistently decline products with high rewards, we infer that their α_i is low. These arguments suggest that the model is non-parametrically identified, and with sufficient choice variation, it is. However as we have relatively few choices per individual, we rely on distributional assumptions for the cross-person heterogeneity.

5.3 Simple evidence that vaccine hesitant find some vaccines appealing

Before presenting our demand estimates, we report simple results from the conjoint experiment. Figure 4 shows the fraction of participants choosing to vaccinate, in each vaccine intention group, for three different choice sets. Panel (a) shows that, when faced with a realistic choice set, stated vaccination rates are low, and near zero for people who do not intend to vaccinate.[22] At the best choice set in our data, however, vaccination rates are much higher, approaching 100 percent for people who say they may or may not vaccinate, and 45 percent for people who do not intend to vaccinate. Thus a suitably appealing vaccine option can induce vaccine hesitant to change their mind. One concern, however, is that the best choice set includes an option with no side effects. Vaccine hesitant participants might interpret side effects to include scientifically unwarranted ones. We therefore show in panel (c) the take-up rates in a choice set where all vaccines have a positive side effect rate. Even here, vaccine take-up is fairly high among the uninsured.

5.4 Preference estimates and simulations

Preference estimates: As our interest is in cross-group differences in preferences, we report mean preferences by group in Table 3; Appendix Table A.4 contains information on the distribution of preference parameters. The table shows stark differences in vaccine preferences by vaccine intention group. The average weight on each characteristic is roughly ordered by vaccine intention. Reward sensitivity is almost twice as high for the "do not intend to vaccinate" group as for the "already vaccinated group," and the weight on vaccine effectiveness is only a third as large. All groups are sensitive to side effects, but the most vaccine hesitant place the greatest weight on side effects. No group is especially sensitive to needles or convenience.[23]

[22] The choice sets in this figure are 0, 9, and 8. Appendix Table A.3 reports the characteristics of the offered choices. The vaccination rate may seem surprisingly low in choice set 0 for people already vaccinated or intending to vaccinate, but this is because they believe existing vaccines are much better than the realistic vaccine we offer.

[23] The low weight on convenience might be surprising. Banerjee et al. (2010); Duflo and Banerjee (2011) report that regularly available "vaccination camps" had a dramatic effect on pediatric vaccination rates, arguably reflecting the importance of convenient. Their intervention, however, occurs in a very different context, where the baseline option was to walk miles to a clinic where staff might or might not be present. A likely explanation for the low weight on convenience in our setting is that for many participants, the baseline vaccination site is already fairly convenient.

The "do not intend to vaccinate" group is the only group with a negative average sensitivity to needles, and this sensitivity is small: participants in this group are on average indifferent between a needle and a $3.50 reward, or a a 3 percentage point increase in vaccine effectiveness. Likewise for convenience: for the most sensitive group, a convenient vaccine is equivalent to a roughly $1.25 extra payment on average.

Our preference estimates help rationalize a puzzle from prior findings. Prior work uses stated preference data like ours to back out a willingness to pay (WTP) for vaccination, and uses this willingness to pay to obtain an implied value of a statistical life (VSL; Carlin et al. (2022)). The puzzle is that, at least when applied to the COVID-19 vaccine, this approach yields a very small VSL. This approach, however, ignores all non-financial costs of vaccination, for example side effects or needle aversion.

Accounting for these alternative costs ends up raising the implied VSL substantially. We show this in Table 4. The first column reports the median WTP for vaccination, defined as mean utility scaled by marginal utility of income MU, i.e. $X_i'\beta_i/\mu$.[24]

Here we evaluate utility at perceived characteristics, setting effectiveness and side effects to the average belief in each intention group.[25] Median WTP declines from $69 for the already vaccinated to -69 for those who do not intend to vaccinate, with an overall median of $6.2. WTP for vaccination reflects both the value that participants place on effectiveness as well as the costliness of side effects, discomfort, and inconvenience. Prior work, however, has lacked any measure of the costliness of side effects, and so interpreted WTP as coming *only* from the value of effectiveness, and in particular only from mortality reduction. Under this interpretation, we can back out an implied value of a statistical life:[26]

$$VSL_i = \frac{WTP_i}{VE \cdot Pr(\text{die from flu if unvaccinated})}.$$

[24] The standard way to estimate MU would be by α, the reward sensitivity parameter from the conjoint. However, stated preference data often under estimate price sensitivity, a fact we document in our data. We therefore estimate MU using revealed preference data—wave 2 vaccination rate—combined with our randomly assigned incentives to vaccinate to estimate MU. See Appendix E for details.

[25] We set reward and convenient to 0 and needle to 1, corresponding to prevailing flu vaccines.

[26] This expression also assumes that vaccine effectiveness at preventing death is equal to effectiveness at preventing flu, an assumption we maintain throughout. The conditional death rate is not widely reported but we back it out from the unconditional death rate, the vaccination rate, and vaccine effectiveness using the law of total probability. We assume a 50 percent effectiveness.

In columns (2) and (3) of Table 4, we report median VSL assuming WTP for vaccination is due only to vaccine effectiveness. Overall median VSL is only $145,000, much lower than values commonly used by policy makers or found in the literature.[27]

This apparently low VSL, however, is driven by aversion to side effects and other vaccine characteristics, as well as low perceived flu risk, rather than low value of effectiveness *p*er se. To show this, we calculate WTP for full vaccine effectiveness ($100\beta_{VE}/MU$), and report its median in column (3). Median WTP for effectiveness is $392 overall. This implies a median VSL of $6.39 million, which ranges from $10 million for the most vaccine confident to $2.82 million for the least. These values are much more in line with both policy guidelines and typical estimates from the literature These calculations show the importance of side effects in understanding preferences for flu vaccines. The calculations also show intuitive patterns of risk sorting: people with low VSLs have low demand for flu vaccines.

Explaining vaccine intentions: We use our estimates to decompose the cross-intention-group difference in vaccination rates into beliefs, preferences, and their interactions. For each group, we simulate baseline vaccination probability as

$$\frac{1}{N_g}\sum_{i\in\mathcal{I}_g}\frac{1}{N_d}\sum_{d}Pr(vacc|X_g,\theta_{id}).$$

Here N_g is the number of people in group g, $\mathcal{I}_g$ is the set of people in group g, N_d is the number of draws from the posterior distribution, and $Pr(Vacc|X_g,\beta_{id})$ is the probability of vaccinating given preference vector draw θ_{id}, assuming participant i faces a choice of a single vaccine with characteristics given by their

[27] For example, Banzhaf (2022)'s meta-analysis of meta analyses yields a range of $2-$14 million, and the EPA requires a VSL of $11 million (in 2024) dollars, `https://www.epa.gov/environmental-economics/mortality-risk-valuation`.

group's mean, X_g. [28] Specifically,

$$Pr(Vacc|X_g, \theta_{id}) = \frac{\exp u(\theta_{id}, X_g)}{1 + \exp u(\theta_{id}, X_g)},$$
$$u(\theta_{id}, X_g) = \beta_{0,id} + \beta_{Needle} + \beta_{VE,id} X_{VE,g} + \beta_{SE} X_{SE,g}.$$

This simulation assumes a choice set with two options: no vaccination, or a vaccine with no reward, not convenient, needle delivery, and the effectiveness and side effects profile perceived by their intention group.

We show the simulated vaccination rate in the baseline panel in Figure 5. We see a sensible pattern of predicted vaccination rates: people already vaccinated and intending to vaccinate are predicted to have vaccination rates around 90 percent, while participants who may or may not vaccinate have a predicted rate of 30 percent, and people not intending to vaccinate have a predicted rate of 3 percent.

In the remaining panels, we decompose these large cross-group differences in vaccination rates into differences in preferences, beliefs, and their interactions. In the next panel we change beliefs but hold fixed preferences. To do so, we resimulate vaccination rates, but assuming that all groups have the beliefs (i.e. perceived effectiveness and side effect) of the already vaccinated group, that is, we re-simulate vaccination rates but using X_1 rather than X_g for all groups. Doing so increases vaccination rates by 30 percent for the may-or-may-not group and by 14 for the do-not-intend-group.

In the "change preferences" panel, we change preferences but not beliefs. To do so, we shift each preference parameter distribution to match that of the "already vaccinated" group. In practice we take advantage of our assumption of normality and of common variances. These assumptions imply that we if shift each preference draw θ_{id} from group g to match the posterior mean of the already vaccinated group ($\bar{\theta}_1$), then we will align the preference distributions of the groups.[29] preferences alone produces a larger effect than changing beliefs alone, raising vaccination rates by 28-42 percent for the most hesitant groups.

[28] We use group mean beliefs rather than individual beliefs because the individual level beliefs are measured with error, likely a consequence of both our elicitation method and imperfect numeracy among some participants.

[29] In fact we need to be a bit more careful, because some parameters are log-normally distributed. For these parameters we shift in log space and then exponentiate.

This change in vaccination rates is due to changes in both preferences for side effects and effectiveness, and changes in β_0, a blackbox residual capturing overall dislike of vaccines. However, we show in the next panel that very little of the increase in vaccination from changing preferences is due to changing β_0, despite the fact that β_0 is much lower for the do-not-intend group than the already vaccinated group. Conversely, the next panel shows that changing just preferences for side effects and vaccines, but not other aspects of preferences, has a large effect on vaccination rates, similar to changing beliefs. Changing these preferences and also changing beliefs, shown in the last panel, results in very high vaccination rates across all groups. This combined effect is large because the already vaccinated both have high positive preferences and positive beliefs, and the two interact.

Overall, this decomposition shows that the cross-group difference in vaccination rate is nearly fully explained by differences in beliefs about and preferences for effectiveness and side effects. Other characteristics,such as needle or inconvenience, play at most a small role quantitatively. Likewise scientifically unwarranted beliefs, which are captured by β_0, also play a quantitatively small role.

5.5 Validating the stated preference data

Our analysis so far relies on stated preference data to make inferences about vaccine choices. Because the stated preference tasks are not incentivized and not actual vaccination choices, one might worry that our demand estimates are not informative for actual choices. To address this concern, we provide two kinds of evidence on the validity of the stated preference data: first, by showing that stated preferences are predicted of actual vaccine take-up, and second by showing that they are predictive of incentivized choices (which are not used in estimation).

Predicting vaccination rates: We begin with a standard test of conjoint validity: showing we predict market demand accurately. Table 5 shows that our model predicted vaccination rate is close to the vaccination rate both in our experiment sample and in the 2023 BRFSS, a nationally representative sample. Our experiment sample is not representative in multiple respects (Table 2). Nonetheless, re-weighting the BRFSS to match our sample's age-income-education-sex distribution produces a similar

overall vaccination rate.[30]

Thus at an aggregate level, our model predicts vaccination rates fairly well. Our stated preference data are also predictive of actual vaccination decisions at an individual level. Appendix Table A.2 shows that, even conditional on beliefs, the individual-specific simulated vaccination rate from the stated preference data is highly predictive of actual vaccination rates.

Validity checks from incentivized questions: This evidence shows that stated preferences predict individual vaccination choices. However we use the conjoint to understand individual preference parameters, not just overall take-up. Here we show that stated preferences data on reward sensitivity are informative about reward sensitivity in an incentivized task.

Specifically, in wave 1 after the conjoint, we gave participants a series of incentivized tasks resembling a multiple price list. We told participants they would be invited to a follow-up survey in two weeks. They would get a payment for the follow-up survey, and a bonus payment if they showed proof of vaccination. We asked them to choose between different reward structures, and told them we would pick on of their choices to implement. Participants could pick a payment of $3 regardless of vaccination, or a payment of $2 plus a bonus of $$b$, with $b \in \{1.5, 3, 10, 20, 50\}$. Thus to choose a given bonus b, participants give up $1 of certain payment.[31]

We use the estimated vaccine preferences to predict choices in this incentivized setting. We emphasize that this is an out-of-sample prediction: we do not use the incentivized choices in estimation. To predict incentive take-up, we recognize that our bonus payment scheme gives people an *option* to vaccinate. In

[30] The reason this reweighting does not change the vaccination rate much is that there are two offsetting forces: our sample is younger than average, which pulls down the vaccination rate, but it is better educated and higher income than average, which pulls it up.

[31] To make sure participants understood the incentives, we first asked them about bonus that hypothetical participants would receive in different scenarios, and we correct misunderstanding. We randomized the order of the b and we did not force monotonicity. Nonetheless, of the 1,757 participants who answered the incentive questions, 1,702 of them did not have a preference reversal.

our logit utility framework, the value of the option when the bonus is b and beliefs are X_g is: [32]

$$V(b;\theta,X_g) = 2\alpha + \log\left(1+\exp(u(\theta,X_g)+\alpha b)\right),$$

and the value of the base payment is

$$V_0(\theta,X_g) = 3\alpha + \log\left(1+\exp(u(\theta,X_g))\right).$$

Thus given θ, a person choose the contingent payment if $V(b;\theta,X_g) > V_0(\theta,X_g)$. The probability a particular participant i chooses reward b is

$$P_i(b,X_g) = \frac{1}{N_d}\sum_d 1\left\{\theta 1(V(b;\theta_{id},X_g) > V(0;\theta_{id},X_g)\right\}$$

where $1(\cdot)$ is the indicator function and $G_i(\theta)$ is the posterior distribution of θ for i.

We report the actual and model-predicted take-up of incentives in Figure 6. Averaging across all incentive levels, the model predicts a take-up rate of 45 percent, similar to the observed take-up rate of 40 percent. This overall accuracy, however, masks, import differences across incentive levels: the model over-predicts take-up at low levels and under-predicts at high level. Thus, actual price sensitivity is greater than what is implied by stated preference data.

Nonetheless, individual level price sensitivity parameters (α) are highly predictive of take-up of incentives. To show this, we estimate regressions of the following form for the probability that i chooses the reward arm when the reward is b:

$$Pr(i\text{chooses reward}|b) = \gamma_{0r} + \gamma_{1r}\bar{P}(\theta)_i + \gamma_{2r}\bar{\alpha}_i + \varepsilon_i.$$

Here $P(\theta)_i$ is the posterior mean predicted vaccination rate for i and $\bar{\alpha}_i$ is her posterior mean reward

[32] To derive these equations, recall that expected utility before ϵ is known in a logit model is $\log\left(\sum_j \exp(u_j)\right)$ (Train, 2009). We model the fixed payment of \$2 or \$3 as shifting both inside and outside option by 2α or 3α, and the bonus payment as shifting the inside option only.

sensitivity. If the stated preferences are informative for choices in an incentivized context, we would expect $\gamma_{1r} > 0$ for all r (since the reward option is more appealing the higher is the likelihood of vaccination), and γ_{2r} to increase with r (since reward sensitivity should make high r rewards especially appealing).

This is exactly what we see in the data. Table 6 reports the estimated γs for each reward b.γ_{1r} is large and highly significant for all rewards. γ_{2r}, however, is small and insignificant at low rewards, and increases quickly as the reward rises from $3 to $10 to $20, before leveling off. Thus, individual-level preference parameters in the stated preference data predict individual-level choices in the way we would expect.

6 Belief accuracy

So far we have shown that participants' beliefs about flu risk and vaccine characteristics differ sharply by vaccine intention, and differences in believed vaccine characteristics are important for explaining differences in vaccination rates. Do these belief differences reflect private information, or mistakes? Understanding the difference is important because mistakes but not private information might justify interventions. Here we investigate the accuracy of these beliefs.

6.1 Participants often overestimate flu risk

To investigate biases in beliefs, we compare believed and experienced flu rates, among participants who did not vaccine, in Table 7. We focus on participants who did not vaccinate because our flu risk measure asks about risks conditional on not vaccinating. We measure experienced flu in wave 3, where less than 1 percent of participants perceived a connection to wave 1, so this comparison does not suffer from any desire to be consistent. Our experienced flu measure allows participants to choose a range of responses: did not experience flu, unsure, probably had the flu, or confident/definitely had the flu. To accommodate this range of response, we make a series of indicators for "at least unsure", "at least probably got," and so on.

Table 7 shows that the believed flu rate often exceeds the experienced flu rate. For Participants who intend to vaccinate, the average believed flu rate is 30 percent, 11 percentage points higher than the fraction

of participants reporting that they are at least unsure they got the flu. For Participants with less clear intention to vaccinate, beliefs exceed experienced rates by a smaller margin, but even the "do not intend" group shows a 12 percentage point gap between the believed flu rate and the fraction of people who are confident they got the flu.

Despite the gaps between actual believed and experienced flu rates, beliefs contain private information, in the sense that they are useful for forecasting individual flu gap. The gap between actual and experienced flu risk shows that this private information is imperfect, or imperfectly used. To show the private information in beliefs, we regress actual flu on expected flu, as well as indicators for vaccine intention groups. After correcting for measurement error, we find a coefficient on believed risk of 0.26 (se: 0.08).[33]

6.2 Participants often overestimate vaccine side effects

While Participants overestimate flu risk, a bias them towards vaccination, they also overestimate side effects of vaccination, a biasing away from vaccination. As with bias in flu risk, our approach to measuring bias in beliefs about side effects of vaccination is to compare beliefs with actual experiences. This approach is challenging for side effects, however, because the side effect rate of a vaccine is a treatment effect, i.e. the impact of the vaccine on the experience of severe side effect symptoms. Overcoming this challenge requires exogenous variation in vaccination rates, for the same sample for which we elicit side effect beliefs.

Our experiment provides this variation via randomized incentives to vaccinate. Recall that the third component of wave 1 was a series of incentivized task where participants chose between a certain follow-up payment of $3, or a certain payment of $2, plus an additional reward of $b if they showed they were vaccinated in the follow-up survey. Participants made 5 choices, corresponding to b = $1.50, $3, $10, $20, $50. We incentivized this task by picking one of the five choices at random to implement. This random

[33] In this regression we measure experienced flu with an indicator for "probably or confident had the flu." To correct for measurement error, we instrument for believed flu risk with our likert scale responses for believed flu risk, e.g., "would definitely get the flu," "would probably get the flu", and so on.

incentive provision also creates exogenous variation in the incentive to vaccinate.

The effect of our incentives on vaccination varies depending on baseline vaccine intentions and on responses to the incentive question. If a person always selects the certain payment of $3, then we never vary their incentive to vaccinate. If a person always selects the bonus payment, even the low bonus of $1.50, then it is likely that they will vaccinate regardless of the randomly selected incentive. Thus our pre-analysis plan specified that we look at people who switch from certain payment to bonus at $10 or $20. These switchers are marginally vaccine hesitant: their choices indicate that they do not intend to vaccinate without a reward, but they have a chance of vaccinating at a reward of $10 or $20.[34]

To estimate the effect of vaccination on experienced side effects for this sample, we define "incentive" as offering an incentive at the switching point or greater, and we estimate the following regression using these switchers:

$$y_i = \gamma_0 + \gamma_1 incentive_i + \epsilon_i.$$

We estimate separate regressions for the $10 switchers and $20 switchers, as well as pooled regressions.[35] Our outcomes are wave 2 vaccination and wave 2 experience of severe side effect symptoms. We report the wave 1 expected side effect rate, for each group.

Table 8 shows the results. The first row shows that the incentives increased self-reported vaccinations, raising them by 11 percentage points for the $10 switchers, and by (a statistically insignificant) 6 percentage points for the $20 switchers. The next row show that these induced vaccinations did not increase experienced side effect symptoms—in fact they reduced side effects symptoms by 18-25 percentage points. This reduction is statistically significantly different from both zero and from the switchers' ex ante believed side effect rate, of 3-10 percent. Thus these marginally vaccine hesitant participant overestimate the side effect consequences of vaccination: they expect vaccination to increase side effects and, if anything, it re-

[34] In fact we found that people selecting a bonus of $1.50 were not especially likely to vaccinate at that level, but did respond to the incentive, so we also estimate specifications with the fuller set of participants who were randomized to incentives.

[35] When we pool both samples, we include an indicator for $20 switchers. This is important because our design generates a 50 percent chance of incentives for $10 switchers, but only a 25 percent chance for $20 switchers.

duces side effects. This overestimation is consequential for vaccine decisions. Appendix Table A.5 shows simulated vaccination rates by intention group if side effect rates were zero. For the marginally vaccine hesitant—roughly the compliers in our experiment—eliminating perceived side effects would increase vaccination rates by 20 percentage points.[36]

We note that the reduction in experienced symptoms is very large, especially in light of the smaller "first stage" increase in vaccination. We do not have an explanation for this large decrease, but we can rule out a number of important concerns. First, we emphasize that we pre-specified this approach. Second, it is unlikely that experimenter demand or question order effects explains this finding. We asked about experienced side effects before asking about vaccination status, and we asked about expected side effects two weeks before asking about experienced side effects. Third, we worried that differential survey timing might explain the difference in experienced symptoms: perhaps our incentivized participants took the follow-up survey sooner than the incentivized participants, exposing them to different seasonal illness patterns. However nearly all participants took the follow-up survey on the first day it was available, and including fixed effects for response day does not change the estimated effect (Appendix Table A.6, columns 1 and 2). We also find similar results—significantly negative effects of incentives to vaccinate on experienced side effects—when we expand the sample to include anyone with a randomly assigned incentive to vaccinate (Appendix Table A.6, columns 3 and 4). We conclude that the reduction in experienced symptoms, though large, is likely not entirely a statistical or experiment design artifact, and instead provides evidence that marginally vaccine hesitant people overestimate side effect consequences of vaccines.

7 Follow-through on vaccine intentions

Our analysis so far decomposes differences in vaccine intentions using stated preference data and beliefs. Here we investigate how these intentions translate into action. Our main findings are that in the short run

[36] Eliminating side effects raises vaccination by only 3 percentage points for people not intending to vaccinate, because few of them are close to the margin of vaccinating. It also reduces vaccination rates slightly for people intending to vaccinate, because they perceive negative side effect rates (rather than zero).

intentions imperfectly predict actions and nudges help align the two, but in the longer run expectations are closer to actions and nudges less effective.

To more precisely measure intentions, we asked participants their expectation of future vaccination at the end of our initial survey. Recall that we conducted an incentivized task among people not already vaccinated, and we randomized incentives to vaccinate. After revealed the chosen incentives, we asked participants whether they intended to vaccinate before the follow-up survey (and hence get the incentive), We asked a qualitative version of this question (answers: yes, unsure, no), and a quantitative version asking about the probability. After asking this question, we randomized a subset[37] of participants to see a nudge encouraging them to sign up for a vaccination appointment at a pharmacy. We provided links to CVS, Walgreens, Walmart, and a CDC website where they could make an appointment.

7.1 People expect to vaccinate more and sooner than they actually do

Our first result is that vaccine expectations outpace actual vaccinations. The first column of Table 9 reports the fraction of participants who say they will vaccinate by the follow-up survey, and the next reports the average self-reported probability of vaccinating by the follow-up survey. Here we limit the sample to the 82 percent of participants who completed both the wave 2 and wave 3 follow-ups, so that we measure intentions and follow through for the same set of people. Roughly 60 percent of participants intending to vaccinate expect to vaccinate by the wave 2 follow up; this fraction falls to 18-34 percent for participants saying they may or may not vaccinate, and 8-10 percent for participants who do not intend to vaccinate.

We measure actual vaccination in our wave 2 and wave 3 follow ups, which occurred two weeks and two months later. The two-week vaccination rate is 39 percent among those intending to vaccinate, and falling to 7 and 3 percent among the more hesitant. For each intention group, expectations are statistically significantly higher than actions at the two week mark. When we look to the two-month mark, however, the gap between actions and intentions closes. For people intending to vaccinate, the actual vaccinate rate

[37] We only randomized participants who chose at least one of the incentives less than $50, and therefore demonstrated some interest in vaccination.

is 55 percent, still lower than the expected probability, but close to the fraction saying they will vaccinate. Similarly for the "may or may not" group, the fraction actually vaccinating is similar to the fraction saying they will vaccinate, but only half the size of the expected vaccination rate. For people not intending to vaccinate, rates remain low and below expectation.

7.2 Nudges accelerate vaccinations for people who expect to vaccinate

Our nudge intervention–encouragement to sign up for a vaccination appointment–is a potential tool to close the gap between intentions and actions, a common motivation for nudges (Thaler and Sunstein, 2021). Such an encouragement might help people who are in principle open to vaccination, but have not yet made a concrete plan or who generally struggle to follow through on plans. Text messages encouraging vaccination sign-ups, or reminding about them, are a common and effective nudge (Milkman et al., 2022; Patel et al., 2023). Our intervention is similar. We differ from prior work in that we elicit subjective expectations of vaccination rates before nudging. This is useful for seeing who responds to the nudge.

We report the effect of this nudge in Table 10. The nudge had a large effect in the short run, increasing wave 2 vaccination rates by 8 percentage points (control mean: 13 percent). This large overall effect is concentrated among people who say they expect to vaccinate, for whom we find an effect of 11 percentage points (control mean: 37 percent). For people they are unsure or do not expect to vaccinate, the effect size is an insignificant 3 percentage points. This large effect on wave 2 vaccination rates entirely dissipates by wave 3, when the effect of the nudge falls to zero for all groups. Although the wave 3 estimates are noisy, we can reject the null hypothesis that the wave 2 and 3 effects are the same overall (p=.02) or for the "will vaccinate" group (p=.03). Thus the nudge *accelerates* vaccination but does not *increase* it for our sample.

8 Conclusions

The CDC recommends that all Americans aged 6 months and older get a flu vaccine, yet fewer than half take-up this recommendation. We have shown that such vaccine hesitancy is primarily due to differences in beliefs about and preferences over vaccine effectiveness and side effects. Vaccine hesitancy is not sub-

stantially explained by needlephobia or extreme beliefs. Both vaccine confident and hesitant participants hold plausible beliefs, in the sense that their beliefs about national flu risk and vaccine effectiveness, and their beliefs about the risk and effectiveness they themselves face, are roughly in line with national averages. Beliefs are not perfectly accurate, however. Most participants overestimate their flu risk, especially vaccine confident. Vaccine hesitant people overestimate the short term side effects of vaccination, and our estimates suggest that correcting this misperception would substantially increase their vaccination rates. Finally, consistent with prior work (e.g., Bronchetti et al. (2015)), we document a substantial gap between intention and action. Among the large share of participants indicating an intention to vaccinate soon, relatively few actually did so, but nudges to increase vaccination are effective in the short term among this group.

We emphasize multiple limitations of our study. Most critically, we focus on flu. Vaccine hesitancy is a broader phenomenon applying to other seasonal diseases—prominently, COVID-19—as well as pediatric vaccines. It is unclear whether our results apply more broadly to these diseases, as different sets of issues may govern vaccine hesitancy in other contexts, such as concerns about child development or newly developed mRNA vaccines. Second, our sample is drawn from participants on an online survey platform. Although they appear similar to the US in flu vaccination rates after reweighting for age, education, and gender, they may be nonrepresentative in terms of views about vaccines or flu risk. Finally, much of our results rely on stated preferences and beliefs. We validate these data with revealed preference questions and follow-up outcome measures, but not all measures can be validated.

Despite these limitations, our results have several implications for vaccine policy and research. On the policy side, our results first suggest that vaccine nudges are welfare-improving. As Allcott et al. (2022) explain, nudges are welfare improving when they correct an internality among the nudged. Our results show that nudges change behavior among people who indicate an intention to vaccinate. Viewing this failure-to-follow-through as a negative internality, our results imply that nudges are well-targeted. Second, our results suggest that information provision interventions must be carefully designed or targeted. As beliefs about flu risk and vaccine effectiveness are in line with objective estimates (or biased in a way

that favors vaccinations), and extreme beliefs play a quantitatively minor role, many forms of information provision appear unlikely to help. However, providing clearer information about short-term side effects may help, as might targeting information to groups with the most pessimistic beliefs. Designing, targeting, and evaluating such information interventions would be a valuable direction for future research.

References

Abrevaya, J. and K. Mulligan (2011). Effectiveness of state-level vaccination mandates: evidence from the varicella vaccine. *Journal of health economics 30*(5), 966–976.

Ahmad, F. B. (2024). Mortality in the united states—provisional data, 2023. *MMWR. Morbidity and Mortality Weekly Report 73*.

Akel, K. B., N. B. Masters, S.-F. Shih, Y. Lu, and A. L. Wagner (2021). Modification of a vaccine hesitancy scale for use in adult vaccinations in the united states and china. *Human vaccines & immunotherapeutics 17*(8), 2639–2646.

Allcott, H., D. Cohen, W. Morrison, and D. Taubinsky (2022). When do" nudges" increase welfare? Technical report, National Bureau of Economic Research.

Allenby, G. M., N. Hardt, and P. E. Rossi (2019). Economic foundations of conjoint analysis. In *Handbook of the Economics of Marketing*, Volume 1, pp. 151–192. Elsevier.

Alsan, M. and S. Eichmeyer (2024). Experimental evidence on the effectiveness of nonexperts for improving vaccine demand. *American Economic Journal: Economic Policy 16*(1), 394–414.

Athey, S., K. Grabarz, M. Luca, and N. Wernerfelt (2023). Digital public health interventions at scale: The impact of social media advertising on beliefs and outcomes related to covid vaccines. *Proceedings of the National Academy of Sciences 120*(5), e2208110120.

Banerjee, A. V., E. Duflo, R. Glennerster, and D. Kothari (2010). Improving immunisation coverage in rural india: clustered randomised controlled evaluation of immunisation campaigns with and without incentives. *Bmj 340*.

Banzhaf, H. S. (2022). The value of statistical life: A meta-analysis of meta-analyses. *Journal of Benefit-Cost Analysis 13*(2), 182–197.

Ben-Akiva, M., D. McFadden, K. Train, et al. (2019). Foundations of stated preference elicitation: Consumer behavior and choice-based conjoint analysis. *Foundations and Trends® in Econometrics 10*(1-2), 1–144.

Bronchetti, E. T., D. B. Huffman, and E. Magenheim (2015). Attention, intentions, and follow-through in preventive health behavior: Field experimental evidence on flu vaccination. *Journal of Economic Behavior & Organization 116*, 270–291.

Campos-Mercade, P., A. N. Meier, S. Meier, D. G. Pope, F. H. Schneider, and E. Wengström (2024). Incentives to vaccinate. Technical report, National Bureau of Economic Research.

Carlin, P., B. E. Dixon, K. I. Simon, R. Sullivan, and C. Wing (2022). How undervalued is the covid-19 vaccine? evidence from discrete choice experiments and vsl benchmarks. Technical report, National Bureau of Economic Research.

Carpenter, C. S. and E. C. Lawler (2019). Direct and spillover effects of middle school vaccination requirements. *American Economic Journal: Economic Policy 11*(1), 95–125.

Center for Disease Control (2025a). Flu diease burden: Past seasons. Accessed: 2025-03-10.

Center for Disease Control (2025b). How cdc estimates the burden of seasonal flu in the united states. Accessed: 2025-03-10.

Center for Disease Control (2025c). Weekly influenza vaccination coverage and intent for vaccination, overall, by selected demographics and jurisdiction, among adults 18 years and older (dataset). Accessed: 2025-03-10.

Centers for Disease Control and Prevention (2024a). About estimated flu burden. Accessed: 2025-02-17.

Centers for Disease Control and Prevention (2024b). Seasonal flu vaccine effectiveness studies. Accessed: 2025-02-17.

Chan, A. (2022). Discrimination against doctors: A field experiment. *Unpublished manuscript*.

Churchill, B. F., L. E. Henkhaus, and E. C. Lawler (2024). Effect of vaccine recommendations on consumer and firm behavior. *Journal of Policy Analysis and Management*.

Diamond, P. A. and J. A. Hausman (1994). Contingent valuation: is some number better than no number? *Journal of economic perspectives 8*(4), 45–64.

Duflo, E. and A. Banerjee (2011). *Poor economics*, Volume 619. PublicAffairs New York.

Freedman, S. M., D. W. Sacks, K. I. Simon, and C. Wing (2022). Direct and indirect effects of vaccines: Evidence from covid-19 in schools.

Freeman, D., S. Lambe, L.-M. Yu, J. Freeman, A. Chadwick, C. Vaccari, F. Waite, L. Rosebrock, A. Petit, S. Vanderslott, et al. (2023). Injection fears and covid-19 vaccine hesitancy. *Psychological medicine 53*(4), 1185–1195.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). Five facts about beliefs and portfolios. *American Economic Review 111*(5), 1481–1522.

Goodkin-Gold, M., M. Kremer, C. M. Snyder, and H. L. Williams (2020). Optimal vaccine subsidies for endemic and epidemic diseases. Technical report, National Bureau of Economic Research.

Green, J. (2025, January 7). I'm the governor of hawaii. i've seen what vaccine skepticism can do. Accessed: February 4, 2025.

Haaland, I., C. Roth, and J. Wohlfart (2023). Designing information provision experiments. *Journal of economic literature 61*(1), 3–40.

Hausman, J. (2012). Contingent valuation: from dubious to hopeless. *Journal of economic perspectives 26*(4), 43–56.

Ho, L., E. Breza, A. Banerjee, A. G. Chandrasekhar, F. C. Stanford, R. Fior, P. Goldsmith-Pinkham, K. Holland, E. Hoppe, L.-M. Jean, et al. (2023). The impact of large-scale social media advertising campaigns on covid-19 vaccination: Evidence from two randomized controlled trials. In *AEA Papers and Proceedings*, Volume 113, pp. 653–658. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Horne, Z., D. Powell, J. E. Hummel, and K. J. Holyoak (2015). Countering antivaccination attitudes. *Proceedings of the National Academy of Sciences 112*(33), 10321–10324.

Humlum, M. K., M. O. Morthorst, and P. R. Thingholm (2024). Sibling spillovers and the choice to get vaccinated: Evidence from a regression discontinuity design. *Journal of Health Economics 94*, 102843.

Kessler, J. B., C. Low, and C. D. Sullivan (2019). Incentivized resume rating: Eliciting employer preferences without deception. *American Economic Review 109*(11), 3713–3744.

Kriss, J. L. (2024). Influenza, covid-19, and respiratory syncytial virus vaccination coverage among adults—united states, fall 2024. *MMWR. Morbidity and Mortality Weekly Report 73*.

Larsen, B. J., T. J. Ryan, S. Greene, M. J. Hetherington, R. Maxwell, and S. Tadelis (2023). Counter-stereotypical messaging and partisan cues: Moving the needle on vaccines in a polarized united states. *Science advances 9*(29), eadg9434.

Larson, H. J., C. Jarrett, W. S. Schulz, M. Chaudhuri, Y. Zhou, E. Dube, M. Schuster, N. E. MacDonald, R. Wilson, et al. (2015). Measuring vaccine hesitancy: the development of a survey tool. *Vaccine 33*(34), 4165–4175.

Lawler, E. C. (2017). Effectiveness of vaccination recommendations versus mandates: Evidence from the hepatitis a vaccine. *Journal of health economics 52*, 45–62.

MacDonald, N. E. et al. (2015). Vaccine hesitancy: Definition, scope and determinants. *Vaccine 33*(34), 4161–4164.

Manski, C. F. (2004). Measuring expectations. *Econometrica 72*(5), 1329–1376.

Milkman, K. L., L. Gandhi, M. S. Patel, H. N. Graci, D. M. Gromet, H. Ho, J. S. Kay, T. W. Lee, J. Rothschild, J. E. Bogard, et al. (2022). A 680,000-person megastudy of nudges to encourage vaccination in pharmacies. *Proceedings of the National Academy of Sciences 119*(6), e2115126119.

Moshary, S., B. Shapiro, and S. Drango (2023). Preferences for firearms and their implications for regulation. Technical report, National Bureau of Economic Research.

Naquin, A. (2024). Laboratory-confirmed influenza-associated hospitalizations among children and adults—influenza hospitalization surveillance network, united states, 2010–2023. *MMWR. Surveillance Summaries 73*.

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

Paget, J., P. Spreeuwenberg, V. Charu, R. J. Taylor, A. D. Iuliano, J. Bresee, L. Simonsen, C. Viboud, et al. (2019). Global mortality associated with seasonal influenza epidemics: New burden estimates and predictors from the glamor project. *Journal of global health 9*(2).

Patel, M. S., K. L. Milkman, L. Gandhi, H. N. Graci, D. Gromet, H. Ho, J. S. Kay, T. W. Lee, J. Rothschild, M. Akinola, et al. (2023). A randomized trial of behavioral nudges delivered through text messages to increase influenza vaccination among patients with an upcoming primary care visit. *American Journal of Health Promotion 37*(3), 324–332.

Philipson, T. (2000). Economic epidemiology and infectious diseases. *Handbook of health economics 1*, 1761–1799.

Reed, C., S. S. Chaves, P. Daily Kirley, R. Emerson, D. Aragon, E. B. Hancock, L. Butler, J. Baumbach, G. Hollick, N. M. Bennett, et al. (2015). Estimating influenza disease burden from population-based surveillance data in the united states. *PloS one 10*(3), e0118369.

Rolfes, M. A., B. Flannery, J. R. Chung, A. O'Halloran, S. Garg, E. A. Belongia, M. Gaglani, R. K. Zimmerman, M. L. Jackson, A. S. Monto, et al. (2019). Effects of influenza vaccination in the united states during the 2017–2018 influenza season. *Clinical Infectious Diseases 69*(11), 1845–1853.

Rolfes, M. A., I. M. Foppa, S. Garg, B. Flannery, L. Brammer, J. A. Singleton, E. Burns, D. Jernigan, S. J. Olsen, J. Bresee, et al. (2018). Annual estimates of the burden of seasonal influenza in the united states: a tool for strengthening influenza surveillance and preparedness. *Influenza and other respiratory viruses 12*(1), 132–137.

Rossi, P. E., G. M. Allenby, and S. Misra (2024). *Bayesian Statistics and Marketing* (2nd ed.). Wiley.

Schneider, F. H., P. Campos-Mercade, S. Meier, D. Pope, E. Wengström, and A. N. Meier (2023). Financial incentives for vaccination do not have negative unintended consequences. *Nature 613*(7944), 526–533.

Smith, J., A. Whalley, and N. Wilcox (2021). *Are Participants Good Evaluators?* WE Upjohn Institute.

Stantcheva, S. (2023). How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics 15*(1), 205–234.

Thaler, R. H. and C. R. Sunstein (2021). *Nudge: The final edition*. Yale University Press.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

U.S. Food and Drug Administration (2024). Fda approves nasal spray influenza vaccine for self- or caregiver-administration. Accessed February 4, 2025.

Ward, C. J. (2014). Influenza vaccination campaigns: is an ounce of prevention worth a pound of cure? *American Economic Journal: Applied Economics 6*(1), 38–72.

White, C. (2021). Measuring social and externality benefits of influenza vaccination. *Journal of Human Resources 56*(3), 749–785.

World Health Organization (2025). Influenza (seasonal). Accessed: 2025-03-10.

Table 1: Sample sizes

Inclusion criteria	Count
Start survey	3,447
Pass attention checks	2,660
Consistent stated preferences	2,572
Only one attempt	2,569
Not missing key variables	2,567
No contraindications (Main sample)	2,357
Not already vaccinated (eligible for incentives and waves 2/3)	1,757
In wave 2	1,681
In wave 3	1,449

Notes: Table reports the sample size as we impose our inclusion criteria.

Table 2: Summary statistics

Characteristic	Analysis sample	2023 BRFSS
Age 18-39	0.51	0.37
Age 40-54	0.29	0.23
Age 55-64	0.09	0.07
Age 65+	0.10	0.32
Female	0.52	0.52
Less than HS	0.00	0.12
High school	0.11	0.27
Some college	0.27	0.30
College or more	0.61	0.31
Income < $25,000	0.12	0.13
Income $25,000 - $50,000	0.20	0.20
Income $50,000 - $100,00	0.36	0.23
Income > $100,00	0.30	0.45
Already vaccinated	0.25	.
Intend to vaccinate	0.24	.
May or may not vaccinate	0.21	.
Do not intend to vaccinate	0.30	.
Flu vaccination	0.46	0.42

Notes: Each cell is the mean of an indicator variable. BRFSS does not ask about vaccine intentions but asks whether subjects received a flu shot in the last 12 months. Our flu vaccination measure is an indicator for either "already vaccinated" or "reports vaccination in wave 3," and is set to missing for people not in wave 3 and not already vaccinated..

Table 3: Mean preference parameters by vaccine intention group

Sensitivity to:	Reward	Vaccine effectiveness	Side effects	Needle	Convenient	Inside option
Already vaccinated	0.103	0.647	-1.14	0.929	-0.0354	-18.1
Intend to vaccinate	0.106	0.486	-1.12	0.814	0.119	-15.7
May or may not vaccinate	0.142	0.343	-1.59	0.0917	0.0881	-15.8
Do not intend to vaccinate	0.178	0.215	-1.57	-0.591	0.0544	-20.8

Notes: Table reports the posterior mean of each preference parameter, by vaccine intention group. The parameters are the coefficients on the indicated characteristics.

Table 4: Low willingness-to-pay for vaccination reflects distaste for side effects not low taste of effectiveness

VSL based on:	WTP for vaccine		WTP for effectiveness	
	WTP ($s) (1)	VSL-objective ($millions) (2)	WTP ($s) (3)	VSL-objective ($millions) (4)
Already vaccianted	248	5.22	867	9.8
Intent to vaccinate	155	5.07	621	10.4
May or may not vaccinate	-36.8	-1.59	324	7.26
Do not intend to vaccinate	-248	-8.7	153	2.82
All	6.2	0.145	392	6.39

Notes: Table reports median VSL estimates derived from vaccine preferences. In columns (1)-(2) we assume that the only aspect of vaccine preference is effectiveness. Column (1) reports the WTP for vaccination, and columns (2) turns this WTP estimate into a VSL estimate using age-specific flu death risks. Columns (3)-(4) repeat the exercise but only using willingness to pay for VE, rather than for the vaccine as a whole.

Table 5: Stated preference data predict overall vaccination rates

	Vaccination rate
Model prediction	0.45
Experiment sample	0.42
BRFSS	0.42
BRFSS, reweight age	0.36
BRFSS, reweight age-income-education-sex	0.41

Notes: The first row of the table reports the predicted vaccination rate from the stated preference data, adjusting for subjects who reported they faced contraindications to vaccinate (who were excluded from estimation). The next row reports the actual vaccination rate in the experiment sample. The next three rows report vaccination rates from the BRFSS, overall, reweighting to match the samples' age distribution, and reweighting to make the samples distribution of age, income, education, and sex.

Table 6: Stated preferences predict incentivized choices

Reward	$1.50	$3	$10	$20	$50
Pr(takeup)	0.37	0.48	0.59	0.55	0.46
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Reward beta	0.01	0.04	0.19	0.40	0.38
	(0.04)	(0.05)	(0.07)	(0.07)	(0.07)
N	1,757	1,757	1,757	1,757	1,757

Notes: Each column is a separate regression. The dependent variable is an indicator for choosing the vaccine-contingent reward payment instead of an extra $1 in non-contingent payment. The independent variables are shown: Pr(takeup) is the posterior mean predicted vaccination rate, and α is the posterior mean reward sensitivity, both derived from stated preferences.

Table 7: Believed flu rate exceeds experienced flu rate

	Believed	Experienced flu rate, less believed rate		
	flu rate	Unsure+	Probably+	Confident
Intend to	0.30	-0.11	-0.15	-0.26
N =213		(0.03)	(0.03)	(0.02)
May get, may not	0.23	0.02	-0.10	-0.19
N =322		(0.03)	(0.02)	(0.01)
Do not intend	0.16	0.03	-0.04	-0.12
N =557		(0.02)	(0.02)	(0.01)

Notes: Column 1 reports the wave 1 believed rate of getting the flu if unvaccinated. Columns 2-4 report the difference between experienced flu and believed flu, for three measures of believed flu, ascertained in wave 3: indicators for at least unsure of getting the flu (col 2), at least probably got the flu (col 3), or confident they got the flu (col 4). The sample is always limited to people who were not vaccinated in wave 3, because we elicited flu risk conditional on not vaccination. Standard errors of the differences in means (between belief and experience) in parentheses.

Table 8: Believed vaccine side effect rate exceeds experienced side effect rate

Group	Switch at $10	Switch at $20	Pool
Y = vaccinated	0.11	0.06	0.10
	(0.04)	(0.06)	(0.04)
Y = experience severe symptoms	-0.18	-0.25	-0.20
	(0.06)	(0.09)	(0.05)
Mean subjective SE rate	0.03	0.10	0.05

Notes: Table reports the effect of an incentive to vaccinate on the indicated outcome (top two rows, robust standard errors in parentheses), and the mean believed side effect rate for the estimation sample (bottom row). The first column consists of people who preferred a $10 incentive to vaccinate to $1 for sure (and did prefer not a $1.50 incentive) and the second column consists of people who preferred a $20 incentive to vaccinate to $1 for sure (and did not prefer a $10 incentive). The incentive effect is the effect of being offered the preferred incentive, relative to $1 for sure. We regress the indicated outcome on a group-specific indicator for offered incentives, plus (in the third column) an indicator for the group membership.

Table 9: Expectations to vaccinate exceed actual vaccinations

	Expect to vaccinate by wave 2?			Actual vaccination	
	Unsure	Yes	Probability	Wave 2	Wave 3
Intend to vaccinate	0.35	0.58	0.63	0.39	0.55
May or may not vaccinate	0.50	0.19	0.34	0.08	0.18
Do not intend to vaccinate	0.13	0.07	0.11	0.03	0.04

Notes: Table reports wave 1 expectations of vaccination with self-reported actual vaccination in waves 2 and 3, by wave 1 vaccine intention group. Each cell in the table is a mean. The sample is limited to respondents who completed waves 2 and 3. The first two columns report the fraction of respondents who say "unsure" or "yes" when asked if they will vaccinate by wave 2. The third column is the self-reported probability of vaccinating over this period. The final two columns show the actual vaccination rate in waves 2 and 3, for the same respondents.

Table 10: A sign-up nudge accelerated but did not increase vaccinations for people expecting to vaccinate

Outcome:	Wave 2 vaccination	Wave 3 vaccination
Full sample		
Nudge effect	0.083	0.001
Std. Error	(0.031)	(0.035)
N	785	779
Sample saying they will vaccinate		
Nudge effect	0.112	-0.014
Std. Error	(0.051)	(0.051)
N	379	380
Sample not saying they will vaccinate		
Nudge effect	0.037	-0.003
Std. Error	(0.030)	(0.041)
N	406	399

Notes: Table reports the effect of a nudge, an encouragement to sign-up for vaccinations on an indicator for vaccinating by wave 2 survey or wave 3 survey. The first set panel consists of all respondents eligible for the nudge (which was randomly assigned in this group). The next panel is limited to respondents who said "yes" when asked if they expected to vaccinate (prior to the nudge), and the final panel is the remaining respondents. Robust standard errors in parentheses.

Figure 1: Experiment design

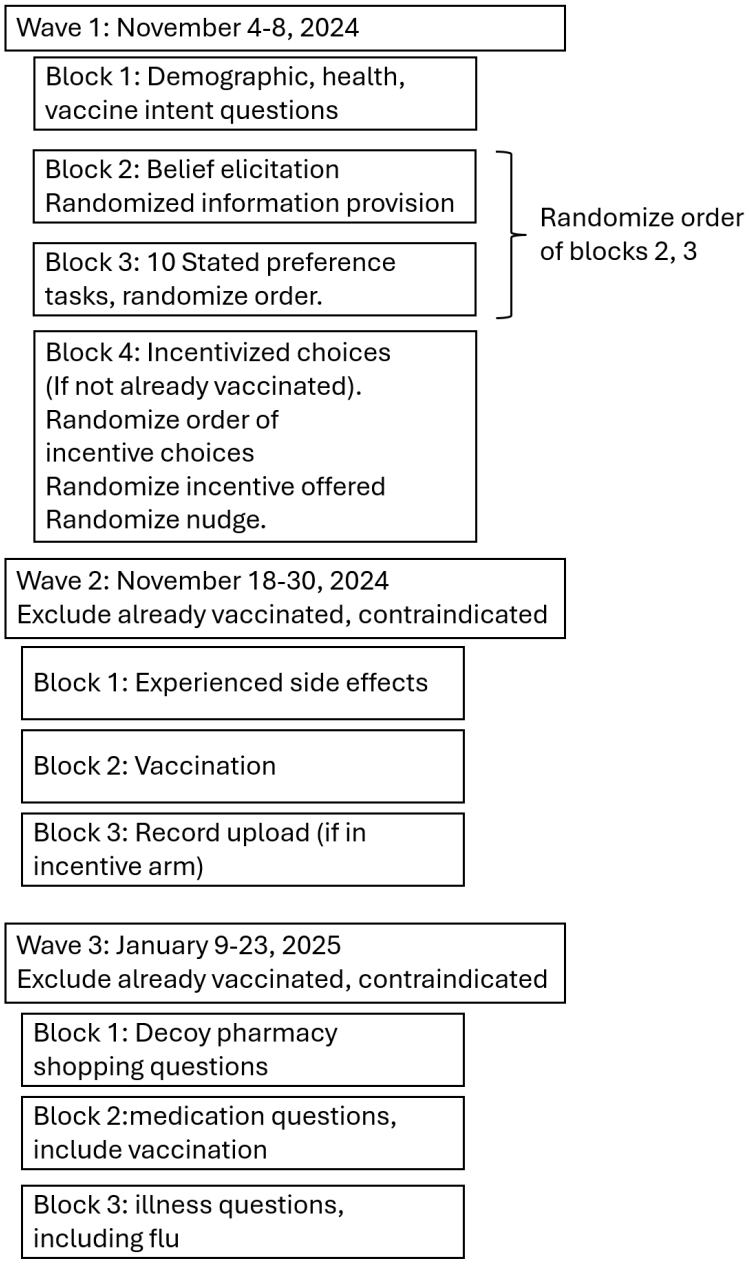
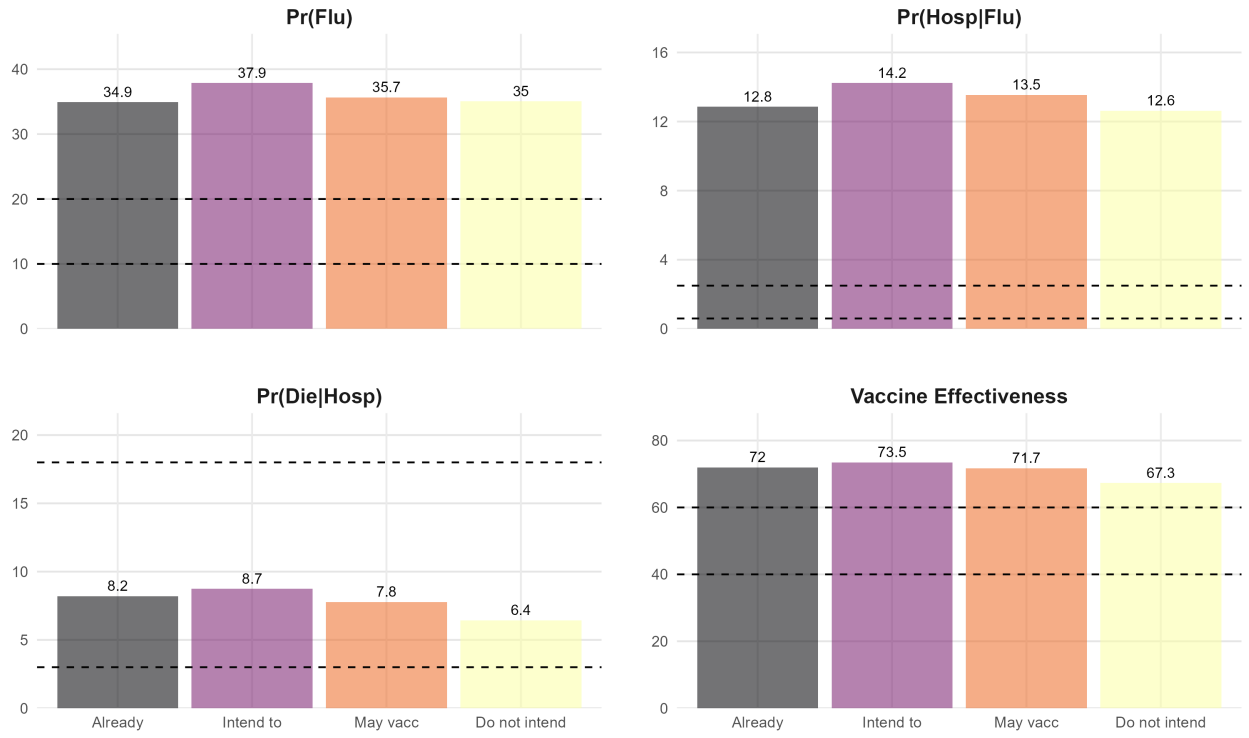
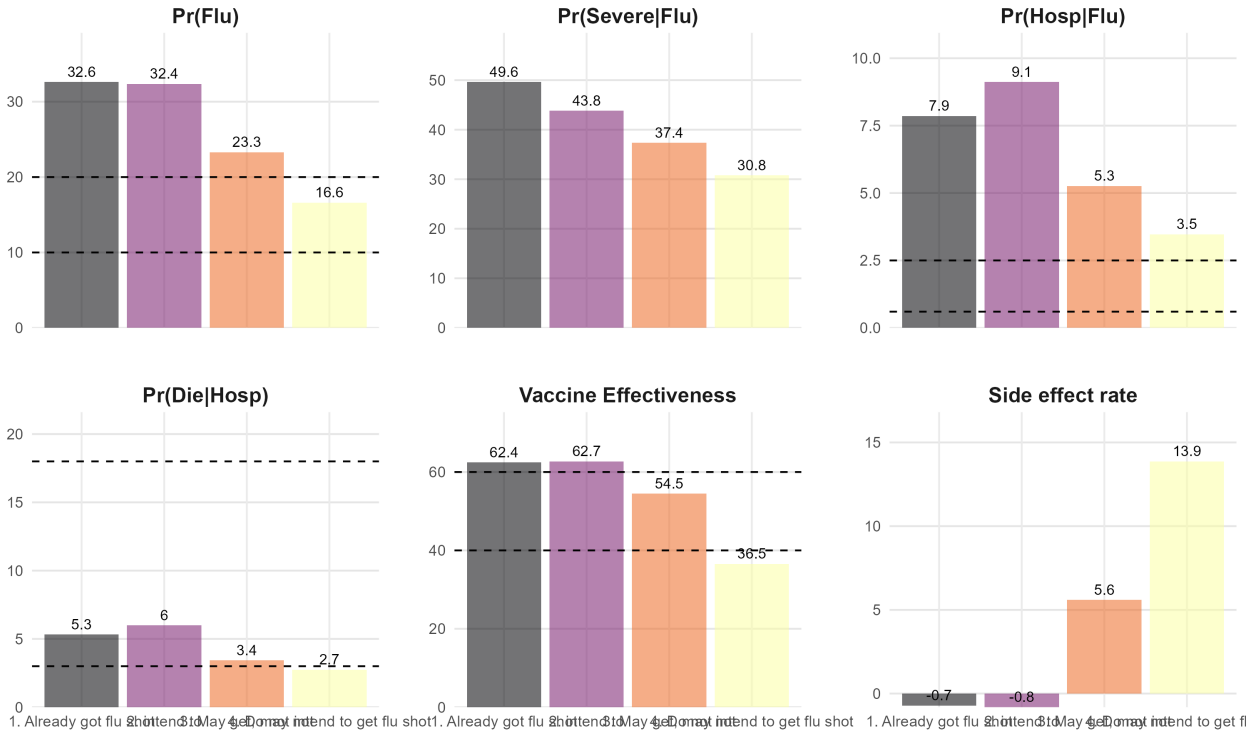


Figure 2: Beliefs about national averages



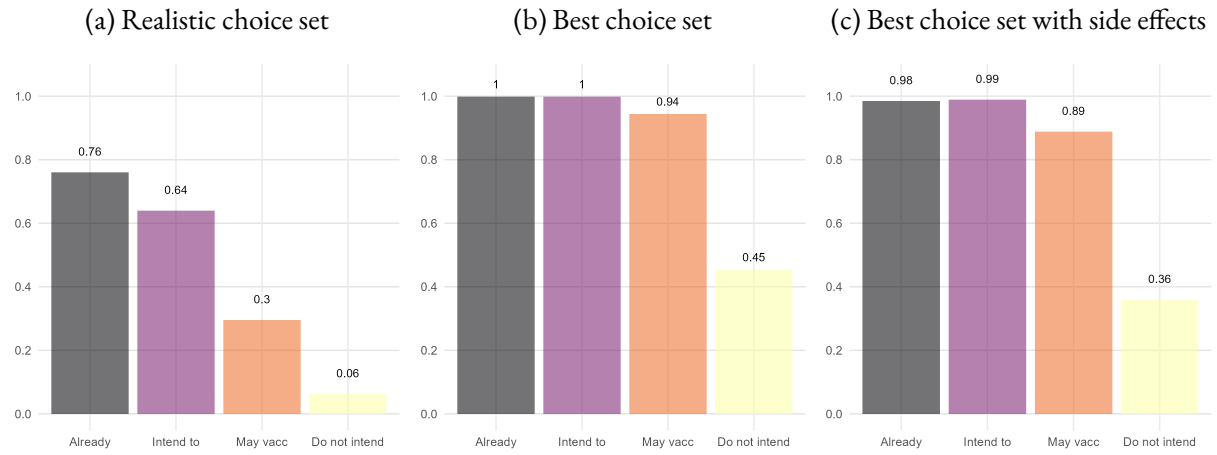
Notes: Figure reports the average belief about the national average value of the indicated object, for each vaccine intention group. The dashed horizontal lines are the values reported by the CDC.

Figure 3: Beliefs about personal risks and vaccine effects



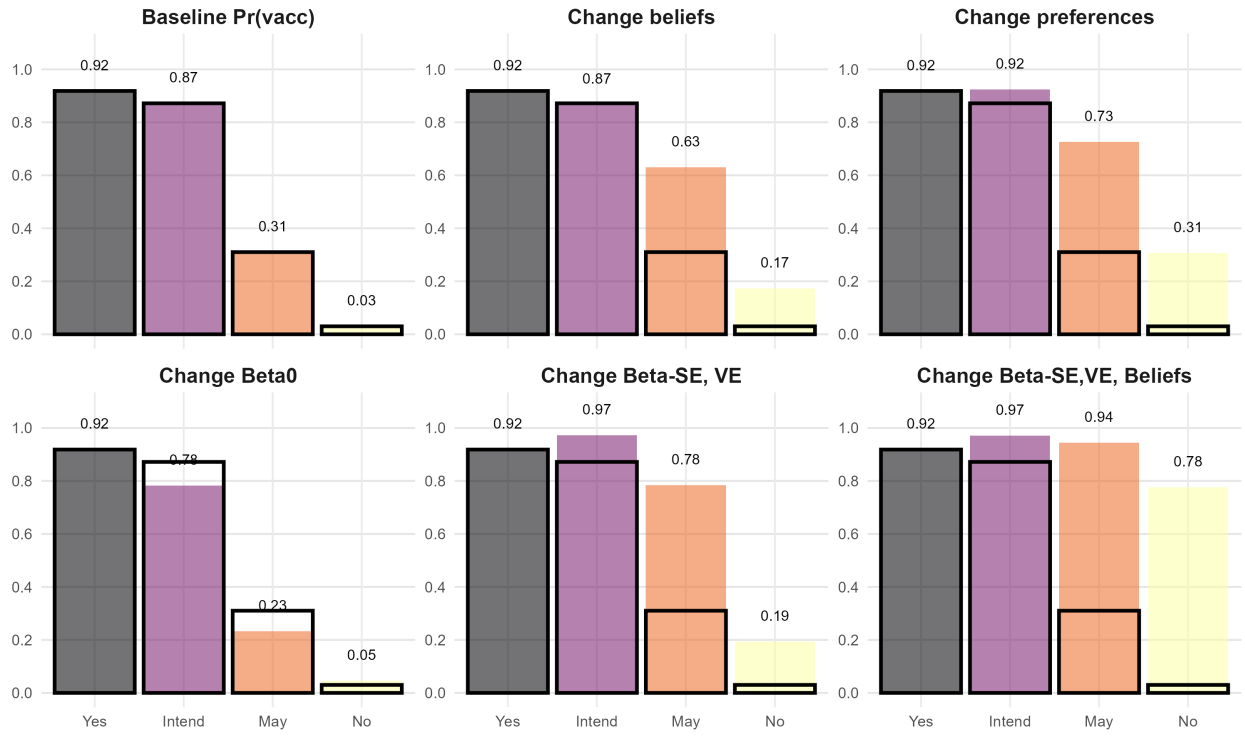
Notes: Figure reports the average belief about the person-specific values of the indicated object, for each vaccine intention group. The dashed horizontal lines are the values reported by the CDC. The CDC does not report values for severe flu or side effects.

Figure 4: Vaccine hesitant report willingness to vaccinate, for sufficiently appealing choices



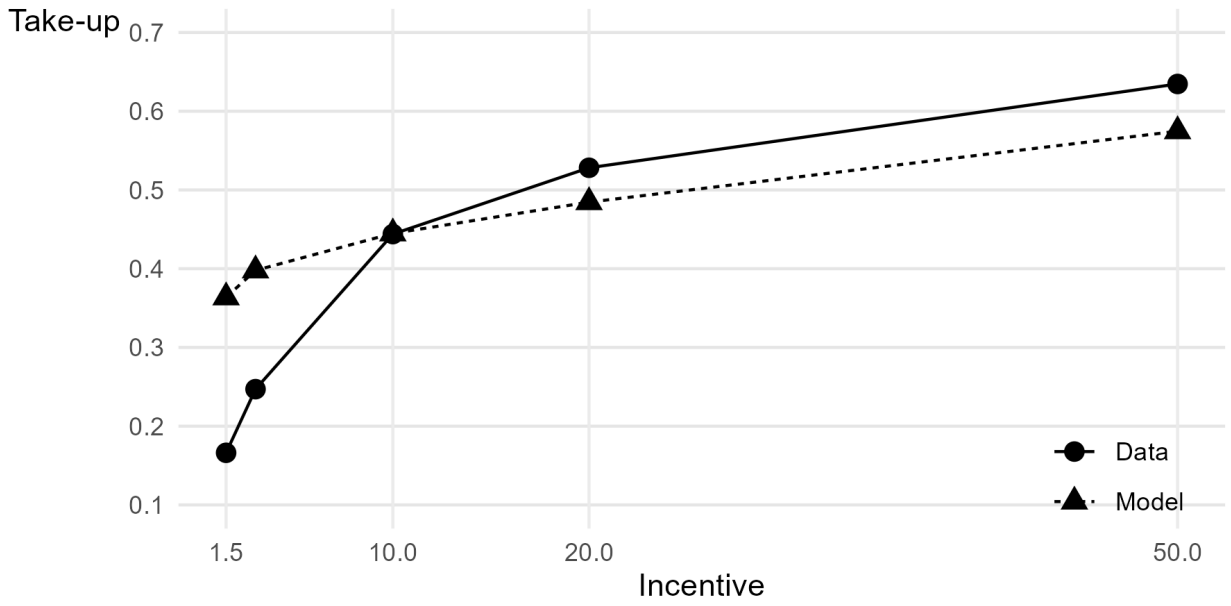
Notes: Each panel reports the fraction of people choosing to vaccinate given their choice set in our stated choice tasks, by vaccine intention group. The figures differ in the offered choice set. In Panel A we offer a realistic choice: no vaccine or a vaccine with no reward 50% effectiveness, 1% severe side effect rate, needle, inconvenient. In Panel B we show the most appealing choice set we offered: vaccine A has a $25 reward, 90% effectiveness, 0% severe side effect rate, needle delivery, convenient; vaccine B has a $5 reward, 50% effectiveness, 1% side effect rate, nasal delivery, and inconvenient. In Panel C, the choice set is the most appealing among those where all products have a positive side effect rate. Vaccine A has a $25 reward, 75% effectiveness, 5% side effect rate, and nasal spray, inconvenient. Vaccine B has a$0 reward, 90% effectiveness, 1% side effects, needle, and convenient.

Figure 5: Decomposition of differences in vaccination rates



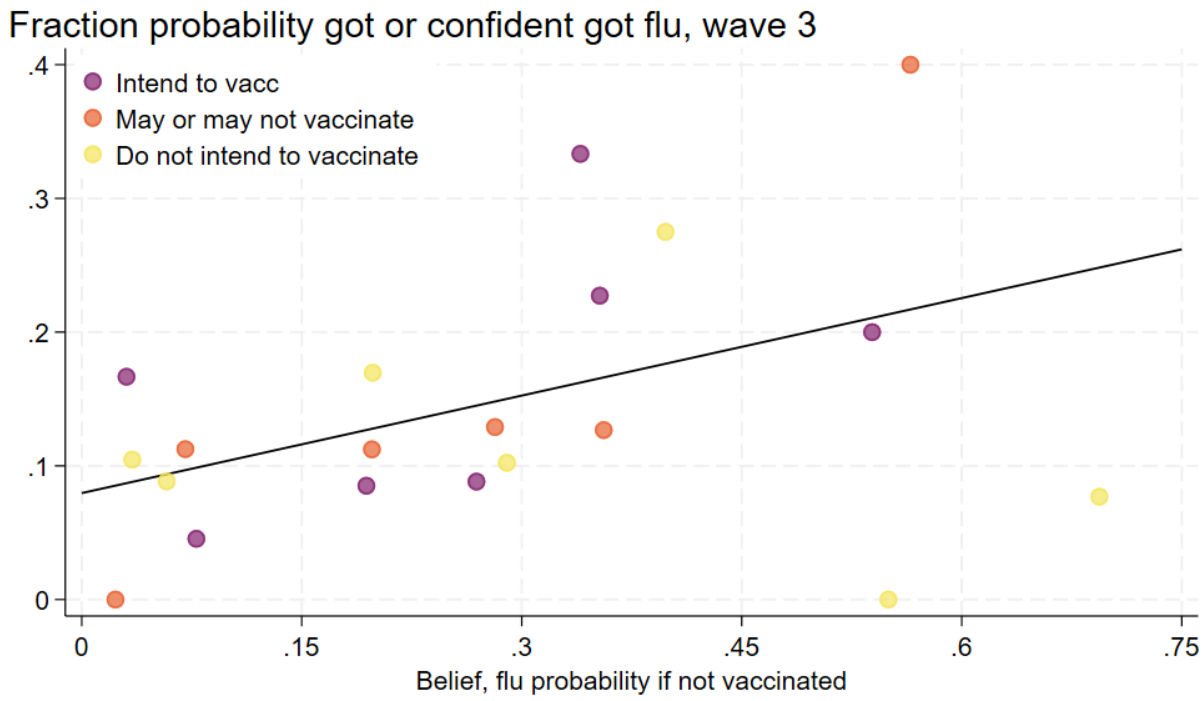
Notes: Each panel in the figure reports the result of a different simulation. The "baseline" panel shows our baseline simulation, using each groups' preferences and beliefs to predict vaccination rates. In "change beliefs" we change each groups' beliefs to be that of the already vaccinated. The dark outline shows the baseline probability in all panels. In "change preferences" we change each groups' preference distribution to match the preferences of the already vaccinated. In "change betao", we change each group's β_0 distribution to match the already vaccinated. In "change beta-SE, VE" we change each group's β_{SE} and β_{VE} to match the already vaccinated. In the final panel we also change each groups' beliefs about vaccine effectiveness and side effects.

Figure 6: Actual and model-predicted take-up of incentives to vaccinated



Notes: Figure reports the actual and simulated fraction of subjects who chose the indicated incentive level.

Figure 7: Beliefs about flu risk predict actual flu



Notes: Figure shows the correlation between beliefs about flu risk if unvaccinated and experienced flu if unvaccinated. Each point in the figure is a cell of vaccine intention and likert scale flu risk (would get the flu, likely to get the flu, etc.). The x-axis is the average believed flu risk in that cell, and the y-axis is the fraction of respondents in the cell who say in our wave 3 survey that they probably got the flu or are confident they got the flu.

Appendices

A Appendix Exhibits 57

B Belief elicitation questions 66

C Predictors of beliefs about flu risk 69

D MCMC Diagnostics 75

E Calibrating price sensitivity 78

F Pre-registration modifications and adjustments 80

A Appendix Exhibits

Table A.1: Vaccine hesitancy measures

Vaccine intention:	All	Already	Intend	May or may not	Do not intend
A. Agreement with "The information I receive from government sources about the flu vaccine is reliable"					
Strongly agree	0.45	0.68	0.58	0.36	0.20
Somewhat agree	0.33	0.26	0.35	0.47	0.30
Neither agree nor disagree	0.11	0.05	0.05	0.12	0.21
Somewhat diagree	0.07	0.01	0.01	0.04	0.18
Strongly disagree	0.04	0.00	0.01	0.01	0.12
B. Agreement with "Generally I do what my doctor recommends about vaccines"					
Strongly agree	0.46	0.71	0.65	0.39	0.12
Somewhat agree	0.33	0.24	0.31	0.48	0.33
Neither agree nor disagree	0.10	0.03	0.02	0.08	0.23
Somewhat diagree	0.07	0.01	0.01	0.04	0.18
Strongly disagree	0.05	0.01	0.01	0.01	0.13

Notes: Table reports the percent of respondents in each group choosing each level of agreement with the indicated statements.

Table A.2: Vaccination is predicted by beliefs and stated preferences

Beliefs:	Flu risk only		Vacc effect only		Both		+Preferences	
Specification:	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Flu rate	0.44	0.77			0.30	0.40	0.20	0.20
	(0.05)	(0.08)			(0.05)	(0.08)	(0.05)	(0.08)
Severe flu given flu	0.15	0.29			0.12	0.23	0.03	0.13
	(0.04)	(0.05)			(0.04)	(0.05)	(0.03)	(0.05)
Hospitalization rate if flu	-0.06	-0.05			-0.17	-0.27	-0.21	-0.18
	(0.12)	(0.21)			(0.12)	(0.20)	(0.12)	(0.18)
Death rate if hospitalized	0.30	0.37			0.31	0.36	0.24	0.13
	(0.13)	(0.26)			(0.12)	(0.24)	(0.12)	(0.22)
Side effect rate			-0.21	-0.26	-0.21	-0.27	-0.12	-0.18
			(0.04)	(0.06)	(0.04)	(0.06)	(0.03)	(0.06)
Vaccine effectiveness			0.63	0.83	0.53	0.71	0.25	0.36
			(0.04)	(0.05)	(0.04)	(0.06)	(0.04)	(0.06)
$\bar{P}(\theta)_i$							0.47	0.42
							(0.02)	(0.03)
N	2,073	2,073	2,074	2,073	2,073	2,072	2,073	2,072

Notes: Table reports results from a regression of vaccination status on subjective beliefs. Vaccination status is an indicator for already vaccinated as of wave 1, or self-reported vaccination in wave 3. Sample excludes people with missing belief measures or missing vaccinated status (which happens for people not vaccinated in wave 1, with no follow up data in wave 3). In the OLS columns, we regress vaccination on quantitative belief measures. In the IV columns, we instrument for quantitative belief measures using multiple choice questions where respondents give beliefs on a scale from "would not happen" to "unlikely" to... "would happen." $P(\beta)$ is each individual's posterior mean predicted vaccination rate given her posterior β, assuming 1 percent side effect rate and 50 percent vaccine effectiveness (i.e. fixing beliefs at roughly their objective level). Robust standard errors in parentheses.

Table A.3: Choice sets used in conjoint

Product	Reward	Effectiveness	Severe side effect rate	Needle	Convenient
Choice set 1:					
1	0	75	0	0	0
2	5	50	1	1	1
Choice set 2:					
1	25	90	5	0	1
2	0	50	0	1	0
Choice set 3:					
1	5	75	0	1	1
2	0	90	1	0	0
Choice set 4:					
1	25	75	1	1	0
2	0	50	5	0	1
Choice set 5:					
1	0	75	1	0	1
2	25	50	5	1	0
Choice set 6:					
1	5	90	0	0	0
2	0	75	5	1	1
Choice set 7:					
1	25	50	1	1	0
2	5	90	0	0	1
Choice set 8:					
1	25	75	5	0	0
2	0	90	1	1	1
Choice set 9:					
1	25	90	0	1	1
2	5	50	1	0	0
Choice set 10:					
1	25	50	0	0	1
2	5	90	5	1	0

Notes: Table reports the characteristics of vaccines shown in the conjoint experiment. The first vaccine listed was always labeled product A and the second product B. Each choice set also included a "do not vaccinate" option. Choices were shown in a random order.

Table A.4: Moments of the posterior preference distribution

Posterior	Mean	SD	2.5 percentile	97.5 percentile
Reward	0.134	0.245	0.00533	0.707
Vaccine Effectiveness	0.416	0.55	0.0306	1.88
Side effect rate	-1.36	5.14	-8.15	-0.0203
Needle	0.272	2.07	-3.85	4.35
Convenient	0.0539	0.79	-1.5	1.61
Inside option	-17.9	13.1	-43.5	7.91

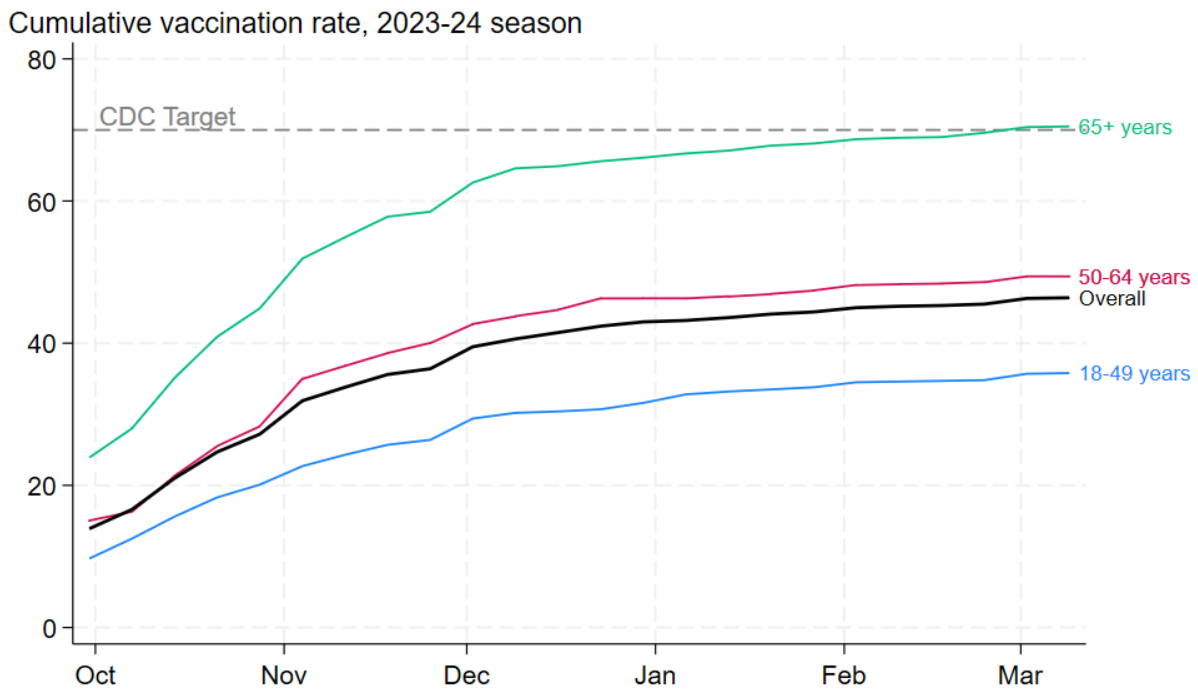
Notes: Table reports the posterior mean, standard deviation, and tail percentiles for the coefficients on the indicated vaccine characteristics. The "inside option" is the coefficient for vaccination relative to not vaccinating.

Table A.5: Caption

Intention group	Vaccination rate, baseline	Vaccination rate, No side effects
Yes	0.918	0.909
Intend	0.872	0.856
May	0.31	0.501
No	0.0302	0.0682

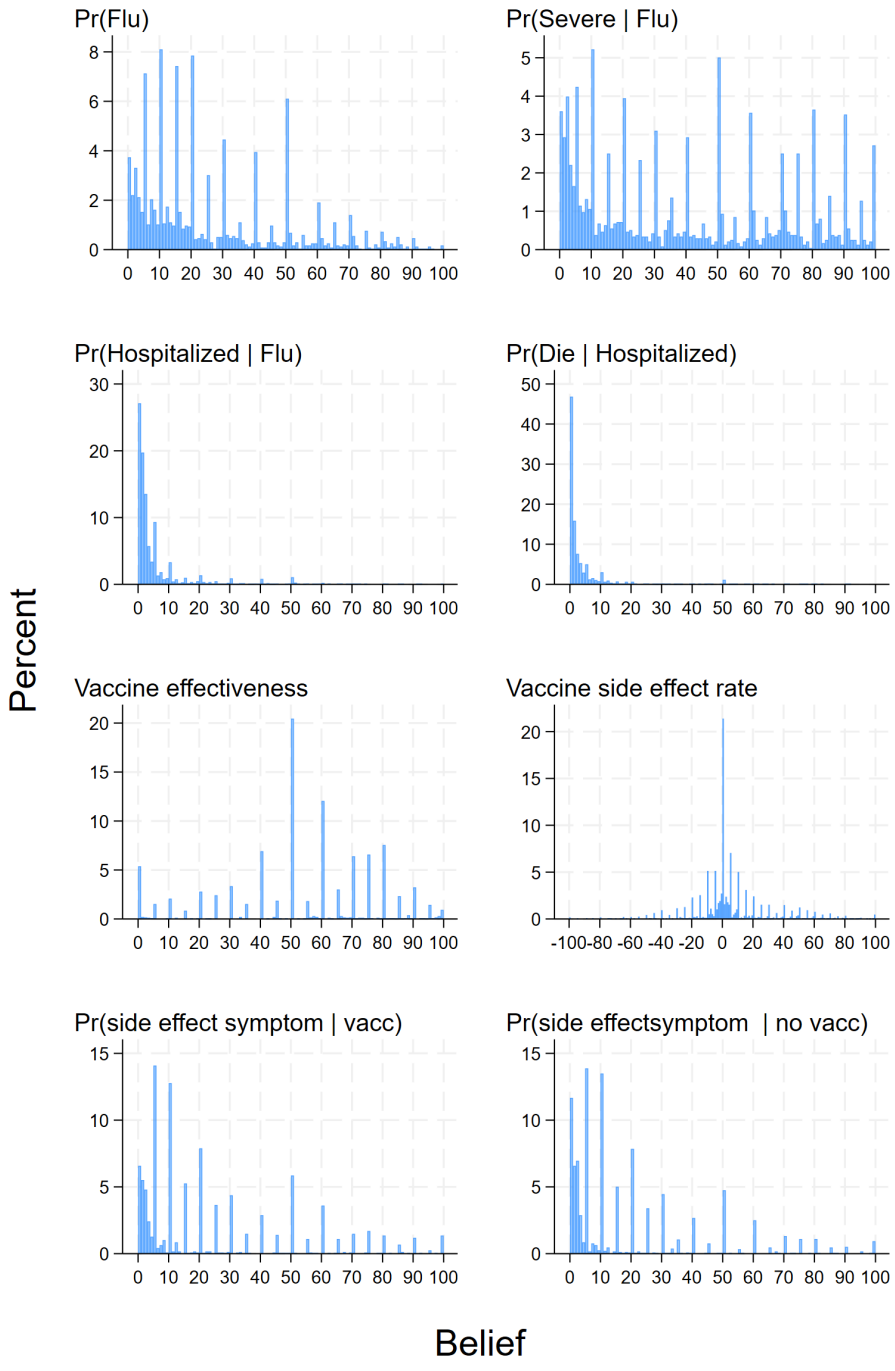
Notes: Table reports the simulated vaccination rates at the baseline available vaccine and when side effects of the vaccine are set to zero, by vaccine intention. This reduces vaccination rates for those already vaccinated and intending to vaccinate because they perceive the side effects of the vaccine to be negative.

Figure A.1: Vaccination rates are below CDC targets



Notes: Figure plots the weekly cumulative vaccination rate for the 2023-24 flu season, overall and by age group. Source: Center for Disease Control (2025c).

Figure A.2: Distribution of subjective beliefs



Notes: Figure plots the percent of respondents in each 1 percentage point bin of beliefs, for our subjective belief measures. The first four belief measures condition on not being vaccinated. The next two are beliefs about the flu vaccine (effectiveness against flu and effect on experienced side effects). The vaccine side effect rate is measured as the difference in likelihood of experiencing side effects symptoms with vaccination relative to without; report these raw measures in the final row.

Table A.6: Alternative specification for estimating side effect rates

Sample	10/20 Switchers	10/20 Switchers	All	All
Day control?	No	Yes	No	Yes
Y = vaccinated	0.10	0.10	0.13	0.12
	(0.04)	(0.04)	(0.03)	(0.03)
Y = experience severe symptoms	-0.20	-0.21	-0.16	-0.16
	(0.05)	(0.05)	(0.04)	(0.04)
Mean subjective SE rate	0.05	0.05	0.03	0.03
N	477	477	1,034	1,034

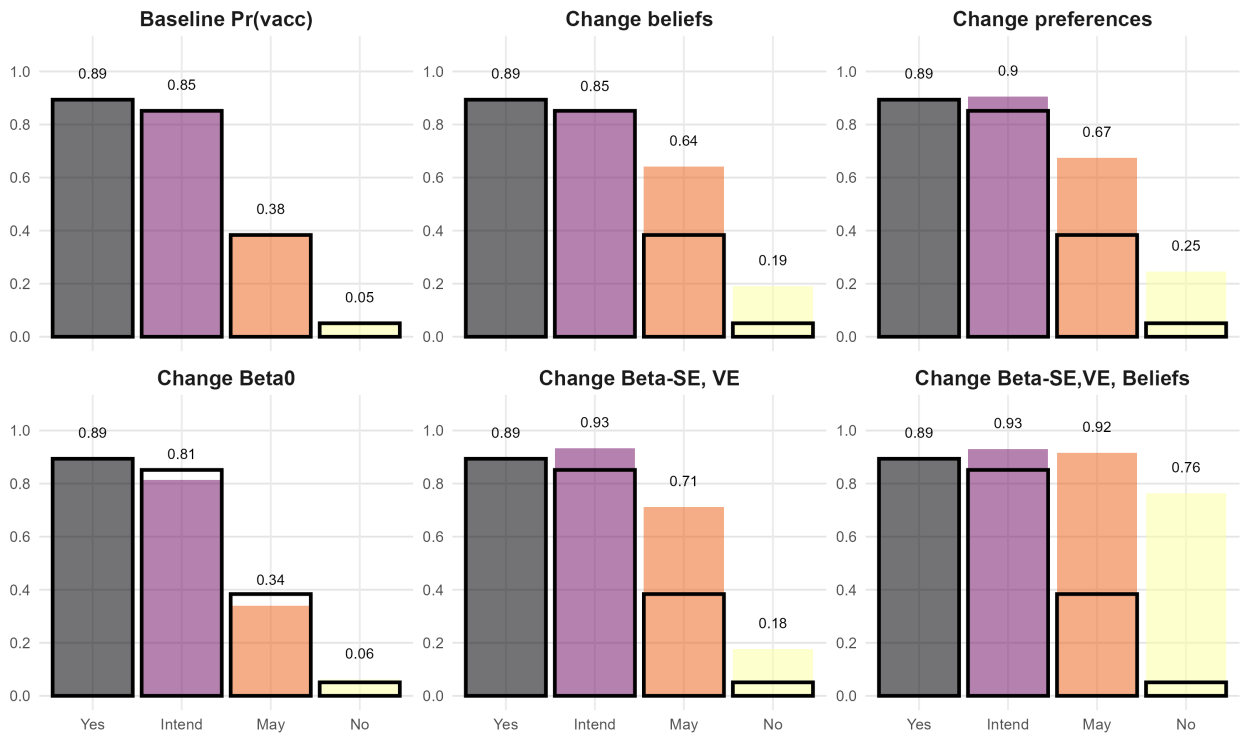
Notes: Table reports the effect of an incentive to vaccinate on the indicated outcome (top two rows, robust standard errors in parentheses), and the mean believed side effect rate for the estimation sample (bottom row). In each regression we regress the outcome on an incentive indicator, equal to 1 if we offered an incentive to vaccinate, along with a set of indicators for each respondents choices of incentives; the reward offer is random conditional on these choices. Where indicated, we also control for a set of indicators for response day. The sample is limited to people who chose at least one incentive to vaccinate, who completed the wave 2 survey, and had no reversals in their incentive choices. Robust standard errors in parentheses.

Figure A.3: Example choice set

Product	A	B
Gift card reward	$25	$0
Effectiveness	75%	90%
Severe side effect rate	5%	1%
Delivery mode	Nasal spray	Injection
Vaccination venue	Usual	Nearby

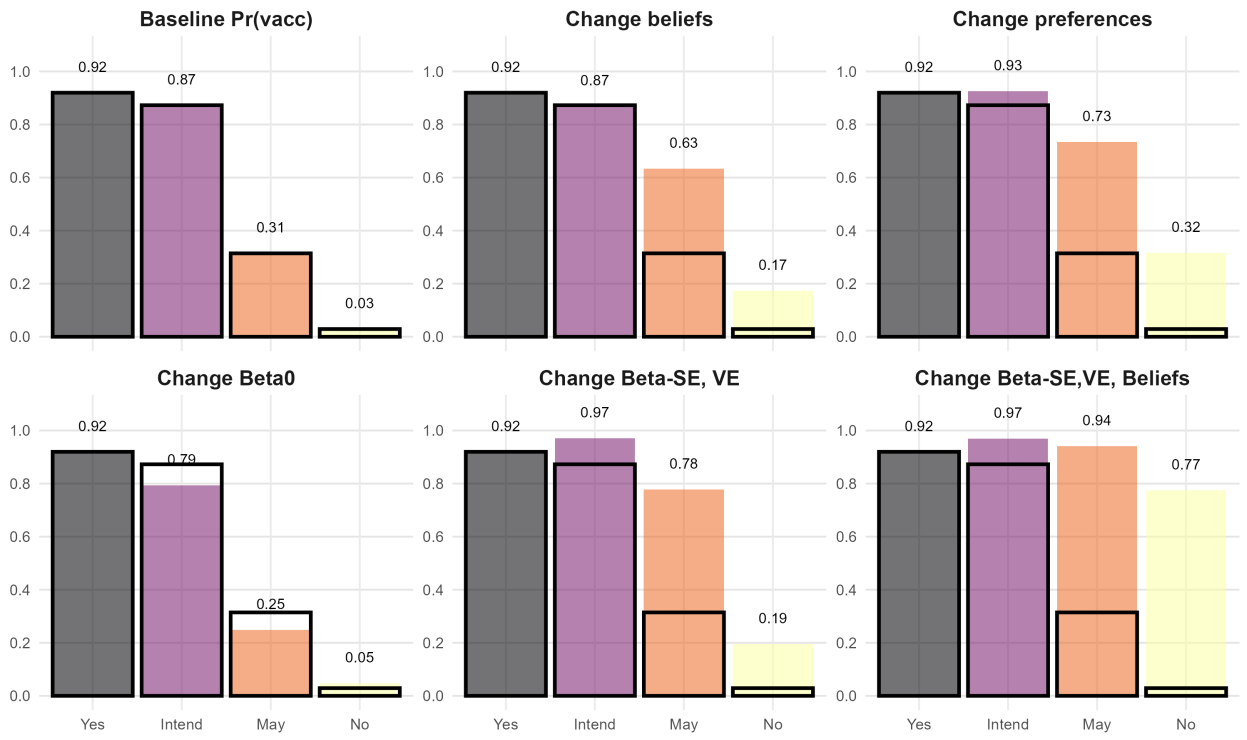
Notes: Figure shows how we displayed product characteristics in the conjoint design. Below this table we gave respondents a button to click for a reminder of the meaning of the different characteristics.

Figure A.4: Simulation robustness: Include pre-specified covariates only



Notes: See notes to Figure 5. This figure is identical except the model allows the mean of all random coefficients to depend only on indicators for above median age, above median education, and above median vaccine hesitancy score, and not on vaccine intention.

Figure A.5: Simulation robustness: Include more covariates



Notes: See notes to Figure 5. This figure is identical except the model allows the mean of all random coefficients to depend on indicators for above median age, above median education, and above median vaccine hesitancy score, in addition to depending on vaccine intention.

B Belief elicitation questions

This appendix describes our belief questions. We asked several sets of belief questions: beliefs about objective rates, beliefs about subjective flu risks if unvaccinated, and beliefs about flu vaccine characteristics. We asked each subjective question twice, first with a set of multiple choice questions, and then with a quantitative response. We report each set of questions in turn.

Objective questions: On a new survey block, we began by saying, "Now we'd like to ask you some questions about the country's experience with the flu virus last year." Then we ask the following questions:

- Flu rate: About what percent of the total U.S. population do you think got the flu in 2023-2024?
- Hospitalization rate: Of Americans who got the flu in 2023-2024, what percent do you think ended up getting hospitalized for the flu?
- Death rate: Of Americans who got hospitalized for the flu in 2023-2024, what percent do you think died of the flu that year?
- Vaccine effectiveness: The Center for Disease Control (or "CDC") measures flu vaccine effectiveness as the percent reduction in the chances of catching the flu. What do you think the CDC says is the effectiveness of the flu vaccine?

For all questions, we used a slider to elicit responses, with values ranging from 0 to 100. We focused on sliders here because we wanted to fix the scale of responses, but we did not want to suggest particular levels. A text entry box would have required that we say something like "write 1 for 1 percent, not .01.

Subjective questions about flu risk - multiple choice: We began by saying "Some people choose to get the flu vaccine, and others choose not to. For the next few questions, we'd like to ask your beliefs about what will happen to you if you choose not to get the flu vaccine." There are always seven possible answers to these multiple choice questions: would not X, unlikely to X, somewhat unlikely to X , neither likely nor unlikely to X, somewhat likely to X, likely to X, would X. We give the full response stem, e.g. X= "get the flu."

We randomly assigned half of respondents to see the CDC's estimates of the objective quantities. For these respondents, we reminded them of the CDC's answers, with text indicated in brackets. (The CDC does not report a severe flu rate, so neither do we.)

We asked the following questions:

- Subjective Flu rate, mc: If you don't get the flu vaccine in the next year, how likely would you say it is that you'd get the flu in the next year? [As a reminder, the CDC estimates that 10-20% of Americans got the flu last year.]
- Subjective severe flu rate, mc: Continue to assume you don't get the flu vaccine in the next year. Also assume you get the flu in the next year. How likely do you think it is you get a case bad enough to make you stay in bed for at least a day?
- Subjective hospital rate, mc: Continue to assume you don't get the flu vaccine in the next year, and also assume you get the flu in the next year. How likely do you think it is you get a severe case,

requiring hospitalization? [As a reminder, the CDC estimates that 0.6% to 3% of Americans who got the flu were hospitalized for it last year.]

- Subjective death rate, mc: Continue to assume you do not get the flu vaccine this year. Now also assume you get a severe case of the flu this year, requiring hospitalization. How likely do you think it is you die from the flu? [As a reminder, the CDC estimates that 3% to 18% of Americans who were hospitalized for the flu died of it last year.]

Subjective questions about flu risk - quantitative: After the multiple choice questions, we re-elicited beliefs, but asking on a quantitative scale. On a new page, we told respondents, "We just asked you about how likely certain flu-related events would be, if you don't get the flu vaccine. We're going to ask you the same questions, but this time we'd like your best guess of the numerical percent for each of them." We had respondents answer with a slider. For each question, the range of the slider was 0-25 if the respondent's multiple choice response was "unlikely to X" or less likely. Otherwise the range was 0-100.

We asked the following questions:

- Subjective Flu rate, quant: Assuming you did not get the flu vaccine in the next year, what do you think is the percent chance you get the flu in the next year? [As a reminder, the CDC estimates that 10-20% of Americans got the flu last year.]

- Subjective Severe rate, quant: Continue to assume you do not get the flu vaccine in the next year. Now also assume you do get the flu in the next year. What do you think is the percent the chance you get a case bad enough to make you stay in bed for at least a day?

- Subjective hosp rate, quant: Continue to assume you do not get the flu vaccine in the next year, and that you do get the flu in the next year. What do you think is the percent chance you would be hospitalized from the flu? [As a reminder, the CDC estimates that 0.6% to 3% of Americans who got the flu were hospitalized for it last year]

- Subjective death rate, quantitative: Continue to assume that you do not get the flu vaccine in the next year. Now also assume that you get a severe case of the flu in the next year, requiring hospitalization. What do you think is the percent chance you would die from the flu? [As a reminder, the CDC estimates that 3% to 18% of Americans who were]

Vaccine characteristic questions: On a new page, we said, "The questions on this page ask how you think getting the flu vaccine will affect you." We then asked:

- Subjective convenience: If you were to get the flu vaccine, how convenient would it be to do so? [Multiple choice question with answers: not an inconvenience, minor inconvenience, worse than a minor inconvenience.]

- Subjective vaccine effectiveness, mc: If you do get the flu vaccine, how much would you say that reduces your chances of getting the flu? [As a reminder, the CDC estimates that the flu vaccine reduces your chances of getting the flu by 40-60%.] [Multiple choice question with possible answers: eliminates, substantially reduces, somewhat reduces, does not reduce, increases.]

- Subjective vaccine effectiveness, quant: Now we'd like you to put a number of this. By what percent do you think the flu vaccine would reduce your chance of getting the flu. 0 means "no reduction" and 100 means "completely eliminates my chance of getting the flu." [Text entry box, required a number, allowed -100 to +100.]

Side effect rate questions: We measure the subjective side effect rate by asking about side effect rate if unvaccinated and if vaccinated, and then taking the difference. We ask multiple choice versions as well as quantitative version of these questions. On one page we ask about side effect rates if not vaccinated, and then on the next we ask about side effect rate if vaccinated. We ask the following questions.

- Subjective side effect rate, no vaccination, mc: After getting vaccinated, sometimes people experience headaches, muscle aches, upset stomach, tiredness, joint pain, fever, or general discomfort. These side effects are severe if they disrupt your daily activities. People experience these symptoms for all kinds of reasons, not just from the flu vaccine. For this question, imagine that you do not get vaccinated in the next two weeks. Thinking about yourself, what do you think are the chances you would experience a case of headache, muscle ache, upset stomach, tiredness, joint pain or fever, severe enough to disrupt your daily activities in the next two weeks?

- Subjective side effect rate, no vaccination, quantitative: Now we'd like you to put a percent on this likelihood. Thinking about yourself, and imagining you do not get the flu vaccine, what do you think is the percent chance that you would experience a severe case of one or more of these symptoms in the next two weeks? Again severe means bad enough to disrupt your daily activity. Please enter a number between 0 and 100, where 0 means you would not experience a severe case and 100 means would experience it for sure. [text entry box]

- Subjective side effect rate, vaccination, mc: Now imagine that you do get the flu vaccine in the next two weeks. After getting vaccinated, sometimes people experience headaches, muscle aches, upset stomach, tiredness, joint pain, fever, or general discomfort. These side effects are severe if they disrupt your daily activities. Of course, people experience these symptoms for all kinds of reasons, not just from the flu vaccine. Continuing to imagine that you get the flu vaccine in the next two weeks, and thinking about yourself, what do you think are the chances you would experience a case of headache, muscle ache, upset stomach, tiredness, joint pain

- Subjective side effect rate, vaccination, quantitative: Now we'd like you to put a percent on this likelihood. Imagining you get the flu vaccine in the next two weeks, and thinking about yourself, what do you think is the percent chance that you would experience a severe case of one or more of these symptoms in the next two weeks? Again severe means bad enough to disrupt your daily activity. Please enter a number between 0 and 100, where 0 means you would not experience a severe case and 100 means would experience it for sure. [text entry box]

C Predictors of beliefs about flu risk

This appendix presents some results on individual-level predictors of beliefs about flu risk. We focus on two sets of predictors which are known to correlate with flu risk and which we measured in our experiment: age and chronic disease. Our age variables are indicators for age 18-39, 40-54, 55-64, and 65 or older. Our disease variables are indicators for asthma, diabetes, heart disease, lung disease, kidney disease, and pre-diabetes. (This last variable is not a known risk category, but pilot surveys indicated that we should include it to avoid confusion with diabetes.) We also allow people to report that they have at least one of these conditions, but they prefer not to say which. We did not ask about other correlates of flu risk (e.g., occupational exposure) nor about prior flu experience.

Appendix Table C.1 reports the means of the risk factor variables. While 21 percent of sample has one or more risk factors, most risk factors are infrequent, with Asthma, diabetes, and pre-diabetes being the most common. Heart, lung, and kidney disease are all rare.

These risk factors predict individual flu risk beliefs, although some of the predictions are noisy. To show this, we regress our dimensions of flu risk beliefs (incidence, and conditional probability of severe flu, hospitalization,and death) on the indicators, in groups and jointly. We report the results of these regressions in Appendix Tables C.2 -C.5. Generally, risk factors are correlated with believed risk in an intuitive way. Older people believe they have a higher risk of hospitalization and death. Having any risk factor is associated with a higher perceived change of getting the flu, and of severe outcomes conditional on the flu. Individual risk factor predictions are fairly noisy, as are the joint predictions, in part because of the rarity of some risk factors and in part because of their correlation with age. One apparent puzzle is that age is negatively correlated with the perceived risk of flu incidence. However this fact aligns with CDC estimates: elders are less likely to get the, but much more likely to experience severe cases of the flu if they get it Centers for Disease Control and Prevention (2024a).

Table C.1: Prevalence of risk factors in our sample

Risk factor	rate
Age 40-54	0.294
Age 55-64	0.095
Age 65+	0.104
Asthma	0.101
Pre-diabetes	0.054
Diabetes	0.049
Heart disease	0.020
Lung disease	0.008
Kidney disease	0.010
Has risk, prefers not to say which	0.005
Has one or more risk factor	0.213

Notes: Table reports means of the indicated risk factor indicators, for our analysis sample.

Table C.2: Predicting believed flu rate

Risk factors	Age	Any risk	Risks	All
Age 40-54	-1.802			-2.205
	(1.090)			(1.093)
Age 55-64	-3.757			-4.269
	(1.608)			(1.614)
Age 65+	-4.447			-5.971
	(1.430)			(1.425)
At least one risk factor		5.341		
		(1.208)		
Asthma			5.001	4.912
			(1.664)	(1.668)
Diabetes			3.146	4.327
			(2.129)	(2.146)
Heart disease			6.087	8.703
			(3.882)	(3.895)
Lung disease			3.273	4.353
			(5.676)	(5.632)
Kidney disease			2.980	4.669
			(4.676)	(4.534)
Has risk, prefers not to say which			-7.556	-6.127
			(4.527)	(4.791)

Notes: Table reports coefficients from a regression of the believed rate of getting flu if unvaccinated, on the indicated variables. Each column is a separate regression. The omitted age category is 18-39. Robust standard errors in parentheses.

Table C.3: Predicting believed severe flu rate

Risk factors	Age	Any risk	Risks	All
Age 40-54	3.854			3.510
	(1.540)			(1.544)
Age 55-64	3.644			3.334
	(2.259)			(2.271)
Age 65+	2.078			0.560
	(2.183)			(2.261)
At least one risk factor		7.392		
		(1.609)		
Asthma			6.835	6.718
			(2.193)	(2.192)
Diabetes			2.082	1.766
			(2.971)	(3.007)
Heart disease			9.723	9.215
			(4.496)	(4.610)
Lung disease			14.413	15.035
			(6.761)	(6.648)
Kidney disease			1.321	0.945
			(6.793)	(6.789)
Has risk, prefers not to say which			-8.024	-8.138
			(8.641)	(8.740)

Notes: Table reports coefficients from a regression of the believed rate of severe flu if unvaccinated and infected with flu, on the indicated variables. Each column is a separate regression. The omitted age category is 18-39. Robust standard errors in parentheses.

Table C.4: Predicting believed flu hospitalization rate

Risk factors	Age	Any risk	Risks	All
Age 40-54	-1.175			-1.514
	(0.539)			(0.532)
Age 55-64	1.244			0.760
	(0.967)			(0.964)
Age 65+	2.269			0.852
	(1.005)			(0.960)
At least one risk factor		3.825		
		(0.724)		
Asthma			2.649	2.789
			(0.926)	(0.925)
Diabetes			4.501	4.351
			(1.479)	(1.501)
Heart disease			6.505	6.058
			(3.125)	(3.080)
Lung disease			3.643	3.210
			(3.824)	(3.804)
Kidney disease			6.335	6.178
			(3.789)	(3.888)
Has risk, prefers not to say which			3.034	2.815
			(4.107)	(4.028)

Notes: Table reports coefficients from a regression of the believed rate of flu hospitalization if unvaccinated and infected with flu, on the indicated variables. Each column is a separate regression. The omitted age category is 18-39. Robust standard errors in parentheses.

Table C.5: Predicting believed flu death rate

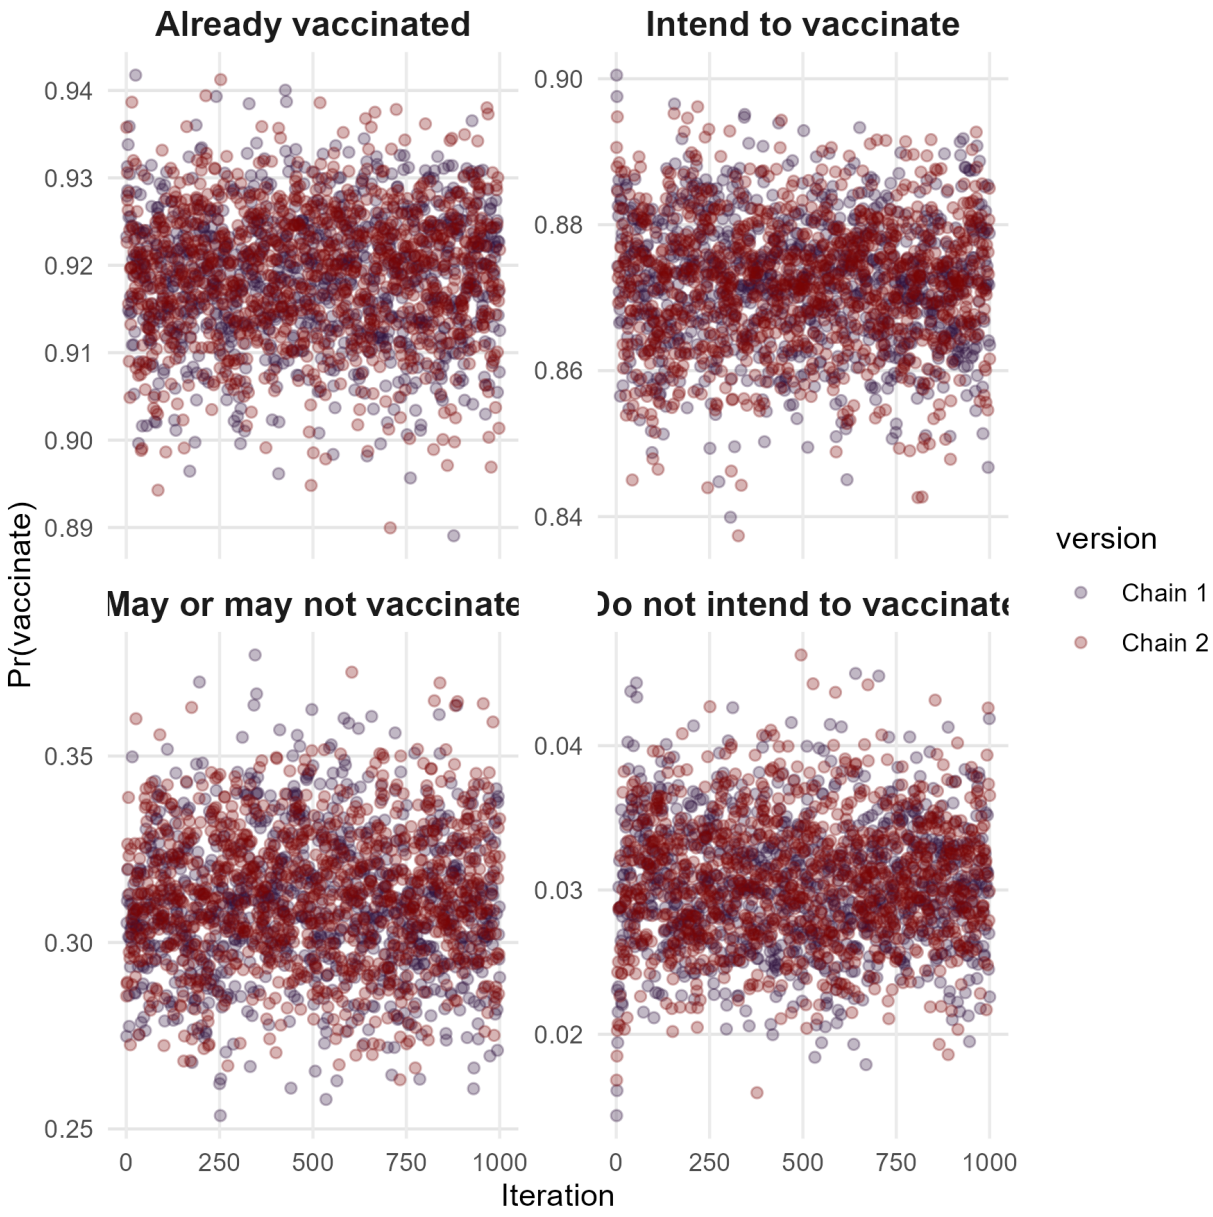
Risk factors	Age	Any risk	Risks	All
Age 40-54	-0.825			-0.921
	(0.453)			(0.456)
Age 55-64	0.651			0.532
	(0.782)			(0.795)
Age 65+	1.981			1.604
	(0.828)			(0.821)
At least one risk factor		1.330		
		(0.551)		
Asthma			0.825	0.931
			(0.747)	(0.746)
Diabetes			1.531	1.269
			(1.084)	(1.077)
Heart disease			1.892	1.220
			(2.136)	(2.081)
Lung disease			1.652	1.080
			(2.993)	(3.051)
Kidney disease			2.698	2.329
			(2.389)	(2.454)
Has risk, prefers not to say which			0.380	-0.014
			(2.532)	(2.503)

Notes: Table reports coefficients from a regression of the believed rate of flu hospitalization if unvaccinated and hospitalized with flu, on the indicated variables. Each column is a separate regression. The omitted age category is 18-39. Robust standard errors in parentheses.

D MCMC Diagnostics

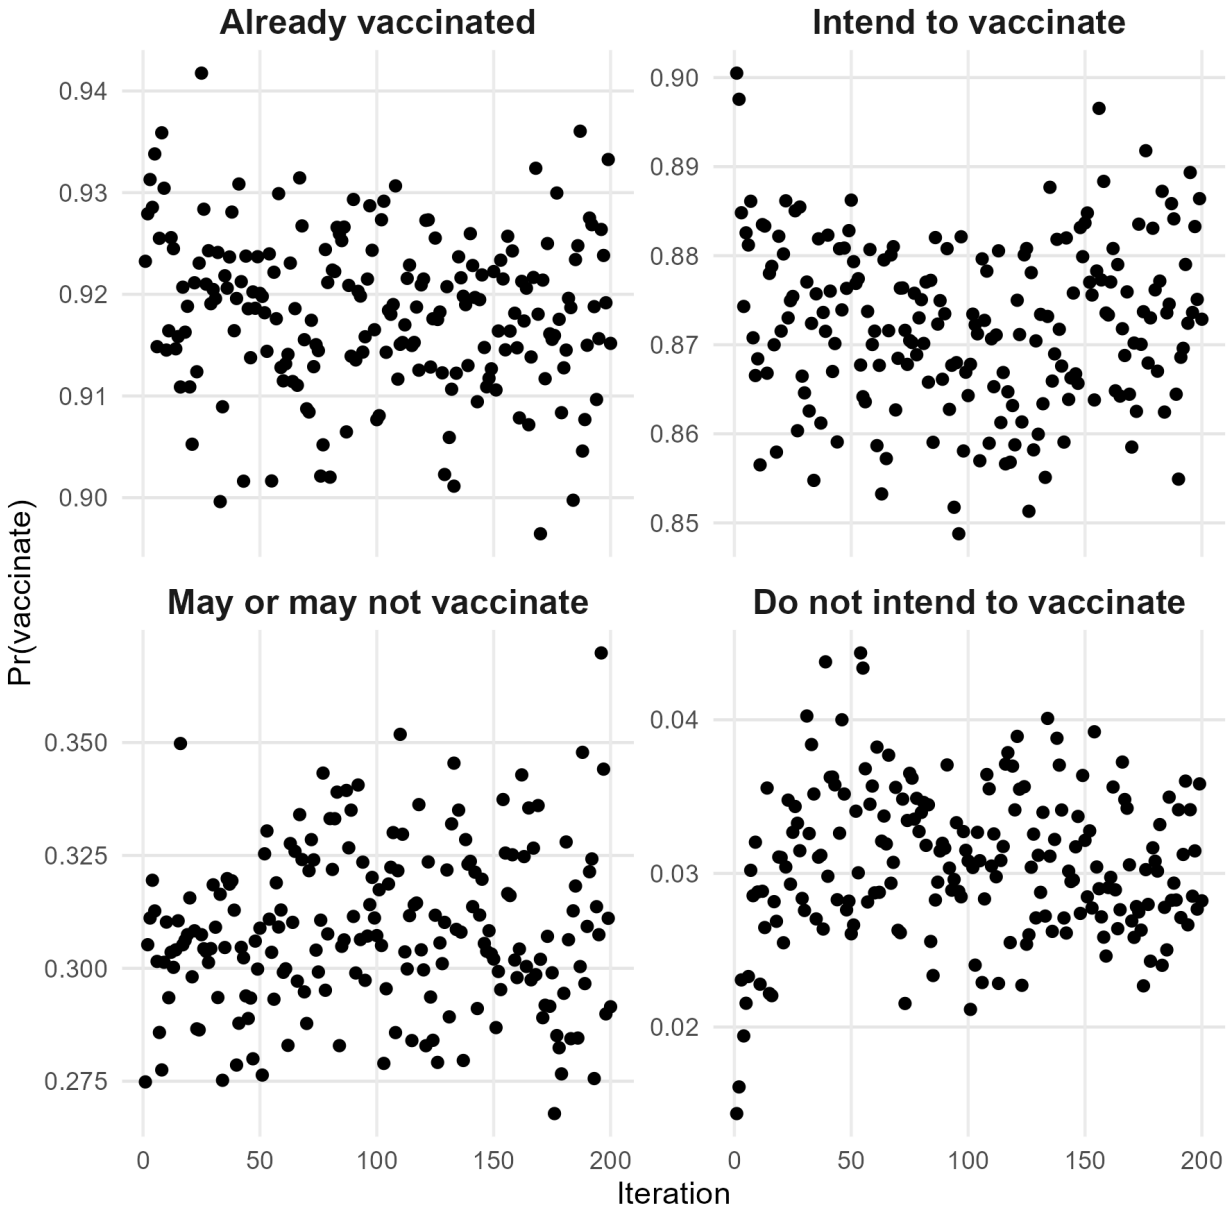
We rely on MCMC methods to sample from the posterior distribution of the model parameters. Here we provide two pieces of evidence on the convergence of our Markov chains. Appendix Figure D.1 shows that our chains exhibit good mixing, and Appendix Figure D.2 shows that they are stationary. To construct Appendix Figure D.1, we draw two independent sets of Markov chains, using different random seeds. Then at each iteration of each chain (after keeping every 300th draw, as in our main specification), we calculate the model-implied mean probability of vaccination, by intention group. If our chains have reached their stationary distribution, then the two chains should closely overlap, and neither should demonstrate any trend. Appendix Figure D.1 shows clear mixing between the two chains. Careful examination shows a bit of a trend in the first few iterations. We zoom in on this area in Appendix Figure D.2, showing probabilities for the first 200 iterations only. Convergence appears rapid: after the first 25 or so iterations (i.e. the first 7500 draws before thinning), the chains stabilize, showing no further systematic trend.

Figure D.1: Markov chains are well-mixed



Notes: Figure plots the simulated vaccination rate, at each iterations of the Markov chain, for two separate starting values.

Figure D.2: Markov chains converge rapidly



Notes: Figure plots the simulated vaccination rate, at each of the first 200 iterations of the Markov chain, after thinning to every 300th draws.

Table E.1: Price sensitivity is higher in revealed preference than stated preference data

Data source	Vaccination probability	Price semi-elasticity
Revealed preference	0.22	0.028
Stated preference	0.51	0.005

Notes: The first row reports the wave 2 vaccination probability and implied semi-elasticity of vaccination with respect to price, which is conditionally randomly assigned. The second row simulates these moments using stated preference data.

E Calibrating price sensitivity

This appendix describes how we adjust our price sensitivity estimates to obtain an estimate of the marginal utility of income used in calculating VSLs. Our unadjusted price sensitivity estimate is estimated from the stated preference data in the conjoint. It is identified by the response of stated vaccine choices to gift card reward amounts. Two features of the data suggest, however, that our price sensitivity estimate is too low. First, when we offer incentivized choices (over vaccination bonuses), we find much more price sensitivity than in our stated preference (Figure 6). Second, we show below that when we use the conditionally randomly assigned vaccination bonus as an instrument for actual vaccination, we find price sensitivity much higher than is by the stated preference data. Here we show that rationalizing the revealed preference data requires roughly tripling our baseline price sensitivity estimate; this is the adjustment we make.

Price sensitivity in revealed preference data: We use the revealed preference data—wave 2 vaccination rates—along with conditionally random assigned vaccination incentives to obtain a revealed preference price sensitivity estimate. Recall that we wave 1 respondents a choice of incentives for wave 2: $3 for sure and a bonus of 0, or $2 plus a bonus b if they show proof of vaccination. We say a respondent is consistent if they always select a higher incentive whenever they select a lower incentive. We do not force respondents to make consistent choices but 97.5 percent do so. We define wta_i as the minimum b for which respondent i chooses the bonus rather than the larger certain reward.[38] After respondents make their choices, we pick one at random to implement. This creates random variation in incentives to vaccinate, conditional on wta.

To use this this information, we estimate the following logistic regression:

$$Pr(vacc_i = 1) = \frac{\exp(\alpha bonus + \theta_{wta})}{1 + \exp(\alpha b + \theta_{wta})}. \tag{2}$$

Here $bonus$ is the selected bonus amount (0 or b), and θ_{wta} is a set of indicators for level of wtp. $bonus$ varies because of our random assignment of choices to implement, and because of people's choices, but θ_{wta} conditions on choices fully. We estimate this regression using the 1,034 wave 2 respondents with consistent choices and $wta_i \leq 50$, for whom there is variation in $bonus$. We use our estimates to obtain a semi-elasticity, $\eta^{RP} = b(1 - \hat{Pr}$, where $\hat{Pr}$ is the predicted probability.

Price sensitivity in stated preference data: To compare our revealed preference semi-elasticity to the stated preference semi-elasticity, we simulate a comparable choice situation. Specifically, among respondents with $wta_i < 50$ and eligible for wave 2, we randomly assign prices of $1.50, $3, $10, $20, and

[38] This is a slight abuse of notation because this switch point merely bounds willingness to accept.

simulate choice probabilities at these prices or $0. We calculate the arc semi-elasticity of going from $0 to each of these prices, and report the overall average.

Stated preference price sensitivity is too low: We report the revealed preference and stated preference results in Appendix Table E.1. The price semi-elasticity is almost six times higher in the revealed preference data than in the stated preference data. (The price sensitivity parameter α is fairly precisely estimated, with a z-statistic of 4.) This reflects higher price sensitivity but also lower take-up, as the semi-elasticity is decreasing in take-up. Actual take-up is lower than predicted, consistent with the short (two-week) follow-through time period, and the action-intention gap that we and others have documented (e.g., Bronchetti et al. (2015)). If revealed preference take-up were as high as predicted by stated preferences, we would have a price semi-elasticity of 0.0.18 rather than 0.028. We conclude that the stated preference price sensitivity is too low by a factor of 3.6 (that is, 0.018/0.005), and to obtain marginal utility of income for calculating VSLs, we scale up reward sensitivity by 3.6.

F Pre-registration modifications and adjustments

We pre-registered our experiment in the in the AEA RCT registry.[39] Here we describe the omissions and deviations from our registration. Our plan omitted some details which are necessary for implementation:

1. Handling missing belief data. We specified we would drop cases with missing demographic data, but we did not specify how we would handle case with missing belief data. We dropped the one observation which otherwise met our inclusion criteria but lacked beliefs about flu vaccine effectiveness.

2. Summary of WTA and beliefs. We pre-specified conjoint-estimated WTA and beliefs about vaccine effectiveness and side effects as two of our primary outcomes. However we did not specify how we would summarize these objects. We report means and percentiles, overall and by groups defined by vaccine intention. We also did not specify whether we would use subjective or objective characteristics in calculating WTA. We report both. We report WTP rather than WTA because that's what the model spits out.

3. We lack follow-up data (including vaccination) for some participants, either because they did not take the wave 2 survey or because we could not link their ids. We did not pre-specify how we would handle this situation, but we treat the data as missing ar random; we drop them from the wave 2 data analysis only.

4. Verifying vaccination. We sent multiple reminders to take follow up survey, and for the incentive group, when people did not verify their vaccination initially, we sent them follow ups to a simple, third survey, to upload vaccination if they had it (and receive their reward for doing so).

5. Estimating treatment effects of vaccination:

   (a) Handling inconsistent answers to incentives. We define complier groups by the "switching point" where people opt in to incentives to vaccinate. We did not specify how we would handle people with inconsistent switching points (e.g., no vacc at $1.50 or $3, vacc at $10, no vacc at $20). We exclude such people from the complier groups.

   (b) First stage. We did not pre-specify that we would report "first stage", effects of incentives on vaccination, but we do so for each complier group.

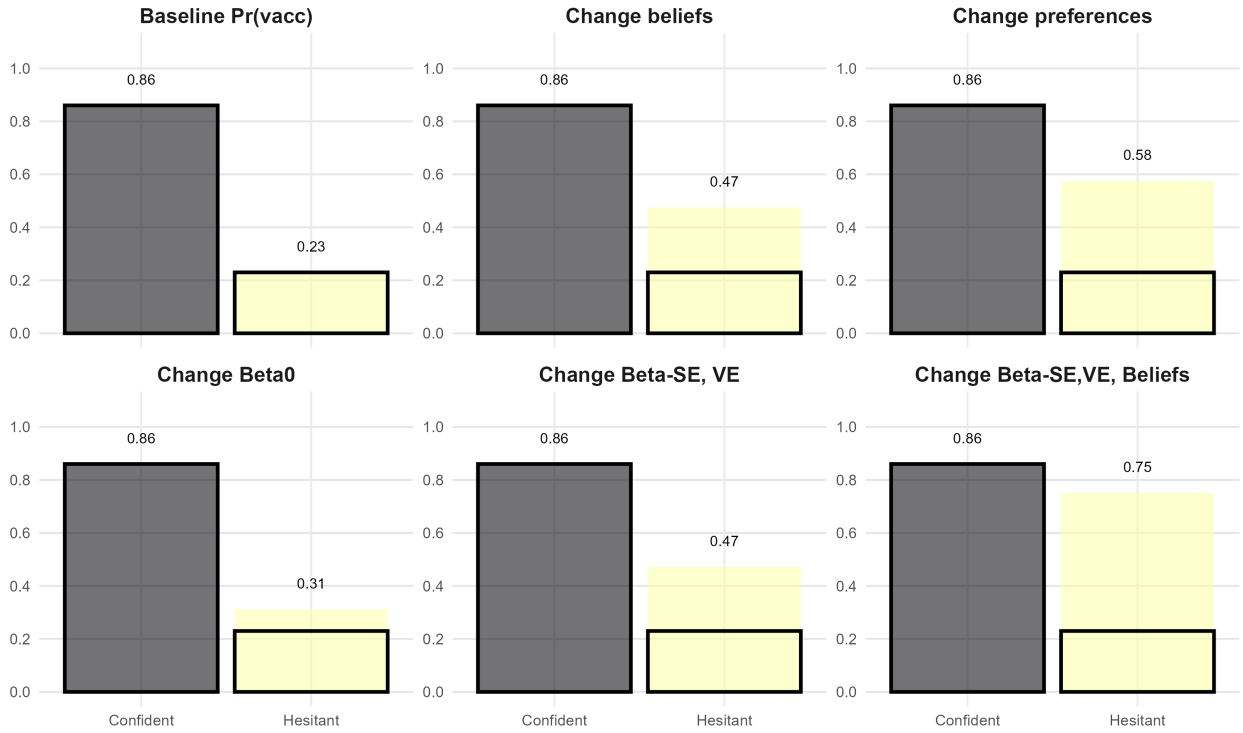
   (c) Specification of IV model. We said we would estimate the TOT of vaccination on experienced side effects using "standard instrumental variables techniques." We imagined that we would have a strong first stage, but in fact the first stage F-statistic is below 10, so we do not estimate the IV model.

Although we report all pre-registered tests in the main text or appendices, we note that our main analyses deviate from the pre-registration in two substantial ways:

1. We let vaccine preferences differ by intention group, and we organize our analysis around vaccine intentions, rather than using vaccine hesitancy questions as we had preregistered. We report results for vaccine hesitancy in Appendix Figure F.1. Here we show simulations analogous to those

[39] Registration here: `https://www.socialscienceregistry.org/trials/14763`.

Figure F.1: Decomposition of differences in vaccination rate by vaccine hesitancy



Notes: See notes to Figure 5. This figure is identical except for two respects. First, it is based on a demand system where the mean of all coefficients depends only on indicators for above median age, above median education, and above median vaccine hesitancy score, and not on vaccine intention. Second, we stratify on vaccine hesitancy or confidence, rather than on vaccine intention.

in the main text, but stratifying on vaccine hesitancy rather than intention. Hesitancy is defined as defined as having above median total answers to our two hesitancy items (roughly, trust in doctors and trust in government; see Appendix Table A.1). As in our main simulations, preferences over and beliefs about vaccine characteristics are quite important in explaining differences in vaccination rates.

2. We conducted a third wave, to re-measure vaccinations with less experimenter demand, and to measure flu experiences. We did not preregister this third wave.