Preferences, Beliefs, and Demand for the Flu Vaccine [*]

Daniel W. Sacks[†] Justin Sydnor[‡]

August 15, 2025

Abstract

Can economic tools help inform the puzzlingly low rate of flu vaccination? Existing interventions focus on misinformation or nudges, not preferences over or beliefs about vaccine characteristics. Using an online experiment, we find a key role for effectiveness and short-run side-effects: equalizing only beliefs and preferences about these characteristics nearly eliminates the vaccination-intention gap between the vaccine hesitant and confident. Fear of needles, inconvenience, and perceived long-run health risks play smaller roles. The vaccine hesitant hold pessimistic but plausible beliefs about effectiveness but greatly overestimate side-effect risks. We estimate the impact of correcting inaccurate beliefs and potential subsidies on vaccination rates.

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

[*]We thank Ashvin Gandhi, Brad Shapiro, Patrick Sullivan, Coady Wing, our discussant Mireille Jacobson, and seminar audiences at ASHEcon, Highland Health Economics Summit, Notre Dame, UCSD Rady for helpful comments. 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).

1 Introduction

Heterogeneity in preferences and beliefs plays a central role in positive and normative economic analysis, because preference heterogeneity determines optimal allocations and belief heterogeneity creates scope for changing behavior (e.g. Berry et al. (1995); Heckman et al. (2006); Dickstein and Morales (2018); Arteaga et al. (2022); Akbarpour et al. (2024); Bates et al. (2025)). Yet recognition of this heterogeneity is rarely part of public health recommendations, which often uniformly encourage or discourage a screening or treatment on the basis of population-average estimates of effectiveness and risks. Moreover, there is typically little information about variation in preferences and beliefs to help researcher better understand health decisions and to inform health policy.

We examine these issues in the context of vaccinations for the seasonal flu. Seasonal diseases, such as Influenza, COVID-19, and Respiratory Syncytial Virus, regularly infect 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). The U.S. Centers for Disease Control and Prevention (CDC) recommends that all Americans aged 6 months and older get the flu vaccine. Despite the CDC recommendation and the fact that vaccines are usually free, in a typical year, fewer than 35 percent of U.S. adults report vaccinating (Kriss, 2024). Low vaccination rates are concerning both because vaccines against infectious disease can prevent disease spread, a classic externality (Ward, 2014; White, 2021; Freedman et al., 2022), and because vaccine choices appear on the surface to imply that people place low value on their own health.[1]

Research on vaccine hesitancy—"the delay in acceptance or refusal of vaccination despite availability of vaccine services" (MacDonald et al., 2015)—has taken two primary approaches. One documents and attempts to correct scientifically unwarranted beliefs, including views that vaccines cause severe illness such as autism or cancer, or even cause the targeted disease itself.[2] An alternative hypothesis centers on the intention-action gap: many people would like to get vaccinated, but fail to follow through. This hypothesis often motivates work to increase vaccination through nudges and other light-touch interventions (for example, Milkman et al. (2022); Patel et al. (2023)). Both approaches essentially assume that vaccine hesitancy is a mistake. An alternative possibility, however, is that vaccine hesitancy reflects a reasoned response to the perceived benefits and costs of vaccinations. That is, for some, vaccines may not be effective enough, or disease not severe enough, to justify the side-effects, inconvenience, or discomfort with needles (Freeman et al., 2023). Indeed, economic epidemiology assumes that vaccination decisions are privately rational—though socially suboptimal in the presence of externalities—given individuals' preferences over health and their beliefs about prevailing risks and vaccine consequences (Philipson, 2000;

[1] Carlin et al. (2022), for example, document very low willingness-to-pay for the COVID-19 vaccine, which in turn implies a value of a statistical life (VSL) orders of magnitude below consensus estimates.

[2] 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).

Goodkin-Gold et al., 2020). People may also engage in this type of decision calculus, but with biases in their perceptions or beliefs about disease risk or side-effect rates. Yet little is known about how people perceive or value vaccine characteristics and their beliefs about disease risk.

We conduct an online experiment that provides evidence on vaccine intentions, beliefs about vaccines, preferences over vaccine characteristics, and the accuracy of people's beliefs. Our central finding is that differences in preferences and beliefs related to core aspects of vaccination – namely, the effectiveness of the vaccine and the rate of short-term side effects – can largely rationalize low vaccination rates. Other aspects, such as scientifically unwarranted beliefs about long-run health effects and failures to follow through on intentions, are quantitatively less important. We are able to further identify that the vaccine hesitant hold biased beliefs, specifically about side-effect rates, that if corrected could significantly boost vaccination rates. Yet we also find that many who do not intend to vaccinate have preferences that suggest vaccination is privately suboptimal for them given current levels of vaccine effectiveness.

The experiment includes a number of components that involve eliciting beliefs and preferences along with randomized provision of incentives and follow-up surveys that help to identify actual vaccination behavior. 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.[3] We measure preferences using a conjoint design that asks subjects to make choices over different combinations of hypothetical vaccine options. Such designs are widely used in marketing and increasingly common in economics (Allenby et al., 2019; Ben-Akiva et al., 2019; Kessler et al., 2019; Moshary et al., 2023; Chan, 2022). The key advantage of the conjoint design is that it lets us vary characteristics which are otherwise fixed in observational settings: vaccine effectiveness, side effects, delivery mode (needle or inhalation), convenience, and price (proxied by gift card rewards). We further elicit incentivized decisions over bonuses to be paid out during follow-up surveys that are potentially contingent on vaccinating. These decisions help us test the validity of the conjoint preference measures and give us an ability to randomize financial incentives for vaccination. Finally, we conduct follow-up surveys that allow us to measure vaccination decisions and experiences of side effects over time.

Our analysis begins by examining beliefs. We find, perhaps surprisingly, vaccine confident and hesitant participants hold 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.

Our first key result, however, is that vaccine confident and hesitant participants have sharply different beliefs about their own flu risk, vaccine effectiveness, and short-run vaccine side effects. Respondents

[3]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. The methods used here could be applied to other vaccination decisions.

intending to vaccinate believe they are twice as likely to get the flu (absent vaccination) as do those 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 50 percent higher than those not intending to vaccinate. While respondents who intend to vaccinate 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 over the next two weeks. Despite their differences, vaccine confident and vaccine hesitant alike hold beliefs that are roughly in line with what the CDC reports. This is possible because the CDC reports no information about short-run side effects, and 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 top of the range.

These belief differences matter for vaccination decisions only if people have strong preferences over vaccine effectiveness and short-run side effects. Our second result is that preferences for these characteristics are strong and heterogeneous, so effectiveness and side effects turn out to be highly influential for stated preferences over vaccines. Differences in preferences for these characteristics, and beliefs about them, account for most differences in vaccine intentions. For example, equalizing preferences for side effects and effectiveness, and beliefs about them, at the level of the most vaccine confident, would raise vaccination rates from 53 percent to 87 percent. Equalizing beliefs alone would have a smaller but still substantial effect, raising vaccination rates to 62 percent. We find that needle aversion and convenience are relatively unimportant. Residual factors, including scientifically unwarranted beliefs, are relevant for a smaller subset who have strong preferences against vaccination, but are not the central explanation for low vaccination rates.

Accounting for side effects and other undesirable aspects of vaccination also explains the apparently low value of statistical life (VSL) implied by vaccine demand. Ignoring all aspects of vaccines except their effectiveness implies a median VSL of $260,000, but accounting for costs of vaccination raises the implied median VSL to $6.8 million.

We estimate vaccine preferences using stated preference data, so a natural question is how meaningful those preferences are for actual vaccinations. We find that stated preferences and beliefs are highly predictive of individual variation in actual vaccination decisions and decisions in incentivized tasks, suggesting that the preferences we measure from the conjoint task have validity. Yet the correlations are not perfect and importantly the model-predicted vaccination rates are around five percentage-points higher than both the in-sample vaccination rate and an appropriately reweighted national-sample vaccination rate. Given these imperfect correlations, we interpret our preference results as reflecting vaccine intentions–the outcome of interest in much vaccine hesitancy research–and imperfectly reflecting follow-through. Thus, interventions suggested by our approach are likely complementary to nudges or other efforts to convert intentions to actions, but also help to clarify that these intention-action gaps likely explain a

quantitatively modest element of the low vaccination rates.

Our estimated preferences and beliefs imply a nuanced view of vaccine hesitancy. While we find that vaccine hesitancy is partially driven by pessimistic beliefs about vaccines, even under optimistic beliefs—the beliefs of the already vaccinated—it is still privately optimal not to vaccinate for 37 percent of participants, given their preferences for effectiveness and side effects. For these participants, raising vaccination rates through mandates or subsidies represents a real cost. External benefits—reducing the spread of infectious disease—may justify those costs. Our estimates are important for quantifying how large the external benefits need to be to justify interventions. For example, simulations indicate that increasing the vaccination rate from 53 percent to 80 percent would require subsidies in excess of $200.

Our results have implications for information provision and vaccine development. On the information provision side, we find that some participants are too pessimistic about side effect risks: they expect vaccination increases their experience of side effect symptoms, but experimentally-varied incentives to vaccinate actually reduce the experience of these symptoms. Simulations from our estimated demand system indicate that correcting these misperceptions would raise vaccination rates by about 5 percentage points overall, and 20 percentage points among marginally hesitant participants.

Vaccine development faces a tradeoff between effectiveness and tolerability. For example, recently-developed at-home nasal sprays offer convenient, needle-free vaccination, which would seem valuable given reports that needle phobia is an important source of vaccine hesitancy (e.g. Freeman et al. (2023)). However, these at-home nasal sprays have lower effectiveness (U.S. Food and Drug Administration, 2024), which our estimates suggest will be an important factor in vaccination decisions. There have also been recent policy shifts encourage developing vaccines with greater effectiveness but worse side effects (Subbaraman and Whyte, Subbaraman and Whyte). Our results provide a way of quantify the potential net effect of these investments given these tradeoffs in effectiveness and side effects.

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 recommendations, sibling vaccinations, or disease prevalence (for example, Abrevaya and Mulligan (2011); Lawler (2017); Oster (2018); 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); Alsan et al. (2019); Milkman et al. (2022); Jacobson 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); Chang et al. (2023); Campos-Mercade et al. (2024)). A third strand of this literature uses stated preference experiments and finds very low willingness to pay for vaccines (e.g. Steiner et al. (2002); Carlin et al. (2022)).

Our measures of preferences over vaccine characteristics and beliefs about effectiveness and side ef-

fects are novel to this literature and contribute in many ways. The literature lacks evidence on preferences over vaccine characteristics. While many studies elicit qualitative beliefs about vaccines, asking whether participants believe they work, they lack the quantitative evidence necessary to assess belief accuracy. Our finding of a key role for effectiveness and side effects helps explain why prior work has found that correcting scientifically inaccurate beliefs does not raise vaccine intentions (Nyhan et al., 2014; Nyhan and Reifler, 2015). Our findings also point to possibly more effective information interventions: ones that lower perceived side effects, raise perceived effectiveness, or raise preferences for effectiveness (for example, by increasing perception of disease risk). Our results can help guide vaccine development, because they show that a more effective vaccine could have higher take-up, even if accompanied by worse side effects. Our findings help explain the puzzlingly low willingness to pay for vaccines: it is due to high perceived side effect costs, not low value of effectiveness or health per se.

We also contribute to a rapidly growing literature on using survey experiments to measure economically important but difficult to observe objects such as beliefs and preferences. Early work observed the difficulty of separately identifying beliefs and preferences from choice data, and showed that surveys could be used to measure beliefs (e.g., Manski (2004)). Recent work has shown how subjective beliefs affect, for example, risky sexual behavior, home selling, consumption and savings, job switching, and stock trading (Delavande and Kohler, 2016; Bottan and Perez-Truglia, 2025; Chopra et al., 2025; Taubinsky et al., 2024; Jäger et al., 2024; Jiang et al., 2024). A related strand of the literature shows the value of survey experiments for measuring not only beliefs but also preferences, for example over social position or policies such as redistribution, trade, and climate change (Hvidberg et al., 2023; Dechezleprêtre et al., 2025; Stantcheva, 2023). Within this literature, our work is most closely related to Delavande (2008), who models contraceptive choice as a function of subjective beliefs about the effectiveness of different methods. Like Delavande, we show the importance of subjective beliefs about treatment characteristics. However, we go beyond effectiveness, showing that beliefs about side effects are highly important, and we document bias in side effect beliefs, implying opportunities for welfare-improving information interventions. For this reason, finally, our work also contributes to a literature on experience goods, where it is challenging to use choice data to measure preferences and beliefs about quality (Erdem and Keane, 1996; Israel, 2005; Dickstein, 2021; Allcott et al., 2025); we show the value of using direct belief elicitation to measure perceived quality.

2 Background

We study vaccines to protect against the influenza virus ("flu"), which annually infects roughly a billion people worldwide, including millions in the United States, where flu season runs from September through March (Centers for Disease Control and Prevention, 2024a; World Health Organization, 2025).

Although most cases are mild, a small fraction of infections are severe enough to warrant hospitalization. The CDC estimates that during the 2023-24 flu season, Americans experienced 40 million cases of symptomatic flu illness, resulting in 470,000 hospitalizations and 28,000 deaths (Center for Disease Control, 2025a). Importantly for individuals making vaccine decisions, 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.[4]

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—face higher infection risk. Thus family and occupational characteristics influence flu risk, as do personal behaviors such as frequency and mode of socialization.

Vaccination mitigates the harms of the flu but does not perfectly eliminate flu risk because the flu vaccine is not fully effective. Vaccine effectiveness refers to the percent reduction in flu illness among the vaccinated, relative to the control (Centers for Disease Control and Prevention, 2024b).[5] The CDC reports vaccine effectiveness rates by season. Reported values range from a low of 19 percent to a high of 60 percent, although most years are between 38 and 60 percent. 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 (Frutos, 2024). 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 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 recommended 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 steadily to about 45 percent by January, where it levels off. Although vaccination rates increase in age, they are below the CDC's target of 70 percent even among older adults. These low vaccination rates are

[4]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.

[5]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.

commonly interpreted as 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 beliefs about 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). Under the WHO SAGE definition, however vaccine hesitancy can reflect a misinformed decision based on, for example, incorrect views about vaccine safety or effectiveness. But it can also reflect a fully informed tradeoff: some people may be unwilling to accept moderate short-term side effects or other vaccine costs for imperfectly protection 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 effects, and trust in vaccination programs and health care providers (Larson et al., 2015). [6] The scale can classify people as hesitant for a range of reasons: if they refuse to vaccinate at all, but also if they (correctly) view themselves as low risk of disease, or judge a vaccine as ineffective based on a 60% effectiveness rate. Because these different sources of vaccine hesitancy reflect fundamentally different concerns, they call for different responses. To understand these concerns, we develop quantitative measures of beliefs, rather than qualitative agree/disagree scales, and we measure preferences over vaccine characteristics.

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 individuals' 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 brief follow-ups to measure vaccine take-up, experienced side effects, and experienced flu. We ran 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, and pre-registered, as discussed below.

We begin in wave 1 by asking a set of questions about demographics, health, trust in doctors and the government, and vaccine intentions. The health module includes questions about flu risk and vaccine contraindications. Blocks 2 and 3 contain our belief elicitation and stated preference tasks, which we de-

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

scribe 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). We then ask 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, as an information treatment. In the stated preference task, we ask respondents to choose between vaccine A and 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 with logically inconsistent responses (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 payable upon showing proof of vaccination in the follow-up survey. Bonus amounts vary across choices: $1.50, $3, $10, $20, and $50. For incentive compatibility, we explain that we will pick one of their answers and actually implement it. Responses to these choices 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 with a contraindication for vaccination (the latter group is excluded 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[7] 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 unvaccinated, non-contraindicated participants are eligible for wave 2. We sent multiple reminder messages encouraging completion of wave 2. We ask about experienced side effects (as we explain in Section 7) and vaccination. Then, for participants offered a reward to vaccinate, we ask for a record upload. To earn their reward, participants upload a pharmacy receipt or a screenshot of their vaccine record. Wave 3, opened two months after wave 1, involved the same set of participants as wave 2. Here we again measure vaccination, and also ask about flu experiences. To avoid experimenter demand effects, we obfuscate our intentions (as in, for example, Haaland and Roth (2020)) by opening with questions about pharmacy shopping and include questions about non-vaccine medications, and non-flu illness.

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

3.2 Managing induced demand and expressive answering

Two natural concerns with surveys are that participants may respond to perceived experimenter demand or survey enumerator effects, and that they may answer questions, particularly belief questions, to express their identity rather than to fully communicate their actual beliefs. An online, anonymous survey like ours 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 perceive vaccinations as especially sensitive subject. Nonetheless we address these concerns in several ways.

In the experiment design, we followed best practices to minimize induced demand (Stantcheva, 2023; Haaland et al., 2023). Our recruitment material did not mention the study's purpose but emphasized that we are "non-partisan researchers."[8] 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 used unrelated questions to minimize the connection to our earlier survey, and indeed virtually no respondents thought the survey was about vaccine preferences or beliefs 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, if participants were answering expressively, we would 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 essentially uncorrelated with vaccine intentions. Finally, we show that beliefs and preferences are individually predictive of actual vaccinations (Appendix Table A.3), and beliefs about personal flu risk are predictive of actual flu incidence (Appendix Figure A.2), meaning that our belief and preference measures capture distinct, behaviorally-relevant information which corresponds to choices and experiences.

3.3 Analysis sample

We limit the participant 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, like other online samples, skews young, we imposed an age quota requiring that 10 percent of our sample be 65 or older.[9]

Appendix Table A.1 shows how the sample size changes with successive inclusion criteria. 3,447 par-

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

[9] Filling this quota was difficult because our elder subjects took longer to complete the survey and so saw a lower hourly reward. We increased compensation for this group accordingly.

ticipants started wave 1. Our main analysis sample excludes respondents with any failed attention check or inconsistent stated preference,[10] multiple survey attempts, and non-missing key variables. We further exclude the 9 percent of participants who report a contraindication to the flu vaccine. We prespecified that we would exclude these participants because we cannot ethically offer them incentives to vaccinate. However 9 percent is a high contraindication rate; rates of anaphylaxis following vaccination are approximately 1.3 per million (McNeil et al., 2016). It is likely that some participants are reporting concerns about vaccine safety rather than actual contraindications. We therefore conducted a sensitivity analysis where we replicated our main analyses, including participants reporting contraindications. Our qualitative and quantitative conclusions are quite similar; see Appendix B.

After excluding participants reporting contraindications, our main analysis sample consists of the remaining 2,357 participants. Of these, 1,757 are unvaccinated at baseline 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 82 percent at wave 3 (about two months later).

We summarize the demographic characteristics of the main analysis sample in Table 1. 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 younger, higher educated, and higher income participants, and our sample's vaccination rate is a few percentage points higher than the national average. However, we show below (Table 6) that after accounting for contraindications and these different demographics, the vaccination rate in our sample becomes similar to the vaccination rate in BRFSS. Thus our sample is roughly representative on this key metric. 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 our analysis around participants' 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 a comparable 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]

[10] 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.

[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 questions about perceived risks and effectiveness of vaccines. We do not ask those questions because they are redundant with our 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.2.

[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 over and beliefs about vaccines, as well as biases in perceived side effects) and our interventions. It also describes our sample size rule, which was dictated by our budget. We pre-registered 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. required decisions we failed to anticipate).

We emphasize here two departures from the pre-registration. First, wave 3 was not preregistered. Two factors motivated this additional wave. We realized we had an opportunity to measure whether flu experiences aligned with flu beliefs. Beyond that, 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 unrelated 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.

Our second departure is that we simulate choice probabilities using mean beliefs by intention group, rather than individual level beliefs. This simulation makes it straightforward to decompose differences in vaccination rates across intention groups into differences in preferences and beliefs.[14]

4 Beliefs about flu risk and vaccine characteristics

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

4.1 Measuring beliefs about flu risk and vaccine effects

Measurement target: Our goal is to measure people's 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

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

[14] The problem stems from the nonlinearity of vaccination rates with respect to beliefs, and differences in the belief distribution across groups. Roughly, the decomposition requires that we swap the distribution of beliefs between intention groups, but hold fixed the conditional distribution of preferences given intentions and beliefs. However, the preference distributions do not have common support, so the conditional preference distribution is not identified.

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 instead 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.[15]

Measuring beliefs about flu risk: The exact text of all our elicitations is in Appendix C; 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.**" (note the face-saving language, ala Stantcheva, 2023). Then we ask about likelihood of 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.[16] 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-

[15] 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.

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

ask our questions, but asking for the "percent chance" 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 conditional 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). However, this definition does not perfectly capture effectiveness, which typically measures reductions in the experience of particular flu consequences (symptomatic illness, hospitalization, or mortality).

Eliciting beliefs about side effect rates is more challenging. Side effects rates, like effectiveness, are causal contrasts, but the language "side effect rate" does not evoke a causal contrast, and may refer to the expected level of an outcome.[17] In our pilot studies, we found that asking about side effect rates of vaccination produced very high side effect rates.[18] 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

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

[18] 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.

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?

We again give a multiple choice version of the question, and then an open numerical response version. 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 $rate_i(1) - rate_i(0)$.

Objective benchmarks: While our main interest is on beliefs about personal risks and vaccine consequences ("subjective beliefs"), we also elicit beliefs about the 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.

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. [19]

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 mea-

[19] The alternative approach to inferring beliefs—recovering them from choices—requires a rational expectations assumption to separate beliefs from preferences (Manski, 2004). This assumption is especially strong in the vaccine context. 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.

sures 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 average 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.3, 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.[20] We also show in Appendix Table A.3 that beliefs strongly predict vaccination rates, in the way we would expect, showing there is some signal in the belief measures. This table also shows that beliefs predict vaccination rates conditional on our preference estimates,[21] indicating that beliefs and preferences are distinct objects, each with their own explanatory power. We further show in Appendix D that beliefs are predicted by individual level risk factors (age and chronic conditions), although small sample sizes make these predictions noisy at times. 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 (Appendix Figure A.2).

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).[22] 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 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 distinct from intentions.

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

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

[21] 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.

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

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 averages in two important ways. First, people generally 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 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 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 already vaccinated 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, 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 incentivized, 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 fully 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 is 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, they do not describe an overall side effect rate. Instead, they list the rate of many individual side effects, typically 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 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.[23]

Finally, to operationalize convenience, we ask participants to imagine that the vaccines are available either at a "usual" location, "wherever people in your community usually get vaccinated," or at a "nearby" location, where "there is a vaccination campaign happening near your home, with appointments available at all hours" (the convenient option). We note that this approach captures location and scheduling aspects of convenience, but not all aspects. For example, a person may find it especially convenient to vaccinate if they are already in a health care setting, and indeed bundling pediatric vaccinations with pediatric primary care appointments has a substantial impact on vaccination rates (Worsham et al., 2020; Anderson et al., ming).

We give each participant the same 10 choice sets, in a random order. We generated choice sets to maximize D-efficiency.[24] Appendix Table A.4 reports the options available in each choice set, and Appendix Figure A.4 shows how we present the choices to the participants. Before beginning the conjoint, we explain the meaning of each characteristic. In case respondents forget, we include click-to-expand explainers of the characteristic in each choice scenario. Finally, before the conjoint, we include a baseline choice that resembles the prevailing vaccine. The baseline choice asks participants to choose between a vaccine with

[23] 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.

[24] D-efficiency is a measure of design efficiency that captures expected precision of the experiment. We use the MktEx Macro in SAS to maximize D-efficiency.

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 $j \neq 0$ in choice set t:

$$u_{ijt} = -\alpha_i \text{Reward}_{ijt} + \beta_{VE,i} \text{VE}_{ijt} - \beta_{SE,i} \text{SE}_{ijt} \tag{1}$$

$$+ \beta_{N,i} \text{Needle}_{ijt} + \beta_{C,i} \text{Convenient}_{ijt} + \beta_{0i} + \epsilon_{ijt}, \tag{2}$$

$$u_{i0t} = \epsilon_{i0t}. \tag{3}$$

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:

$$(\ln \alpha, \ln \beta_{VE}, \ln \beta_{SE}, \beta_N, \beta_C, \beta_0) \sim N(\mu_i, \Sigma), \epsilon_{ijk} \sim T1EV. \tag{4}$$

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.

Our primary, pre-registered specification for μ_i lets μ_i vary with age (above/below 55), education (less than college, college or more), and vaccine hesitancy (above/below median on our totaled score). In robustness tests we let μ_i vary directly with vaccine intentions, as well as a richer set of covariates.

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, resulting in $N_d (= 900)$ draws from the posterior distribution of individual-specific preference parameters. We denote draw d for individual i $\theta_{id} = (\alpha_{id}, \beta_{VE,id}, \ldots, \beta_{0,id})$. We use these extensively in our simulations. We report diagnostics on the resulting Markov chains in Appendix E; 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. This simple cut was not pre-registered. 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.[25] 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 people 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 those not intending to vaccinate at baseline.

5.4 Preference estimates

We find strong preferences for effectiveness and side effects, and, perhaps surprisingly, weaker preferences for needles and convenience. We report the distribution of preferences in Table 2. On average, participants are willing to tolerate one percentage point greater side effect if it increases effectiveness by 2.15 percentage points. However there is a great deal of heterogeneity in this trade-off; for example, the median participant requires only a 0.6 percentage point increase in effectiveness to tolerate an additional

[25] The choice sets in this figure are 0, 9, and 8. Appendix Table A.4 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.

percentage point likelihood of side effects. By contrast, participants are essentially unwilling to trade effectiveness or side effects for needles or convenience. At the mean, the coefficient on needle is positive, and equivalent to half a percentage point of effectiveness. The coefficient on convenient is essentially zero.

As our interest is in cross-group differences in preferences, we report median preferences by group in Table 3. 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 willingness to pay for effectiveness is less than a quarter as large. All groups are sensitive to side effects. Vaccine confident have a higher willingness-to-pay to avoid side effects than do vaccine hesitant. However, vaccine hesitant are most willing to trade off effectiveness to avoid side effects. No group is especially sensitive to needles or convenience.[26] The "do not intend to vaccinate" group is the only group with a negative average sensitivity to needles, and this sensitivity is small: the median participants in this group is on average indifferent between a needle and a $2.3 reward, or a a 3 percentage point increase in vaccine effectiveness.[27] Likewise for convenience: for the most sensitive group, a convenient vaccine is equivalent to a roughly $0.46 at the median.

Interestingly, we find relatively small differences in the value of the "inside option" across intention groups. All groups place a large, negative weight on the inside option coefficient, indicating a general baseline vaccine aversion. The weight is actually most negative for vaccine confident rather than vaccine hesitant. This baseline vaccine aversion therefore likely does not reflect extreme views about vaccines (such as a concern that they cause cancer), because such views are likely more common among vaccine hesitant.

5.5 Explaining low VSLs and vaccine hesitancy

Rationalizing low VSLs: Our analysis of preferences for vaccine characteristics helps 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,

[26] 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 convenience. 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.

[27] Some qualitative evidence helps explain the apparently low rate of needlephobia. One participant expressed discomfort with needles, but also said they liked a nasal spray even less.

for example side effects or needle aversion.

As in prior work, we find that WTP for vaccines implies low VSLs in simple calculations. Accounting for vaccine characteristics beyond effectiveness 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, i.e. $X_i'\beta_i/\alpha$. Here we evaluate utility at perceived characteristics, setting effectiveness and side effects to the average belief in each intention group.[28] Median WTP declines from $254 for the already vaccinated to -230 for those who do not intend to vaccinate, with an overall median of $10.2. 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:[29]

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

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 $260,000, much lower than values commonly used by policy makers or found in the literature.[30]

This apparently low VSL, however, is not driven by low value of life, but by aversion to side effects and other vaccine characteristics, as well as low perceived flu risk. To show this, we calculate WTP for full vaccine effectiveness ($100\beta_{VE}/\alpha$), and report its median in column (3). Median WTP for full effectiveness is $410 overall. This implies a median VSL of $6.8 million, which ranges from $10 million for the most vaccine confident to $3.1 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.[31]

Preferences for and beliefs about side effects and effectiveness explain vaccine intentions: We use

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

[29] 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.

[30] 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`.

[31] An alternative interpretation of the low VSL among vaccine hesitant is that it is driven by low flu risk, given their age. We assume that flu risk is homogenous within age, but if vaccine hesitant participants have low flu risk given their age, we could understate their VSL.

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 group's mean, X_g. [32] 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 the first column of Table 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 5 percent.

In the remaining panels, we decompose these large cross-group differences in vaccination rates into differences in preferences, beliefs, and their interactions. In column (2) we change preferences for effectiveness and side effects, so that the marginal prior distribution of β_{VE} and β_{SE} in each group matches group 1's distribution.[33] Changing these preference parameters results in much higher vaccination rates for the vaccine hesitant, raising the marginally hesitant's rate by 33 percentage points, and the strongly hesitant by 12 percentage point. Preferences over effectiveness and side effects play a key role in vaccine intentions.

In the next two columns, we show that beliefs also play a key role, and that together beliefs about and preferences over side effects and effectiveness explain vaccine intentions. In column (3) we give each group

[32] We use group mean beliefs rather than individual beliefs to facilitate a straightforward decomposition of group differences in vaccination rates into preferences and beliefs.

[33] To do so, we shift the log of β_{VE} by $\overline{\ln \beta_{VE1}} - \overline{\ln \beta_{VEg}}$, then exponentiate back. Here $\overline{\ln \beta_{VEg}}$ is the group average log β_{VE}. We shift β_{SE} similarly. Because the posterior distribution is approximately log normal, this shift means that the entire distributions match.

the same beliefs about effectiveness and side effects as the "already vaccinated", but let their preferences differ. Changing beliefs also raises vaccination rates, although the overall effect is not as large as changing preferences alone. Finally in column (4) we change both beliefs and preferences, and the result is that all groups' vaccination rates exceed 75 percent. Beliefs and preferences are complementary because vaccine confident participants both value effectiveness more, and believe that it is greater.

Thus differences in preferences over and beliefs about effectiveness and side effects explain vaccine intentions. This leaves a quantitatively small role for the remaining preference parameters, which govern preferences over needles, convenience, and the inside option. We show this explicitly in the remaining column, where we allow beliefs and β_{VE} and β_{SE} to differ by group, but set the distribution of β_0, β_N and β_C to the "already vaccinated" group's. Doing so actually reduces vaccination rates slightly. The reason for this, quantitatively, is that these coefficients are fairly similar across groups, and indeed the inside option, β_0, is actually most negative at the median for the "already vaccinated" group. This last result implies that scientifically unwarranted beliefs, or unmodelled vaccine animosity, play a quantitatively minor role in explaining vaccine intentions, because such beliefs or animosity would load onto β_0.

This decomposition is not sensitive to the specification of the utility function. Specifically, letting preference parameters vary directly with intention groups, or using richer sets of covariates, produces very similar preference decomposition, as we show in Appendix Table A.5.

Our results show a much larger role for vaccine effectiveness and side effects than for convenience or needles in explaining vaccine hesitancy. One might worry that this is an artifact of our survey design, since we elicited beliefs about effectiveness and side effects but not convenience or needles. However, our randomization of survey blocks allows us to rule out this concern. We report our main simulations, but stratified on randomly assigned survey order (beliefs first or vaccine choices first). If asking about beliefs had a large effect on preferences, we would see differences in simulated effects of changing preferences or characteristics. However, Appendix Table A.6 shows that the simulation results are identical across survey order.

6 Stated preferences, intentions, and revealed preferences

So far we have used conjoint data to show that vaccine preferences depend strongly on effectiveness and side effects, but not on convenience or needles. Preferences over effectiveness and side effects, and beliefs about these objects, in turn are highly predictive of vaccine intentions. Of course the results so far concern only stated preferences and intentions, not necessarily revealed preference and follow-through. Here we show that our stated preferences are highly but imperfectly predictive of actual vaccination behavior and incentivized choices, and price sensitivity parameters are informative for but imperfectly predictive of choices in incentivized tasks.

6.1 Predicted and actual vaccination rates

We first show the predictive power of simulated vaccination rates for actual vaccination rates. The main result here is that simulated vaccination rates are highly but imperfectly predictive. To show this, Figure 5 shows a binned scatter plot of actual vaccination rates against model-predicted vaccination rates, along with the linear fit and distribution of model-predicted vaccination rates (both at the individual level). To obtain model predictions, we simulate vaccination rates using group-mean beliefs and individual preferences, ignoring incentives to vaccinate, and we use the individual posterior mean vaccination rate. Actual vaccination is equal to 1 for participants who reported being vaccinated at the start of the experiment, and otherwise we use wave 3 vaccination (treating missing values as random). A large mass of participants have a near-zero simulated vaccination rate, roughly 7 percent of them actually vaccinate. As simulated vaccination rates rise, actual vaccination rates also rise quickly, but not one-for-one. Among the many participants with a model-predicted vaccination rates near 1, actual vaccination rates are 70-80 percent. The overall slope of the relationship is 0.72.

Model-predicted vaccination rates are higher than actual vaccination rates, for both our Prolific sample and for nationally representative samples, as Table 6 show. The first row of the table reports the overall model-predicted vaccination rate, scaled by the fraction of people in our sample without contraindications for the flu vaccine. The next row reports the actual vaccination rate in our analysis sample after expanding to include participants with contraindications. The sample rate is 5 percentage points lower than the simulated rate. The sample vaccination rate is similar to the 2023 BRFSS vaccination rate, a nationally representative sample used to measure vaccination rates. This similarity is perhaps surprising because our sample is not representative; after re-weighting the BRFSS to match our sample's age-income-education-sex distribution, we find slightly lower vaccination rates. Thus model-predicted vaccination rates correlate with but overstate actual vaccination rates.

6.2 Incentivized choices

Stated preferences predict individual vaccination choices strongly but imperfectly. This does not necessarily imply, however, that individual preference parameters are informative about real-stakes choices. Here we show that, indeed, stated preferences data on reward sensitivity are informative about reward sensitivity in an incentivized task, although the quantitative predictions are not accurate.

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.[34]

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 our logit utility framework, the value of the option when the bonus is b and beliefs are X_g is: [35]

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

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 a vector of parameters 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. At high incentives, simulated and actual take up of incentives are similar, but at low incentives, the model substantially overstates take-up. Actual incentive sensitivity is much greater than what the stated preference data imply.

Despite the poor performance in predicting incentive take-up rates, 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 sensitivity. If the stated preferences are informative for choices in an incentivized context, we would expect

[34] 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.

[35] 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.

to see a positive relationship between take-up of the reward and the individual's predicted likelihood of vaccinating from the model ($\gamma_{1r} > 0$ for all r), since the reward option is more appealing if you are likely to get vaccinated. If stated preferences are informative, then estimated sensitivity to rewards from the conjoint task should also have a positive relationship with taking up the reward in this task and we would expect 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 7 reports the estimated γs for each reward b.γ_{1r} is large and highly significant for all rewards. The weight on the posterior mean for responsiveness to rewards (γ_{2r}), 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.

7 Implications for information provision and vaccine design

7.1 Correcting side effect beliefs would raise vaccination rates

We show that some participants are too pessimistic in their beliefs about side effects, and correcting this belief would raise vaccination rates. To show excess pessimism, we must show that participants' beliefs are inconsistent with their own experiences. In the case of side effects, this means that their believed treatment effect of vaccination on side effects is inconsistent with their actual treatment effect. We elicit the believed treatment effect directly. We therefore estimate vaccine effects and compare to participant expectations.

To estimate the effect of vaccination on experienced side effects, we rely on 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 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.[36]

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 two-stage least squares regressions using these switchers:

$$SideEffect_i = \beta_0 + \beta_1 \hat{Vacc}_i + \epsilon_i$$
$$Vacc_i = \alpha_0 + \alpha_1 incentive_i + \varepsilon_i.$$

We estimate separate specifications for the $10 switchers and $20 switchers, as well as pooled regressions.[37] Our pre-specified outcome is wave 2 experience of severe side effect symptoms, measured as experience of side effect symptoms over the prior two weeks, severe enough to disrupt daily activity. This measure corresponds exactly to the expectation we elicited, but it is backward rather than forward looking. We pre-specified we would consider two measures of vaccinations, self-reported and verified. We report the wave 1 expected side effect rate, for each group, say $\overline{belief}$. To test the accuracy of beliefs, we test the null hypothesis that $\beta_1 = \overline{belief}$.

Table 8 shows the results. We start with the reduced form in the first three rows. The incentives increase self-reported vaccination rates by 6-10 percentage points in wave 2, but they increase verified vaccination rates only by 3-4 percentage points. We were concerned that participants misreported their vaccination status to earn a reward (though we were clear they would have to provide verification), so in wave 3 we again elicited vaccination rates, with no link to the incentive (not a pre-registered outcome). Here the effect is actually larger than the wave 2 self-reported effect, so we conclude that the incentives raised vaccination rates but people had trouble verifying their status.[38] We note, though, that the first stage in our IV system is weak; the F stats here are 6 or smaller.

The fourth row of the table shows our core result: we estimate that incentives to vaccinate actually reduce experienced side effect symptoms. The estimates are large, negative, and statistically significant, with reductions of 13 to 25 percentage points. These estimates, combined with the weak first stage, imply very large negative effects of vaccination on side effects, more than 100 percentage points. Despite the large standard errors, we can reject the hypothesis that beliefs are accurate. Given the weak first stage and implausible size, we view the exact estimates with skepticism. We report them because we pre-specified that we would do so. However we interpret the results as showing that vaccination does not worsen the

[36] 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 in robustness tests, we also estimate specifications with the fuller set of participants who were randomized to incentives.

[37] 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.

[38] Consistent with this interpretation of generally honest reporting, we found very few instances of apparently fraudulent verifications - four verifications appeared to have been manipulated, incorrect for the circumstance, or AI-generated.

experience of side effects, contrary to beliefs.

Our main finding here - that incentives to vaccinate both increase vaccination and reduce the experience of side effect symptoms - is robust to alternative specifications which are more powerful or better controlled. In particular, we estimate regressions of the following form:

$$y_i = \gamma_0 + \gamma_1 HighIncentive_i + \sum_j \gamma_j 1\left\{\text{Choice}_j = 1\right\}_i + \eta_i.$$

Here our outcomes are vaccination or experience of side effect and $HighIncentive_i$ is an indicator equal to 1 if the incentive selected for i is at least equal to their switching point. This incentive offer is random conditional on the sequence of incentivized choices made by i, so we control for indicators for each of these choices. This specification uses all the variation in our data because it lets us include the large number of people who switch at incentives other than $10 or $20. Appendix Table A.8 shows the results from this alternative specification are quite similar to our baseline result. One final concern is that the incentives may induce people to take the follow-up survey sooner, and the survey timing may independently affect side effect experience. Our estimates however are also robust to controlling for wave 2 survey date.

Thus participants overestimate their own side effect risk. As we have found that vaccine choices are sensitive to perceived side effect, our model implies that correcting these beliefs would raise flu vaccination rates substantially. In Appendix Table A.7, we simulate vaccination rates at baseline and after setting believed side effect rates to zero. Correcting misperceptions in this way raises the overall vaccination rate by 4.6 percentage points, with large gains among the marginally hesitant. Information provision is therefore a potential way to raise vaccination rates.

7.2 Cost of increasing vaccination rates

Our results are also useful for understanding the potential costs of increasing vaccination rates. We determine these costs by simulating vaccination rates under different scenarios and subsidy amounts, in Figure 7. The "base" scenario uses our baseline preference estimates and beliefs. The "no side effect" scenario sets believed side effects to zero, and the 75 percent effective scenario sets believed effectiveness to 75 percent.

The figure shows several important points. First, achieving high vaccination rates is expensive, in fiscal terms and efficiency terms. While subsidies can raise vaccination rates, achieving a high vaccination rate of, say, 80 percent, would require a subsidy of over $200, for a subsidy cost of $160 per person. The deadweight loss from such a subsidy is roughly $27 per person.[39] This calculation of course ignores the positive externalities of increased vaccination; the efficiency cost here should be weighted against the external benefits of vaccinations. Second, subsidies experience diminishing returns. For example, the

[39] This number is based on a Harberger type approximation: $0.5 \cdot \$200 \cdot 0.27$, where 0.27 is the increase in vaccination rate.

first $50 in subsidies raises vaccination rates by about 12 percentage points, but the next $50 only by 8 percentage points. There is a long tail of vaccine aversion, so that even with subsidies of $500, roughly 10 percent of people decline to vaccinate. Finally, improving vaccines via reduced side effects or increased effectiveness does not substantially change the shape of the relationship between subsidy and take-up. The long tail of vaccine aversion remains even with no side effects.

7.3 Implications for vaccine design

Our results give some guidance to vaccine developers and policymakers in navigating the trade-off between effectiveness and tolerability of vaccines. This trade-off comes up, for example, in the development of an at-home, nasal-spray based flu vaccine, which is convenient and needle-free, but less effective (U.S. Food and Drug Administration, 2024). As we find a much stronger preference for effectiveness than for convenience or nasal (relative to needle), our results suggest that at-home, nasal-spray based vaccines with low effectiveness are unlikely to raise vaccination rates, although they may be helpful as a mid-season booster.

The implications for trading off effectiveness and side effects are less obvious, because we find strong and heterogeneous preferences for both characteristics. Understanding this tradeoff is important because recent policy proposals prioritize developing a universal flu vaccine using whole virus particles; such vaccines can be more effective but also produce worse side effects (Subbaraman and Whyte, Subbaraman and Whyte). How the trade-off plays out plainly depends on the effectiveness and side effects of developed vaccines.

To illustrate the trade-off in a general way, we plot the level sets of the simulated vaccination rate as a function of side effects and effectiveness in Figure 8, panel A. We use our baseline preferences estimates and assume correct beliefs, since we lack belief data for counterfactual vaccines. The current flu vaccine has roughly a 1 percent side effect and 50 percent effectiveness, for a vaccine rate of about 50 percent and an immunization rate of 25 percent. A more effective vaccine with worse side effects produces a greater vaccination rate only if the increase in side effects is about a third as great as the increase in effectiveness. For example, an effectiveness of 65 percent would increase vaccination rates to 60 percent at a side effect rate of 2.5 percent, but would reduce vaccination rates to 45 percent at a side effect rate of 9 percent.

Of course, a higher effectiveness confers greater protection at a given vaccination rate, so we might prefer to focus on effective immunity rather than the vaccination rate. Panel B plots level sets for "effective immunity", meaning the product of vaccination rate and vaccine effectiveness. This panel shows that the level sets for immunity are fairly steep in the side effect rate, but nonetheless single-digit side effect increases can offset the gains in immunization. For example, increasing effectiveness from 50 to 55 percent also increases immunity as long as side effects do not increase to 4 percentage points or greater.

8 Conclusions

We use an online experiment to investigate the puzzle of low take-up of the flu-vaccine, despite modest effectiveness, zero-cost, and substantial flu illness. We find that many people do not intend to vaccinate, not because they do not value their health, not because they believe the vaccine produces long-run harms, and not because they believe it is totally ineffective, but because they believe the flu vaccines effectiveness is not high enough to justify its cost in terms of short-run side effects. Accounting for these perceived costs of vaccination, we find that vaccine preferences imply VSLs in line with typical estimates. Critically, low demand for vaccines reflects not only preferences but also beliefs. If people vaccine hesitant participants shared the beliefs about effectiveness and side effects of already vaccinated participants, their vaccination rates would be 15-25 percentage points higher.

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. Also, our conjoint experiment was designed to measure preferences over effectiveness and side effects, but not to measure how disease risk affects these preferences—an important question we leave to future work. Finally, much of our results rely on stated preferences and beliefs. Such data are informative for intentions but imperfectly translate to action.

Despite these limitations, our results have several implications for vaccine policy and research. We find that side effect costs are especially high for vaccine hesitant participants, both because they have a strong aversion to side effects and because they believe that the side effect rate is high—higher than it turns out to be when we experimentally induce them to vaccinate. Thus, correcting misperceptions of side effects may be an especially effective intervention to raise vaccination rates. As we find that many participants value vaccine effectiveness, even more so than side effects, developing vaccines with high effectiveness–even if they also have worse side effects–can increase vaccination rates as well as raise immunity among the vaccinated. A natural direction for research would be to investiage interventions to change beliefs about effectiveness and side effects—it is unclear what interventions can durably shift beliefs, but our model suggests that doing so can shift vaccination rates.

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*.

Akbarpour, M., E. Budish, P. Dworczak, and S. D. Kominers (2024). An economic framework for vaccine prioritization. *The Quarterly Journal of Economics 139*(1), 359–417.

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., J. C. Castillo, M. Gentzkow, L. Musolff, and T. Salz (2025). Sources of market power in web search: Evidence from a field experiment. 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.

Alsan, M., O. Garrick, and G. Graziani (2019). Does diversity matter for health? experimental evidence from oakland. *American Economic Review 109*(12), 4071–4111.

Anderson, M. L., C. Dobkin, D. Gorry, and H.-F. Tseng (forthcoming). Real-world effectiveness of the influenza vaccine in young children. *American Economic Journal: Applied Economics*.

Arteaga, F., A. J. Kapor, C. A. Neilson, and S. D. Zimmerman (2022). Smart matching platforms and heterogeneous beliefs in centralized school choice. *The Quarterly Journal of Economics 137*(3), 1791–1848.

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.

Bates, M., M. Dinerstein, A. C. Johnston, and I. Sorkin (2025). Teacher labor market policy and the theory of the second best. *The Quarterly Journal of Economics 140*(2), 1417–1469.

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.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Bottan, N. and R. Perez-Truglia (2025). Betting on the house: Subjective expectations and market choices. *American Economic Journal: Applied Economics 17*(1), 459–500.

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*.

Chang, T. Y., M. Jacobson, M. Shah, M. Kopetsky, R. Pramanik, and S. B. Shah (2023). Reminders, but not monetary incentives, increase covid-19 booster uptake. *Proceedings of the National Academy of Sciences 120*(31), e2302725120.

Chopra, F., C. Roth, and J. Wohlfart (2025). Home price expectations and spending: Evidence from a field experiment. *American Economic Review 115*(7), 2267–2305.

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*.

Dechezleprêtre, A., A. Fabre, T. Kruse, B. Planterose, A. Sanchez Chico, and S. Stantcheva (2025). Fighting climate change: International attitudes toward climate policies. *American Economic Review 115*(4), 1258–1300.

Delavande, A. (2008). Pill, patch, or shot? subjective expectations and birth control choice. *International Economic Review 49*(3), 999–1042.

Delavande, A. and H.-P. Kohler (2016). Hiv/aids-related expectations and risky sexual behaviour in malawi. *The Review of Economic Studies 83*(1), 118–164.

Dickstein, M. J. (2021). Efficient provision of experience goods: Evidence from antidepressant choice.

Dickstein, M. J. and E. Morales (2018). What do exporters know? *The Quarterly Journal of Economics 133*(4), 1753–1801.

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

Erdem, T. and M. P. Keane (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing science 15*(1), 1–20.

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.

Frutos, A. M. (2024). Interim estimates of 2023–24 seasonal influenza vaccine effectiveness—united states. *MMWR. Morbidity and mortality weekly report 73*.

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. and C. Roth (2020). Labor market concerns and support for immigration. *Journal of Public Economics 191*, 104256.

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

Heckman, J. J., S. Urzua, and E. Vytlacil (2006). Understanding instrumental variables in models with essential heterogeneity. *The review of economics and statistics 88*(3), 389–432.

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.

Hvidberg, K. B., C. T. Kreiner, and S. Stantcheva (2023). Social positions and fairness views on inequality. *Review of Economic Studies 90*(6), 3083–3118.

Israel, M. (2005). Services as experience goods: An empirical examination of consumer learning in automobile insurance. *American Economic Review 95*(5), 1444–1463.

Jacobson, M., T. Y. Chang, M. Shah, R. Pramanik, and S. B. Shah (2022). Can financial incentives and other nudges increase covid-19 vaccinations among the vaccine hesitant? a randomized trial. *Vaccine 40*(43), 6235–6242.

Jäger, S., C. Roth, N. Roussille, and B. Schoefer (2024). Worker beliefs about outside options. *The Quarterly Journal of Economics 139*(3), 1505–1556.

Jiang, Z., H. Liu, C. Peng, and H. Yan (2024). Investor memory and biased beliefs: Evidence from the field. Technical report, National Bureau of Economic Research.

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). Counterstereotypical 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.

McNeil, M. M., E. S. Weintraub, J. Duffy, L. Sukumaran, S. J. Jacobsen, N. P. Klein, S. J. Hambidge, G. M. Lee, L. A. Jackson, S. A. Irving, et al. (2016). Risk of anaphylaxis after vaccination in children and adults. *Journal of Allergy and Clinical Immunology 137*(3), 868–878.

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.

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

Oster, E. (2018). Does disease cause vaccination? disease outbreaks and vaccination response. *Journal of health economics 57*, 90–101.

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.

Steiner, M., L. C. Vermeulen, J. Mullahy, and M. S. Hayney (2002). Factors influencing decisions regarding influenza vaccination and treatment: a survey of healthcare workers. *Infection Control & Hospital Epidemiology 23*(10), 625–627.

Subbaraman, N. and L. E. Whyte. RFK Jr. Bets $500 Million on Universal Vaccines in Shift From Covid-19 Funding. *The Wall Street Journal*. Updated 2:49 pm ET.

Taubinsky, D., L. Butera, M. Saccarola, and C. Lian (2024). Beliefs about the economy are excessively sensitive to household-level shocks: Evidence from linked survey and administrative data. Technical report, National Bureau of Economic Research.

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.

Worsham, C., J. Woo, and A. B. Jena (2020). Birth month and influenza vaccination in children. *New England Journal of Medicine 383*(2), 184–185.

Table 1: 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
High school or less	0.11	0.27
Some college	0.27	0.30
College or more	0.61	0.31
Income < $25,000	0.12	0.12
Income $25,000 - $50,000	0.20	0.19
Income $50,000 - $100,00	0.36	0.22
Income > $100,00	0.30	0.44
Already vaccinated	0.25	NA
Intend to vaccinate	0.24	NA
May or may not vaccinate	0.21	NA
Do not intend to vaccinate	0.30	NA
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 2: Preference estimates

Preference parameter	Mean	SD	Median	2.5 percentile	97.5 percentile
Sensitivity to...					
Reward	0.13	0.25	0.06	0.01	0.7
Vaccine Effectiveness	0.45	0.62	0.25	0.03	2.14
Side effect rate	-1.32	4.9	-0.4	-7.89	-0.02
Needle	0.27	2.12	0.27	-3.96	4.46
Convenient	0.04	0.79	0.04	-1.51	1.61
Inside option	-17.38	14.08	-17.31	-45.06	10.05
MRS Effectiveness for side effects	2.15	7.36	0.62	0.03	13.54

Notes: Table reports the posterior mean, standard deviation, median, and tail percentiles for the coefficients on the indicated vaccine characteristics, from Equation 1. The "inside option" is the coefficient for vaccination relative to not vaccinating. The MRS row reports the distribution of the marginal rate of substitution between effectiveness and side effects, $-\beta_{VE}/\beta_{SE}$.

Table 3: Median preferences parameters by vaccine intention group

Parameter:	Reward	Willingness to pay for...				
	Sensitivity	Effectiveness	Side effects	Needle	Convenience	Inside opt
Already vaccinated	0.0467	8.77	-8.89	8.78	-0.0104	-279
Intend to vaccinate	0.053	6.41	-7.88	7.27	0.329	-222
May or may not vaccinate	0.0655	3.58	-6.31	1.5	0.461	-172
Do not intend to vaccinate	0.0813	1.69	-4.24	-2.33	0.403	-199

Notes: Table reports the posterior median of each preference parameter, by vaccine intention group. Willingness to pay is the coefficient for a characteristic is the coefficient on that characteristic divided by reward sensitivity. The final column, "inside option", can be interpreted as the willigness-to-pay for a vaccine with no effectiveness, no side effects, no needle, and not convenient.

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 vaccinated	254	5.52	877	10
Intent to vaccinate	169	5.67	641	10.7
May or may not vaccinate	-26.4	-1.13	358	7.93
Do not intend to vaccinate	-230	-8.07	169	3.1
All	10.2	0.26	410	6.78

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: Vaccine intentions are explained by preferences over and beliefs about effectiveness and side effects

Scenario	Simulated vaccination rate for...				
	Already Vacced	Intend to vacc	May or may not	Don't intend	All
	(1)	(2)	(3)	(4)	(5)
Baseline	0.89	0.85	0.38	0.051	0.53
Equalize prefs for VE and SE	0.89	0.93	0.71	0.18	0.65
Equalize beliefs about VE and SE	0.89	0.85	0.64	0.19	0.62
Equalize prefs and beliefs about VE and SE	0.89	0.93	0.92	0.76	0.87
Equalize prefs for needle, conv., vacc intercept	0.89	0.81	0.35	0.07	0.52

Notes: Table reports the simulated vaccination rate at baseline and in the indicated scenario. In each non-baseline scenario, we change the indicated parameters to match the mean among the "already vaccinated. In column (2) we change preferences for effectiveness and side effects, in column (3) we change beliefs about effectiveness and side effects, in column (4) we change preferences and beliefs, and in column (5) we instead change the remaining preference parameters (for the inside option, needles, and convenience).

Table 6: Model-predicted and actual vaccination rates, overall

	Vaccination rate
Model simulation, adjusting for contraindication rate	0.48
Analyis sample plus contraindicated sample	0.43
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 7: Stated preferences predict incentivized choices

Reward	$1.50	$3	$10	$20	$50
Pr(takeup)	0.36	0.48	0.59	0.54	0.46
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Reward beta	0.05	0.06	0.22	0.42	0.39
	(0.05)	(0.05)	(0.07)	(0.07)	(0.06)
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 8: Exccess pessimism about vaccine side effects

Group	Switch at $10	Switch at $20	Pool
Flu vaccine	0.11	0.06	0.10
	(0.04)	(0.06)	(0.04)
Verified vaccine	0.04	0.03	0.04
	(0.04)	(0.05)	(0.03)
Flu vaccine (wave 3, not pre-registered)	0.12	0.17	0.13
	(0.05)	(0.10)	(0.05)
Severe symptoms	-0.18	-0.25	-0.20
	(0.06)	(0.09)	(0.05)
TOT, Side effects	-1.70	-3.92	-1.75
	(0.81)	(3.98)	(0.67)
Sample expected effect	0.03	0.10	0.05
p-val (two sided)	0.03	0.31	0.01
N	319	146	465

Notes: Table reports the effect of an incentive to vaccinate on the indicated outcome (top four rows, robust standard errors in parentheses), the effect of vaccination on side effects among the vaccinated on the vaccinated ("TOT, Side effects" row), the sample mean believed side effect rate, and a p-value of the test that believed side effect rate equals the TOT. 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.

Figure 1: Experiment design

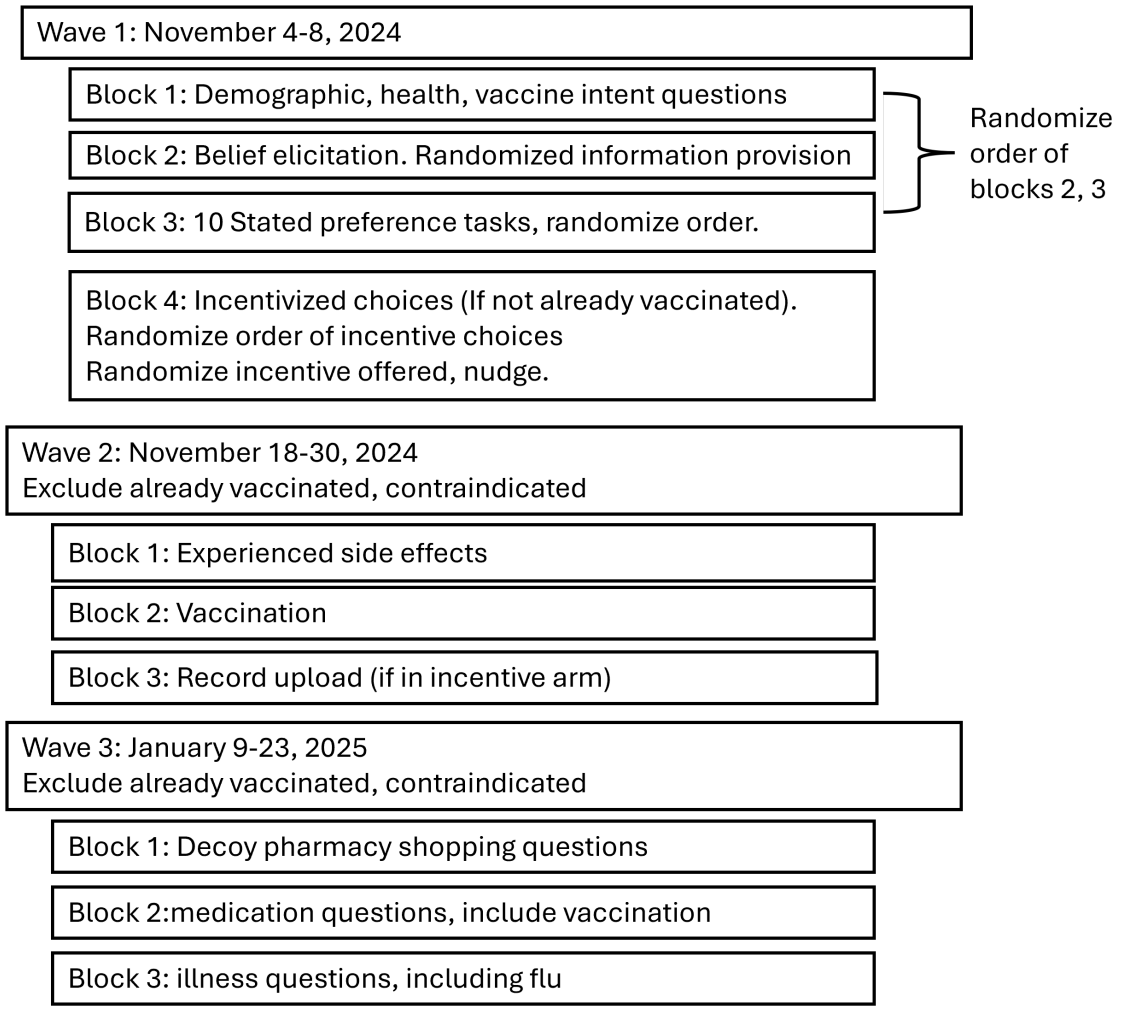

Figure 2: Beliefs about national averages

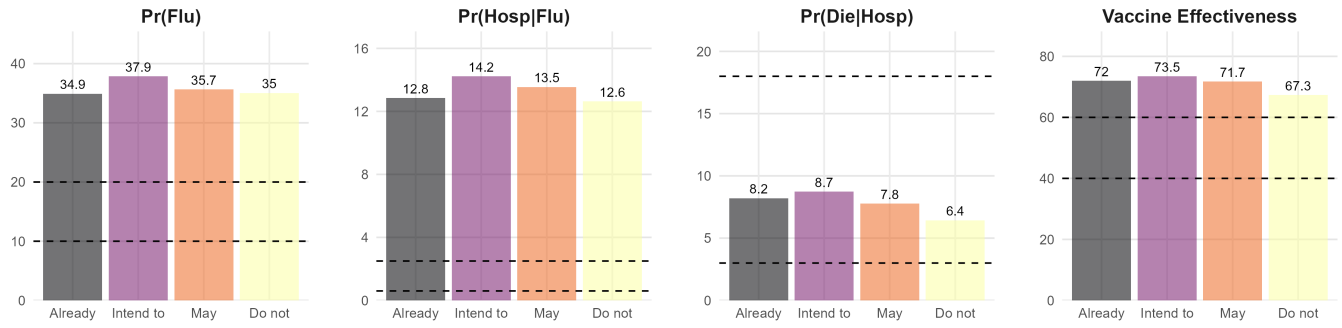

Notes: Figure reports the average belief about the national average value of the indicated object, for each vaccine intention group. Dashed lines are the values reported by the CDC.

Figure 3: Beliefs about personal risks and vaccine effects

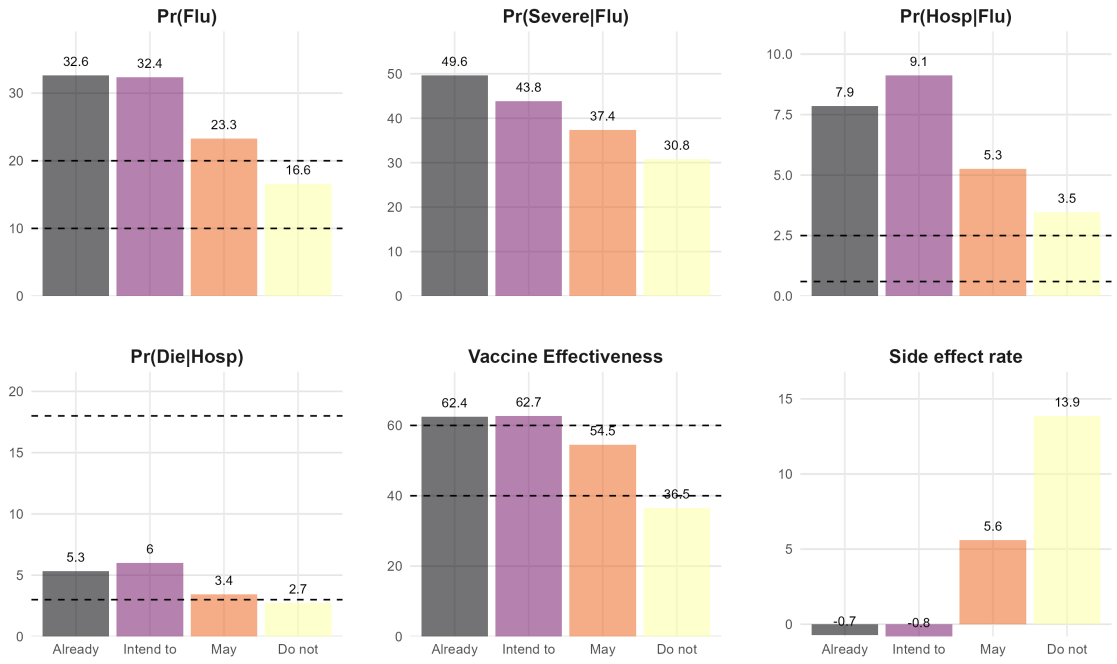

Notes: Figure reports the average belief about the person-specific values of the indicated object, for each vaccine intention group. Dashed lines are the values reported by the CDC (severe flu or side effects not reported).

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

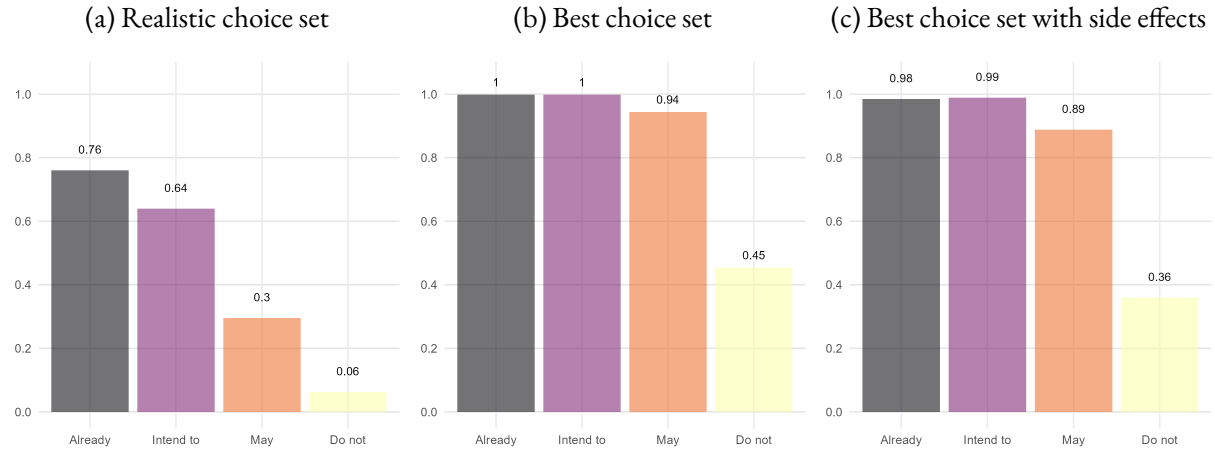

Notes: Each panel reports the fraction of people choosing to vaccinate given their choice set in our stated choice tasks, by vaccine intention group. The figures differ in the offered choice set. In Panel A we offer a realistic choice: no vaccine or a vaccine with no reward 50% effectiveness, 1% severe 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: Actual and model-predicted vaccination rates

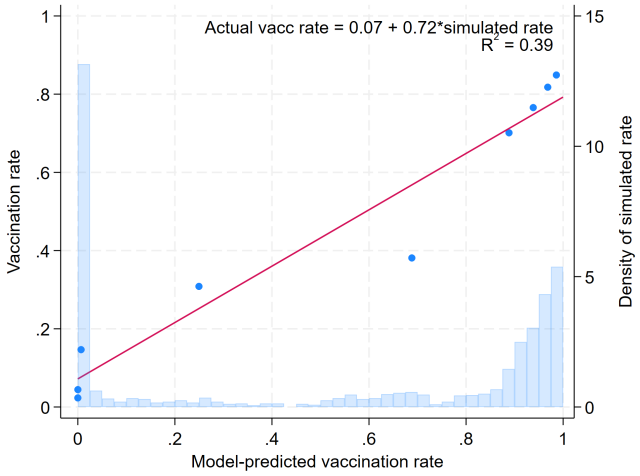

Notes: Figure shows the relationship between simulated and actual vaccination rates. The x-axis is the individual (posterior mean) model-simulated vaccination rate. We plot the distribution of this simulated rate in the background. We show a binned scatter of actual vaccination rate (measured as of wave 3) against simulated vaccination rate, along with the linear fit. We use 9 bins with equal numbers of observations in them.

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

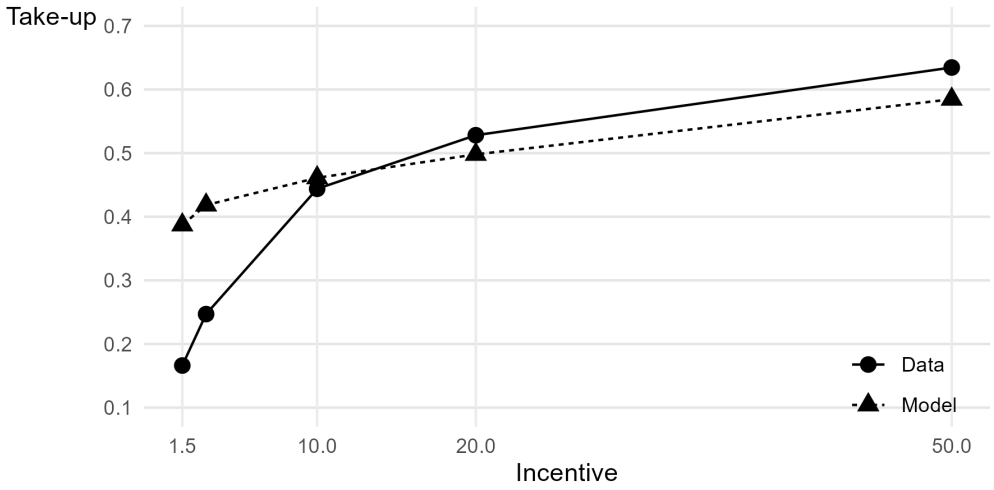

Notes: Figure reports the actual and simulated fraction of subjects who chose the indicated incentive level.

Figure 7: Predicted vaccination rates at alternative subsidies

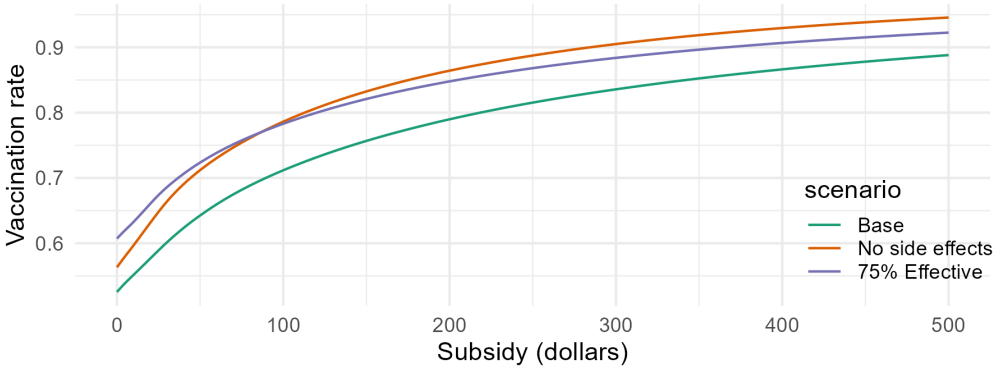

Notes: Figure reports the simulated vaccination rate as a function of vaccine subsidies. The base scenario corresponds to estimated preferences and group mean beliefs. The other scenarios set believed side effects to zero or believed effectiveness to 75 percent.

Figure 8: Vaccination rate and effective immunity vary with effectiveness and side effects

(a) Levels sets of vaccination rate

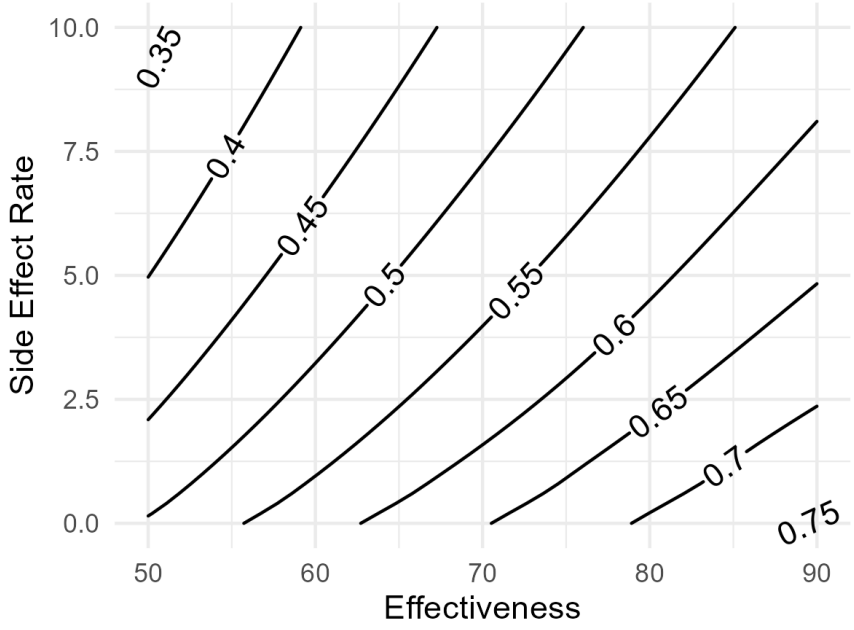

(b) Levels sets of effective immunity

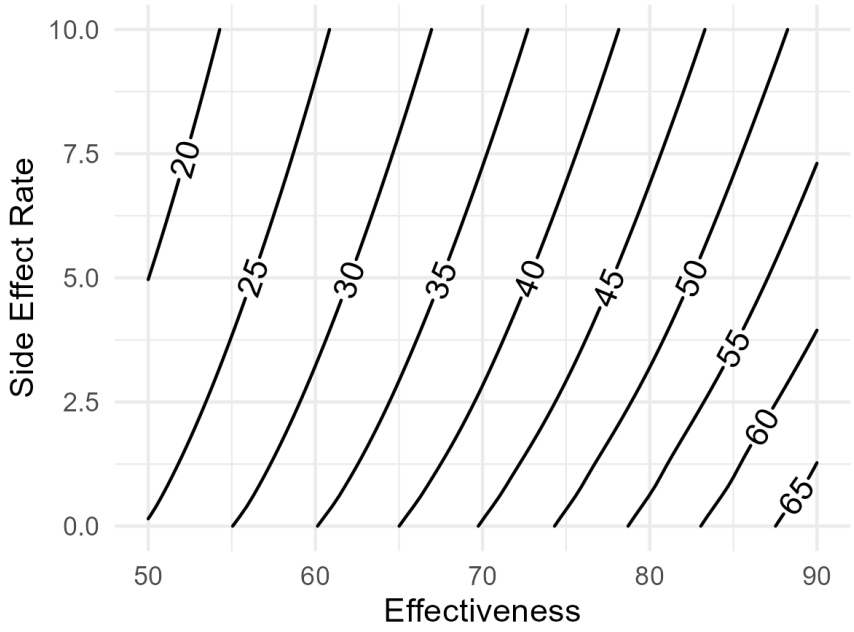

Notes: The figure shows level sets of vaccination rate (top panel) or effective immunity (bottom panel, defined as the product of vaccination rate and vaccine effectiveness). The level sets are derived from simulating vaccination rates under different effectiveness-side effect pairs, using the preference estimates from Section 5.

Appendices

For online publication only

A Appendix Exhibits

Table A.1: Sample sizes

Inclusion criteria	Count
Start survey	3,447
Pass attention checks	2,658
Consistent stated preferences	2,570
Only one attempt	2,569
Not missing key variables	2,567
No contraindications (analysis sample)	2,357
Not already vaccinated (eligible for incentives and waves 2/3)	1,757
Participated in wave 2	1,681
Participated in wave 3	1,449

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

Table A.2: 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.3: 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	0.22	0.24
	(0.05)	(0.08)			(0.05)	(0.08)	(0.05)	(0.08)
Severe flu given flu	0.15	0.29			0.12	0.23	0.04	0.14
	(0.04)	(0.05)			(0.04)	(0.05)	(0.03)	(0.05)
Hospitalization rate if flu	-0.06	-0.05			-0.17	-0.27	-0.20	-0.18
	(0.12)	(0.21)			(0.12)	(0.20)	(0.12)	(0.19)
Death rate if hospitalized	0.30	0.37			0.31	0.36	0.24	0.16
	(0.13)	(0.26)			(0.12)	(0.24)	(0.12)	(0.23)
Side effect rate			-0.21	-0.26	-0.21	-0.27	-0.14	-0.19
			(0.04)	(0.06)	(0.04)	(0.06)	(0.04)	(0.06)
Vaccine effectiveness			0.63	0.83	0.53	0.71	0.27	0.40
			(0.04)	(0.05)	(0.04)	(0.06)	(0.04)	(0.06)
Conjoint-implied vaccination probability							0.42	0.37
							(0.02)	(0.02)
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.4: 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.5: Robustness of preference-belief decomposition

Scenario	Simulated vaccination rate for...				
	Already Vacced	Intend to vacc	May or may not	Don't intend	All
	(1)	(2)	(3)	(4)	(5)
A. Baseline estimates					
Baseline	0.89	0.85	0.38	0.051	0.53
Equalize prefs for VE and SE	0.89	0.93	0.71	0.18	0.65
Equalize beliefs about VE and SE	0.89	0.85	0.64	0.19	0.62
Equalize prefs and beliefs about VE and SE	0.89	0.93	0.92	0.76	0.87
Equalize prefs for needle, conv., vacc intercept	0.89	0.81	0.35	0.07	0.52
B. Preferences vary by intention group					
Baseline	0.92	0.87	0.31	0.031	0.52
Equalize prefs for VE and SE	0.92	0.97	0.78	0.2	0.69
Equalize beliefs about VE and SE	0.92	0.87	0.63	0.17	0.62
Equalize prefs and beliefs about VE and SE	0.92	0.97	0.94	0.78	0.89
Equalize prefs for needle, conv., vacc intercept	0.92	0.79	0.26	0.059	0.49
C. Preferences vary with many covariates					
Baseline	0.89	0.85	0.38	0.05	0.53
Equalize prefs for VE and SE	0.89	0.93	0.69	0.17	0.64
Equalize beliefs about VE and SE	0.89	0.85	0.64	0.19	0.62
Equalize prefs and beliefs about VE and SE	0.89	0.93	0.91	0.72	0.85
Equalize prefs for needle, conv., vacc intercept	0.89	0.8	0.35	0.074	0.51

Notes: See notes to Table 5. Panel A. is identical. In panel B, we report simulations based on an alternative preference specification in which the mean (or log mean) of all preference parameters depends on vaccine intention group but not other covariates. In panel C, we instead let the mean (or log mean) of all preference depend on vaccine intention group, age (above/below 55), vaccine hesitancy score (above/below median), education (college / less), and beliefs (effectiveness, side effects, flu risk, hospital risk, death risk). In Panel D, we keep the preference specification the same, but stratify the simulated probability based on the (randomly assigned) order of completing the survey (belief elicitation first or vaccine choices first).

Table A.6: No survey order effects in preference-belief decomposition

	Simulated vaccination rate for...	
Scenario	Conjoint first	Beliefs first
	(1)	(2)
Baseline	0.53	0.52
Equalize prefs for VE and SE	0.64	0.66
Equalize beliefs about VE and SE	0.62	0.62
Equalize prefs and beliefs about VE and SE	0.86	0.87
Equalize prefs for needle, conv., vacc intercept	0.52	0.51

Notes: See notes to Table 5. This table is identical butwe stratify the simulated probability based on the (randomly assigned) order of completing the survey (belief elicitation first or vaccine choices first).

Table A.7: Side effects influence vaccination rate

Intention group	Vaccination rate, baseline	Vaccination rate, No side effects
Already vaccinated	0.89	0.89
Intend to vaccinate	0.85	0.85
May or not vaccinate	0.38	0.54
Do not intend to vaccinate	0.051	0.098
Overall	0.53	0.57

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.

Table A.8: 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.1: Vaccination rates are below CDC targets

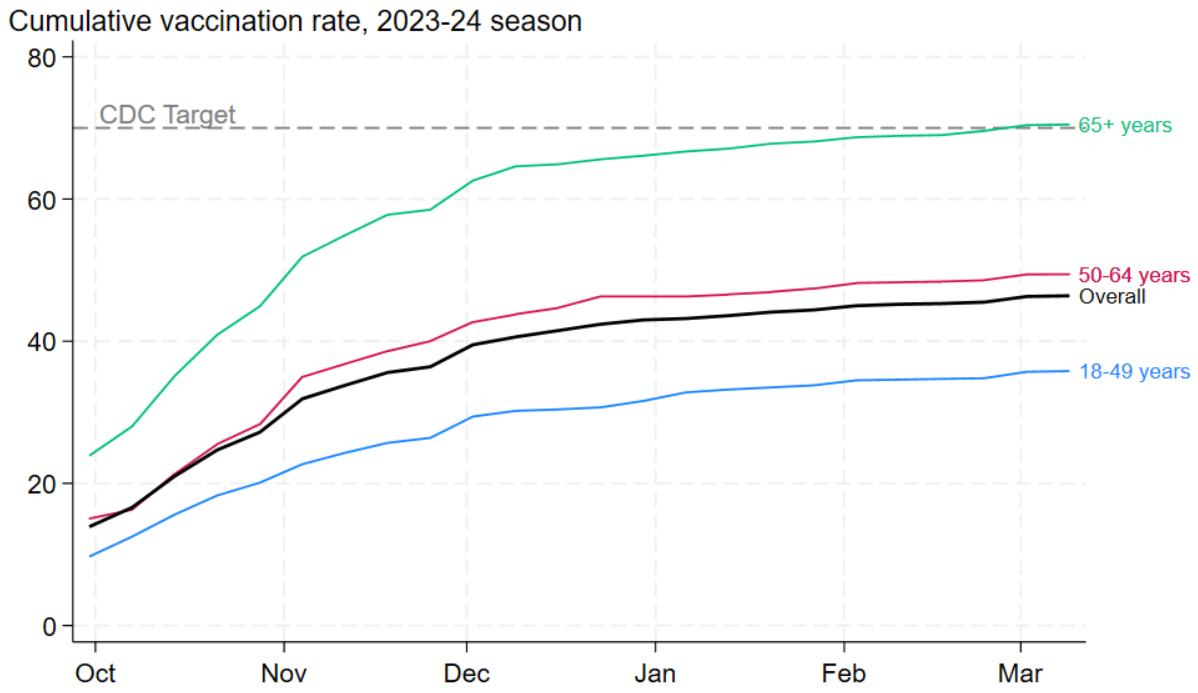

Notes: Figure plots the weekly cumulative vaccination rate for the 2023-24 flu season, overall and by age group. Source: Center for Disease Control (2025c).

Figure A.2: Actual flu experience correlates with expected flu risk

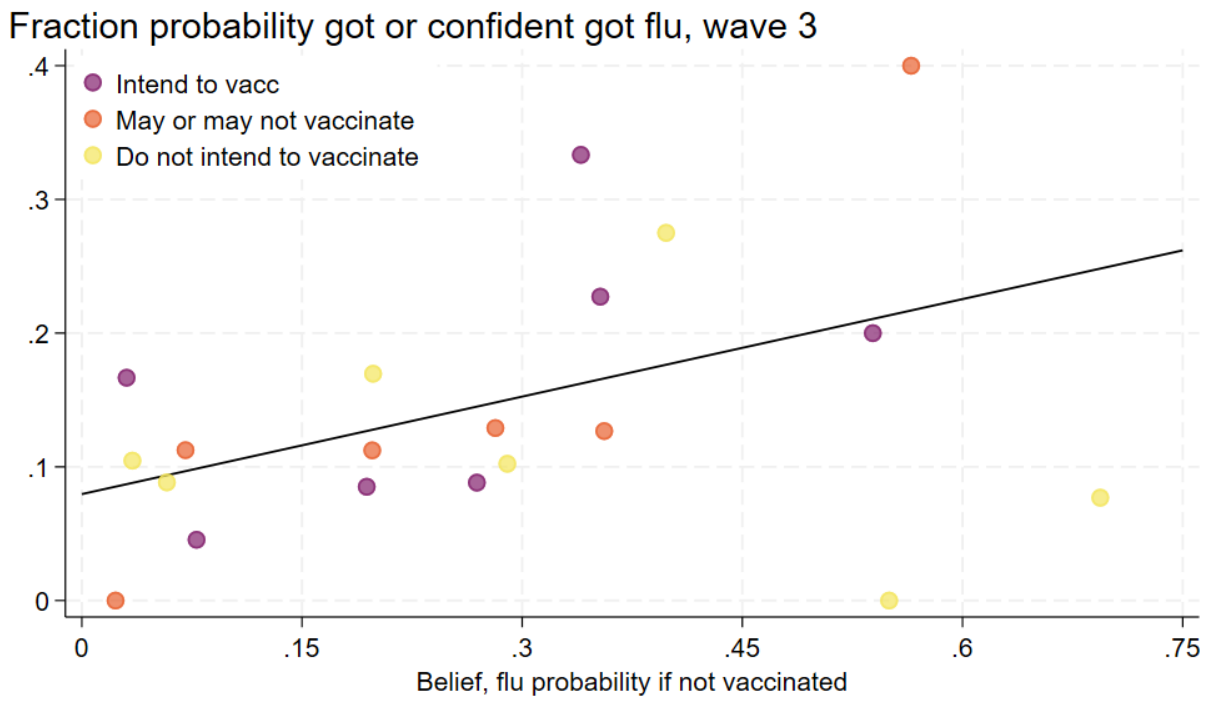

Notes: Figure plots actual flu rate against expected flu rate, among unvaccinated people completing the wave 3 follow up. We limit to unvaccinated because our expected flu rate question conditions on not vaccinating.

Figure A.3: Distribution of subjective beliefs

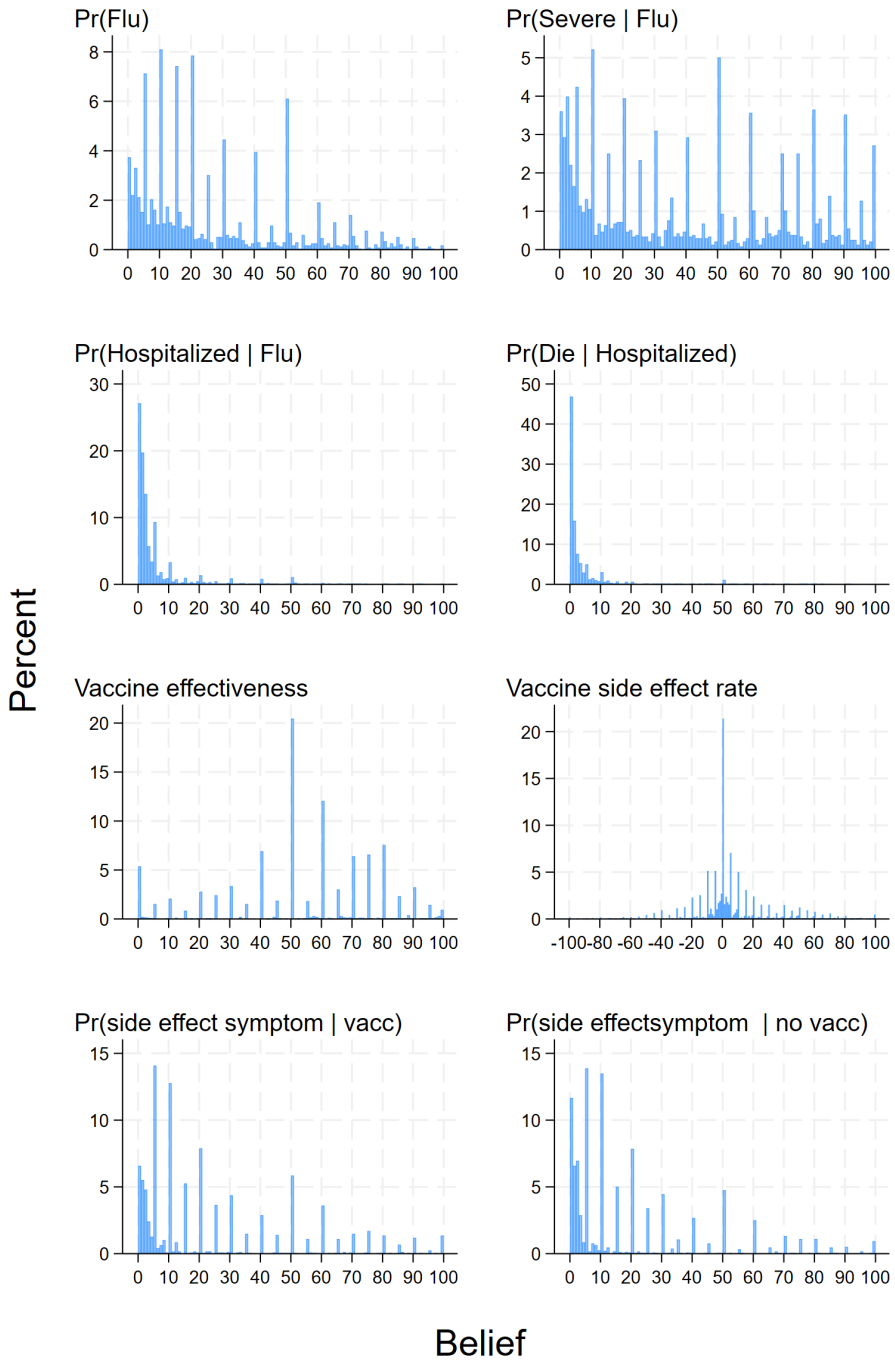

Notes: Figure plots the percent of respondents in each 1 percentage point bin of beliefs, for our subjective belief measures. The first four belief measures condition on not being vaccinated. The next two are beliefs about the flu vaccine (effectiveness against flu and effect on experienced side effects).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

Figure A.4: 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.

B Sensitivity analysis - contraindicated participants

A surprisingly large fraction of our survey participants report a contraindication for the flu vaccine: 8 percent in our sample, versus 1.3 allergic reactions per million flu vaccine doses (McNeil et al., 2016). It is likely that some of these respondents misunderstood our contraindication question, and others used the question to express a general concern about vaccines. We therefore re-analyze our data, using a broader sample that includes participants who reported a contraindication. As these participants were not offered incentives to vaccinate and were not eligible for follow-up surveys, we limit the re-analysis to demographics, intentions beliefs, and preferences.

This re-analysis largely confirms our main finding, that effectiveness and side effects are central to the vaccine decision, with beliefs and preferences both playing large roles. However it also shows some important differences between participants who report a contraindication and those who do not. These differences suggest that respondents interpret the contraindication questions in multiple ways, with some appearing to interpret it as asking whether vaccines would be harmful to them, and others interpreting it as asking whether the flu would be especially We describe these differences and also our main results for this sample.

We report summary statistics for our main sample and for participants who report contraindications in Appendix Table B.1.[40] We allow respondents to indicate either "unsure" or "yes" for contraindications, and we report separate means for each group. Participants reporting contraindications appear distinct from our main sample. They are younger, less educated, and lower income, especially the "unsure group." They have much lower intentions to vaccinate, as expected. However, consistent with some misunderstanding of the question, 25-30% indicate already or intending to vaccinate.

Contraindications are associated with both greater perceived flu risk and more pessimistic views about vaccines. We show this in Table B.2. The table reports the results from regression each belief measure on indicators for the unsure or yes for contraindication, as well as indicators for vaccine intention and for demographics (age level, female, education level). The "unsure" group perceives similar flu risk but lower effectiveness and higher side effects, conditional on intentions, relative to no contraindications. The "contraindicated" group perceives both greater flu risk as well as lower effectiveness and higher side effects.

Finally, to examine the role of preferences and beliefs in flu vaccine take-up for the broader population that includes respondents with reported contraindicaitons, we re-estimate our demand system including contraindicated participants. We allow the mean of all preference parameter distributions to include an indicator for any contraindication, but otherwise adopt an identical specification. Then we re-simulate the main counterfactuals of interest, showing how flu vaccine take-up changes when we change (sets of) preferences or beliefs. We report these simulations in Appendix Table B.3.

Our results are not substantially affected by including contraindicated participants in the preference estimation and simulations. Comparing columns (1) to columns (2) and (3) shows that changing preferences for effectiveness and side effects (column 2) or beliefs about them (column 1) to match the most optimistic group yields large increases in vaccination rates. Changing both beliefs and preferences (over side effects and effectiveness) raises vaccination rates to 87 percent, very similar to our baseline simulation results. Changing other preference parameters has essentially no effect.

[40] We impose the same exclusion criteria to participants with contraindications as to our main sample: pass both attention checks, one survey attempt only, consistent conjoint answers. Additionally one respondent with contraindications did not answer one conjoint question; we drop this respondent.

Table B.1: Summary statistics by contraindication

Characteristic	Analysis sample	"Unsure" contraindication	"Yes" contraindication
Age 18-39	0.51	0.50	0.51
Age 40-54	0.29	0.32	0.26
Age 55-64	0.09	0.12	0.12
Age 65+	0.10	0.05	0.10
Female	0.52	0.55	0.59
High school or less	0.11	0.17	0.06
Some college	0.27	0.39	0.28
College or more	0.61	0.43	0.66
Income < $25,000	0.12	0.26	0.15
Income $25,000 - $50,000	0.20	0.22	0.24
Income $50,000 - $100,00	0.36	0.36	0.34
Income > $100,00	0.30	0.15	0.26
Already vaccinated	0.25	0.07	0.14
Intend to vaccinate	0.24	0.18	0.15
May or may not vaccinate	0.21	0.27	0.20
Do not intend to vaccinate	0.30	0.47	0.51
Number of participants	2,357	137	80

Notes: Each cell is the mean of an indicator variable. The main sample excludes participants who report contraindications, or are unsure if they have contraindications.

Table B.2: Subjective beliefs by vaccine contraindication

Belief about	Flu risk	Severe flu	Hospitalization	Death	Effectiveness	Side effects
Unsure if contraindicated	-1.45	1.99	4.46	4.70	-7.67	5.91
	(1.86)	(2.91)	(1.58)	(1.44)	(2.15)	(2.67)
Contraindicated	8.86	9.30	14.37	13.37	-5.72	6.50
	(2.69)	(3.61)	(2.90)	(2.98)	(2.68)	(3.89)
Overall mean belief	25.73	40.17	6.89	4.91	52.02	5.76

Notes: Table reports the coefficients on indicators for "unsure of contraindications" and "yes contraindication" from a regression of the indicated subjective belief on those indicators, plus indicators for vaccine intentions, female gender, and age categories. Robust standard errors in parentheses.

Table B.3: Simulation results are robust to including contraindicated respondents

Scenario	Baseline	Change...			
	Preferences & Beliefs	β_{VE}, β_{SE}	Beliefs	β_{VE}, β_{SE}, Beliefs	β_0,
	(1)	(2)	(3)	(4)	
Baseline	0.9	0.85	0.37	0.049	
Equalize prefs for VE and SE	0.9	0.95	0.72	0.16	
Equalize beliefs about VE and SE	0.9	0.85	0.64	0.2	
Equalize prefs and beliefs about VE and SE	0.9	0.95	0.92	0.76	
Equalize prefs for needle, conv., vacc intercept	0.9	0.81	0.34	0.07	

Notes: Table reports the simulated vaccination rate at baseline and in the indicated scenario. The table is identical to Table 5 except the sample includes participants with contraindications (who are used in estimation and simulation), and we allow the mean of all preferences to shift with contraindications.

C 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]

D 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 D.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 D.2 -D.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 D.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 D.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 D.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 D.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 D.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.

Figure E.1: Markov chains are well-mixed

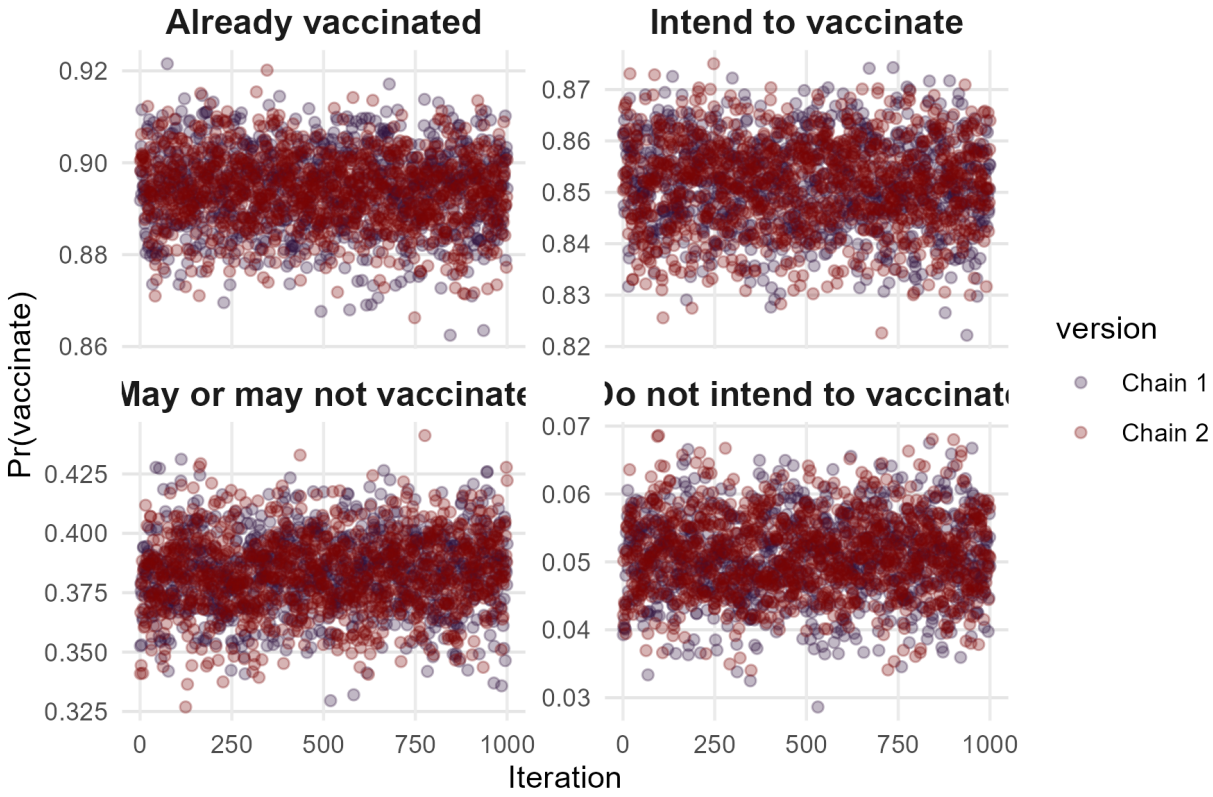

Notes: Figure plots the simulated vaccination rate, at each iterations of the Markov chain, for two separate starting values.

E 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 E.1 shows that our chains exhibit good mixing, and Appendix Figure E.2 shows that they are stationary. To construct Appendix Figure E.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 E.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 E.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 E.2: Markov chains converge rapidly

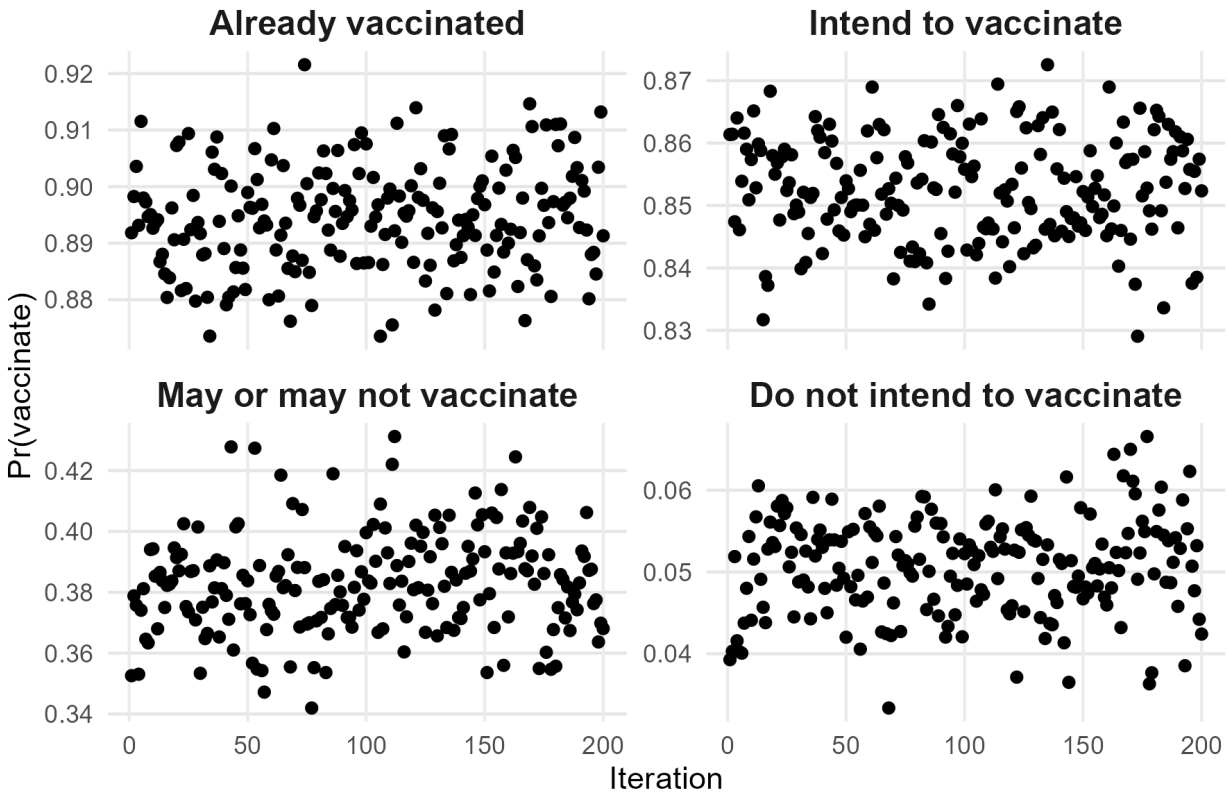

Notes: Figure plots the simulated vaccination rate, at each of the first 200 iterations of the Markov chain, after thinning to every 300th draws.

F Pre-registration modifications and adjustments

We pre-registered our experiment in the in the AEA RCT registry.[41] Here we describe the omissions and mistakes in our registration. By "omissions", we mean some details necessary for implementation which we failed to consider in our plan, as well as one deviation from the preanalysis plan.

The omissions are:

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. Instead of reporting WTA, we report WTP (i.e. the negative of WTA) because that is automatically produced by our model. 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 beliefs or objective characteristics in calculating WTA. We use subjective beliefs, consistent with our calculations throughout the paper.

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 at 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., chose not to vaccinate at $1.50 or $3, chose to vaccinate at $10, chose not to vaccinate 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.

6. Third wave. 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.

We also note two mistakes:

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

1. Preference heterogeneity. We specified both that we would follow Moshary et al. (2023) and that we would allow the covariance matrix of preferences to vary with participant characteristics, which Moshary et al. do not do. Upon realizing our mistake, we decided not to let the preference covariance matrix vary with participant characteristics.

2. We specified that one of demographics would be age ≥ 50. However we actually asked about age in 4 ranges: 18-39, 40-54, 55-64, and 65 and older. We therefore stratify on age ≥ 50.

We deviated from the plan in two ways.

1. We pre-specified that we could compare our sample to the 2024 BRFSS. As of this writing (May, 2025) the 2024 BRFSS are not publicly available so we compare to the 2023 BRFSS.

2. We intended to simulate vaccination rates using individual-level beliefs. However, doing so makes it impossible to cleanly decompose the role of preferences vs. beliefs in explaining differences in vaccination rates across groups.[42] We therefore simulate vaccination rates using group-level mean beliefs.

[42] The problem stems from the nonlinearity of vaccination rates with respect to beliefs, and differences in the belief distribution across groups. Roughly, the decomposition requires that we swap the distribution of beliefs between intention groups, but hold fixed the conditional distribution of preferences given intentions and beliefs. However, the preference distributions do not have common support, so the conditional preference distribution is not identified.