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

Daniel W. Sacks[†] Justin Sydnor[‡]

September 11, 2026

Abstract

Many health-related interventions have modest impacts and are based on presumptions about individual biases and mistakes. Whether intervention is warranted, and in what form, depends on separating biased beliefs and preferences from normatively relevant heterogeneity. We do so for flu vaccination, combining online discrete-choice experiments, belief elicitations, and incentivized validation choices. While the vaccine hesitant are generally well calibrated about average flu risk and vaccine effectiveness in the general population, they believe that their own risk is low and vaccines will not be effective for them. They also overstate side effect rates, a belief inconsistent with their own experience when we randomly incentivize vaccination. Model-based estimates show that these side effect beliefs, along with a low valuation of vaccine effectiveness, are quantitatively key contributors to vaccine hesitancy. Simulations based on our discrete-choice utility model suggest that eliminating potential biases toward side effects and vaccine effectiveness could raise the fraction intending to vaccinate from 55 to 72 percent. However, the remaining 28 percent, two-thirds of the hesitant, may be making a personally optimal decision to not vaccinate. We estimate optimal internality-correcting subsidies and find that they yield only small welfare gains because subsidies target bias poorly.

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

Governments, insurers, health systems, and academic researchers devote significant resources to changing health behaviors, including promoting screening, medication adherence, and vaccinations. Many interventions stem from an assumption that people are making mistakes and target information, encouragement, and incentives to address perceived "internalities". The impacts of interventions, though, are often quite modest. These findings tend to encourage continued testing of alternative interventions aimed at promoting the seemingly desirable health behavior. These teste, however, do not address whether the intervention is warranted in the first place, nor what form and magnitude should it take. As the growing literature in behavioral public economics highlights (see Bernheim and Taubinsky (2018)), distinguishing biases from normatively relevant preferences and determining optimal internality-correcting subsidies requires richer data than most studies of interventions for health behavior produce.

This paper develops and implements an empirical approach to that problem and applies it to intentions to vaccinate against the seasonal flu.[1] We use a combination of discrete-choice experiments, discrete-choice utility modeling, and belief elicitations to identify the extent to which people have biases in their beliefs and preferences toward vaccination as well as potentially welfare-relevant heterogeneity in preferences (e.g., toward hassle costs or disutility of vaccinating) that make it privately optimal for some people not to vaccinate even in the absence of biases. We use these insights to inform the potential effectiveness and directions for interventions targeting beliefs and for the design of optimal corrective subsidies.

We conducted a large-scale multi-wave survey experiment in an online platform. We elicited participants beliefs about flu risk and severity, the effectiveness of vaccination, and the prevalence of side effects from vaccination. To estimate preferences toward vaccination and specific features of the vaccine experience, we asked participants to choose between hypothetical flu vaccine options with different effectiveness, side effect rates, delivery mechanism (needle or nasal spray), and convenience, as well as different incentives in the form of gift card rewards. This type of hypothetical discrete-choice design is 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 discrete-choice experiment is that it lets us vary characteristics which are otherwise fixed in observational settings. We estimate preferences for vaccine characteristics using standard Bayesian MCMC methods, allowing for rich heterogeneity including differences across vaccine hesitant and confident people as well as within-group unobserved preference heterogeneity. The estimated model allows us to simulate demand for vaccination under various

[1] Raising seasonal vaccination rates has been a longstanding goal of health policy and health researchers. Seasonal illnesses such as influenza, COVID-19, and RSV annually kills tens of thousands of Americans and hospitalizes hundreds of thousands (Naquin, 2024; Ahmad, 2024). Tens of millions more Americans suffer pain and disruption from flu, resulting in lost work, education, and leisure (Ward, 2014; White, 2021; Daysal et al., 2021; Cheipesh, 2025). Despite these health costs, many people do not vaccinate seasonally; annual adult flu vaccination rates over the past decade have averaged 37-50 percent (Kahn et al., 2026).

counterfactual vaccine characteristics, belief patterns, and incentive regimes.

Several factors establish the validity of our belief and preference data. Beliefs are highly predictive of vaccination choices, correlated with risk factors, predictive of health outcomes, and distinct from preferences. Stated preferences are predictive of actual vaccination choices, including in incentivized choice experiments in which we give participants choices between sure rewards and rewards contingent on vaccination. Participants with stronger estimated preferences to vaccinate are more likely to choose contingent rewards, and price sensitivity estimated in the stated preference experiment is predictive of sensitivity to contingent rewards. Despite these strong validity results, our predicted vaccination rates are higher than actual vaccination rates. This gap likely reflects the known action-intention gap in vaccination (Milkman et al., 2011; Bronchetti et al., 2015), as stated preferences likely reflect intentions and not necessarily actions. Accordingly, we interpret our simulations as pointing to interventions which can shift intentions. These simulations are an important complement to existing research, which has mainly documented interventions to increase follow-through conditional on intentions.

We show that vaccine effectiveness and side effects play a key role in vaccination decisions, even among the vaccine hesitant. Model-free evidence shows that among vaccine hesitant people, only 18 percent would choose a vaccine at prevailing effectiveness and short-run side effect rates. But when offered a choice set with highly effective vaccines and side effects similar to the prevailing vaccine, or with no side effects, vaccination rates rise to 50 percent or higher, showing substantial sensitivity to side effects and effectiveness. This finding, new to the literature, provides clear causal evidence that vaccine take-up is sensitive to short-run side-effects and effectiveness.

Our results on beliefs help explain why information campaigns may struggle to increase vaccine intentions, and point towards potentially welfare-improving alternative interventions. We find that both those who do and do not intend to vaccinate have reasonably well calibrated beliefs about the prevalence and severity of the flu in the general population and about the overall average effectiveness of vaccination. If anything, beliefs about these average are slightly biased in favor of vaccinating, suggesting information campaigns about average flu risk and vaccine effectiveness are unlikely to increase vaccination. However, consistent with prior research on vaccine hesitancy (for example, Chapman and Coups (1999); Blue and Valley (2002); Shahrabani et al. (2009); Shahrabani and Benzion (2010); Tsutsui et al. (2012)), we find substantial variation in subjective beliefs about ones' own flu risk and the personal effectiveness of vaccination. Those who are hesitant to vaccinate believe they have much lower risk from the flu and that vaccines will be significantly less effective for them. It is difficult to determine whether this belief heterogeneity reflects bias or could reflect real underlying heterogeneity in flu risk and vaccine efficacy. Where we do detect what appears to be a clear bias, however, is in subjective beliefs about side effect rates. Vaccine hesitant participants believe that vaccines have 8-15 percent severe side effect rates, much higher than clinical trials find. We show, in a novel result, that this belief is also inconsistent with participants' own

experiences: we provide randomized incentives to vaccinate and show that their effect on experienced side effects is much lower than participants expect. To the extent that there is scope to correct misperceptions, these results suggest those efforts should likely focus on side effects. Simulations suggest that if it were possible to correct side-effect misperceptions, it could raise overall vaccine intentions by about 4 percentage points, on a base of 55 percent, driven entirely by an 11 percentage point increase among vaccine hesitant participants.

Our discrete-choice model reveals substantial heterogeneity in preferences over vaccine characteristics, suggesting both role for further bias but also a significant share of the population for whom vaccination may not be privately optimal. Vaccine hesitant participants value vaccine effectiveness much less than vaccine confident participants; they are much less willing to pay to increase effectiveness and much less willing to tradeoff effectiveness against side-effects. At age-specific average flu mortality rates, the most vaccine hesitant participants have an implied value of statistical life of only $4 million, significantly below the $12-$14 million for those willing to vaccinate, which is in line with standard VSL benchmarks. Their low value on effectiveness likely stems from their low perceived flu risk: the most vaccine hesitant participants perceive a flu mortality risk that that is about a tenth of objective (age-specific) average risk. If one believes that those beliefs are biased, then the low valuation of vaccine effectiveness may also be a bias. However, if one takes the subjective beliefs as welfare relevant, then the implied subjective VSL for the vaccine hesitant is in line with conventional benchmarks and there may be no bias toward effectiveness. Using our discrete-choice model we simulate that if each group's median VSL (evaluated at age-specific average mortality) changed to $12 million (by changing health preferences directly, or changing mortality-risk beliefs), vaccine intentions would increase by 6 percentage points. Changing both preferences for effectiveness and beliefs about side effects would raise vaccination intentions by 17 percentage points, to 72 percent. This suggests that up to around a third of those who do not intend to vaccinate at baseline could be biased and making a mistake. Yet if we respect the remaining preference heterogeneity as potentially normatively relevant, we estimate that at least 28% of our subjects (or about two-thirds of the baseline vaccine hesitant) are making a privately optimal decision not to vaccinate.

We use these estimates to examine the use of subsidies to correct biases in flu vaccination choices. Allcott (2015) shows that the optimal internality-correcting subsidy is equal to the average bias among people who are marginal to the subsidy, i.e. just indifferent between purchasing or not at the prevailing subsidy, and therefore determining the optimal requires obtaining the joint distribution of biases and demand. We extend the approach in Allcott and Taubinsky (2015) and show how to recover this distribution from our preference estimates under assumed sources of bias (beliefs about side effects, the implied VSL, or both). We estimate that optimal subsidies would not achieve first-best levels of vaccination, sometimes undershooting and sometimes overshooting, and generally produce small average welfare gains. The reason is that subsidies are not well targeted, because many biased participants nonetheless have negative

normative willingness to pay, so subsidies end up inducing inefficient as well as efficient vaccinations. These results suggest that while subsidies may increase vaccination rates, for the flu vaccine the welfare justification for such an approach would likely need to rest on estimates of the externality benefits of vaccination rather than correction of internalities.

Our work contributes to several literatures. First is the large literature on vaccine hesitancy, much of which has been rooted in the health belief model in which intentions to vaccinate depend on beliefs about the likelihood of contracting illness, the severity of illness if contracted, the effectiveness of vaccination, and the costs of vaccination, especially side effects (e.g. Janz and Becker (1984)). A substantial body of evidence shows that attitudes towards vaccines and intentions to vaccinate strongly correlate with these beliefs, with vaccine hesitant people holding benign beliefs about disease risk and more pessimistic beliefs about vaccine safety and effectiveness (for example, Chapman and Coups (1999); Blue and Valley (2002); Shahrabani et al. (2009); Shahrabani and Benzion (2010); Tsutsui et al. (2012)). Our discrete-choice experiments provide experimental variation suggesting that these previously documented correlations reflect underlying causal impacts of beliefs about vaccine effectiveness and side effects. We further contribute by providing the first evidence that short-run side effect beliefs are biased in the sense of overstating experienced side effects.

Within the literature on vaccine hesitancy, our results also speak to the limited effectiveness of many information and "nudge-style" interventions at encouraging vaccination. For example, Patel et al. (2023) report the results of a mega-study that tested 19 different variations of text-based messages encouraging flu vaccination at an upcoming primary care visit. The average effect across interventions was a very modest 1.8 percentage-point increase in vaccination rates and most of the interventions did not have a statistically significant impact.[2] Much of the content of the messages in this mega study focused on the average severity of the flu and benefits of vaccination, but our work points to the importance of individual subjective assessments of risk and effectiveness rather than biased beliefs about average effects.[3] Moreover, none of the interventions addressed beliefs about side effects, which our analysis suggests are a significant driver of vaccine hesitancy. The results here highlight that more directly studying the nature and variance in beliefs and attitudes toward vaccination may point toward different directions for interventions.

Our results similarly speak to research within the broader vaccine-hesitancy literature that has focused on the impacts of financial incentives to vaccinate. A number of studies have documented positive impacts of moderate financial incentives on vaccination rates for both the influenza and COVID-19 vaccines (e.g., Bronchetti et al. (2015); Alsan et al. (2019); Chang et al. (2026); Campos-Mercade et al. (2021); Thiru-

[2] Large-scale public health messaging and advertising campaigns also had limited to no impact on COVID-19 vaccination (Athey et al., 2023; Larsen et al., 2023; Ho et al., 2023)

[3] Interestingly while attempting to correct some myths about vaccines have not increased vaccine intentions (Nyhan et al., 2014; Nyhan and Reifler, 2015), focusing on highlighting the severity of illness among children for measles vaccines has increased vaccinations (Horne et al., 2015).

murthy et al. (2022)). The methods we use in this study highlight the feasibility and value of determining the optimal level of internality-correcting subsidies and considering not only the impact on vaccination rates but the underlying welfare impact of incentives.[4] In our case we estimate that there is enough potentially normatively welfare-relevant vaccine hesitancy and subsidies target bias poorly enough that the internality-based welfare benefits of raising influenza vaccination rates through subsidies is likely small.

The second major literature we contribute to is the growing literature on behavioral public finance, which considers taxes, subsidies, and other instruments to correct behavioral distortions such as unawareness of program eligibility, biased beliefs, and self-control problems (e.g., (Allcott and Taubinsky, 2015; Baicker et al., 2015; Allcott et al., 2019, 2022; Chandra et al., 2024). Our results connect in particular to Allcott et al. (2025), who show that the welfare effects of nudges depend not only on the average distortion they correct but on how they change the variance of choice distortions across consumers. We find a counterpart of this insight on the price side: because the optimal subsidy is set to the average bias among marginal consumers, it corrects the mean of the distortion but not its dispersion. In our estimates, that dispersion is large relative to the average marginal bias and only weakly related to normative willingness to pay. Taken together, these results suggest that how well an instrument is targeted on the distribution of bias governs how much of the available welfare improvements corrective policy recovers.

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 potential 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.

[4] Goodkin-Gold et al. (2022, 2024) develop formulas for optimal subsidies to correct externalities, but not internalities.

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). There is considerable uncertainty in these estimates: the 95 percent confidence interval for 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.[5]

Flu risks are heterogeneous. 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).[6] 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. The flu vaccine, like others, often produces short run discomfort at the vaccination site. It has a low rate of severe (i.e. grade 3) short-run

[5]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. The uncertainty in these estimates derives from the fact that the CDC estimates population prevalence from data on a network of reporting hospitals (Center for Disease Control, 2025b; Rolfes et al., 2018; Reed et al., 2015).

[6]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.

side effects, however; clinical trials reported on package inserts indicate the side effect rate is less than a percentage point higher in the vaccine than in placebos (e.g. U.S. Food and Drug Administration (2026)).

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 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).

3 Experiment Design

3.1 Overview

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 four blocks making up the the core of our experiment. We begin by asking a set of questions about demographics, health, trust in doctors and the government, and vaccine intentions. The trust in doctors and government questions are on 5-point scales and are used as part of assessments of vaccine hesitancy (as in Larson et al. (2015)). The health questions includes questions about flu risk factors (e.g. asthma) and vaccine contraindications. Next are blocks containing the discrete choice experiment and eliciting beliefs and flu risks and flu vaccine consequences. The final block offers incentivized choices around the flu vaccine. We randomized the order of belief elicitation and discrete choice experiments, and for a random subset of participants we provided information about average flu risk and vaccine effectiveness from the CDC. We included attention checks in the belief elicitation and discrete choice experiment modules.

Waves 2 and 3 are brief follow-ups to measure vaccine take-up, experienced side effects (for those vaccinating and not), 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.

3.2 Discrete choice experiment details

We use a discrete choice experiment to measure preferences over vaccine characteristics. Our experimental set up is common in the marketing literature (e.g., Allenby et al. (2019)) and has seen many recent eco-

nomics applications (e.g. Ben-Akiva et al. (2019); Kessler et al. (2019); Moshary et al. (2023); Chan (2022)). In our design, participants make hypothetical choices between two vaccines and a no vaccinate option in 10 choice opportunities. The vaccines differ in their characteristics. By carefully varying the characteristics across the 10 choice sets, we 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 and in particular intentions. We describe several validity tests in Section 7 below.

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.[7]

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.[8] Appendix Table A.4 reports the options available in each choice set, and Ap-

[7] 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.

[8] 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.

pendix Figure A.3 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 explanations of the characteristic in each choice scenario. Finally, before the conjoint, we include a baseline choice that resembles beliefs about the prevailing vaccine. The baseline choice asks participants to choose between a vaccine with no reward, 50 percent effectiveness, 5 percent side effect rate, needle delivery, and not convenient, or no vaccination at all. (A 5 percent side effect rate is higher than actual side effect rates but is close to the average side effect belief we find.)

3.3 Belief elicitation

Conceptualization: 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, encompassing both the likelihood of catching the flu and the severity of illness conditional on disease. We condition on not vaccinating because this risk, along with the consequences of vaccination, determines the incentive to vaccinate.

For consequences of vaccination, we focus on effectiveness and short-term side effects. Effectiveness is the percent reduction in flu risk from vaccination, relative to not being vaccinated. Short-term side effects are the effect of vaccination on the likelihood of experiencing grade 3 adverse events, such as fatigue, muscle soreness, or fever. Grade 3 means disruptive to daily routines and therefore represent a meaningful cost of vaccinations. Clinical trials establishing the safety of vaccines elicit and measure mild and severe–i.e., grade 3—side effects; for people getting the flu vaccine, the grade 3 side effect rate is within a percentage of the placebo rate (U.S. Food and Drug Administration, 2026). We do not attempt to measure beliefs about long-run consequences of vaccination. We focus on effectiveness and short run side effects because we can compare these beliefs to objective benchmarks, and substantial individual level uncertainty exists about these objects (as people experience post-vaccination discomfort, for example, but cannot definitively attribute it to the vaccination or to coincidence).

Measurement: The exact text of all our elicitations is in Appendix B; here we summarize. To elicit beliefs about flu risks, we begin by saying "Some people choose to get the flu vaccine, and others choose not to. For the next few questions, we'd like to ask your beliefs about what will happen to you **if you choose not to get the flu vaccine.**" (note the face-saving language, ala Stantcheva, 2023). Then we ask about likelihood of getting the flu in the next year, as well as the likelihood of a severe case given flu, hospitalization given flu, and 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. We first give respondents a 7-point likert scale (ranging from "would definitely not happen" to "would definitely happen"), and then

we ask for quantitative responses, with a scale set by the multiple choice answers.[9] Asking twice not only lets us set the scale but also lets us correct measurement error in some specifications (as in Giglio et al. (2021)). s be above 10 percent (flu risk).

We measure beliefs about vaccine effectiveness similar to how we measure beliefs about flu risk, first asking on a multiple-choice scale and then on a quantitative scale. When measuring believed effectiveness, we ask "If you do get the flu vaccine, how much would you say that reduces your chances of getting the flu?"[10] However, we must be especially careful in measuring side effect beliefs. 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 .[11] We therefore measure vaccine side effects using a different, two-step approach. First we ask participants their likelihood of experiencing over the next few days a severe case of headaches, muscle aches, upset stomach, tiredness, joint pain, fever, or general discomfort, assuming they do not get vaccinated. Call this $rate_i(0)$. We measure the rate with vaccination, $rate_i(1)$, analogously, and we obtain the believed side effect rate as $rate_i(1) - rate_i(0)$. Note that our elicitation procedure does not impose that $rate_i(1) \geq rate_i(0)$, and indeed because there is overlap in side effect symptoms and flu symptoms, a person who believes flu risk is high and the flu vaccine is effective might report that $rate_i(1) < rate_i(0)$.

Beliefs about objective statistics: Before we elicit beliefs about personal risk and vaccination consequences ("subjective beliefs"), we first measure beliefs about aggregate risks and CDC-reported vaccine effectiveness ("objective beliefs"). Objective beliefs are helpful for understanding whether subjective belief differences reflect disagreement about facts or disagreement about personal risk. Our objective belief measures are analogous to our subjective measures, except they do not condition on vaccination status and they are explicitly framed as asking beliefs about what the CDC reported.

Information provision: Between the objective and subjective belief elicitation, we provided information about the CDC's range of estimates to a random 50 percent of participants. We repeated this information next to each subjective belief elicitation. We hoped to use this information provision as an instrument for personal beliefs. However it appears that participants over-reacted to this information, because in the information provision group, beliefs are clustered at exactly the provided values. (See Appendix Figure A.2.) We therefore do not use beliefs from the information provision arm.

[9] In the quantitative elicitation, we use sliders (despite the possibility of trembling hand measurement error (Stantcheva, 2023)) to communicate an appropriate scale, based on the multiple choice answer. If an outcome would not happen or is unlikely, the slider ranges from 0 to 25; otherwise it ranges from 0 to 100. Narrowing the range is helpful because some outcomes could be less than 1 percent (e.g. hospitalization risk) but others could be above 10 percent (flu risk).

[10] 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).

[11] Smith et al. (2021) notes the importance of clear causal contrasts for measuring beliefs about treatment effects.

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.[12]

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). Our focus instead is on measuring beliefs over manipulatable vaccine characteristics. Second, some participants may lack the numeracy to give meaningful answers to this question. Third and related concern is that our belief measures are not incentivized, and subjects therefore may not provide adequate effort in answering.

In practice these concerns likely have limited impact. 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.2, which shows that respondents use the whole response scale, beliefs cluster towards zero where such clustering is appropriate, and relatively few respondents choose 50% for beliefs, except for vaccine effectiveness, where such a belief is warranted.[13] We also show in Appendix Table A.7 that beliefs strongly predict vaccination rates, in the way we would expect, showing there is signal in the belief measures. This table also shows that beliefs predict vaccination rates conditional on our preference estimates,[14] indicating that beliefs and preferences are distinct objects, each with their own explanatory power. Finally, we show that although beliefs are on average biased for realized risk, they are nonetheless predictive of it: people who believe they face higher likelihood of experiencing side effect symptoms or getting the flu are indeed more likely to report experiencing those symptoms or getting the flu in follow-up surveys (Appendix Table A.6).

[12] 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.

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

[14] 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 and independently predictive of actual vaccination rates.

3.4 Incentivized choices

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[15] 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.[16]

3.5 Follow-up surveys

Wave 1 of our survey 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 to all eligible participants, encouraging completion of wave 2. In wave 2 We ask about experienced side effects over the prior two weeks. Our measure closely parallels our side effect belief elicitation, so we can compare expectations to experiences. (See Appendix B for exact wording.) Next we ask participants if they had the flu over the prior flu weeks and (separately) if they've gotten the flu vaccine this season. Finally, for participants offered a reward to vaccinate, we ask for a record upload. To earn their reward, participants must upload a pharmacy receipt or a screenshot of their vaccine record. Note that we ask about side effects *before* vaccination, to avoid any priming.

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. We measure vaccination a second time because we found that many people in wave 2 reported vaccinating but did not provide a validated record (despite a strong incentive to do so), so we worried their answers reflected experimenter demand. In Wave 3, to avoid experimenter demand effects, we obfuscate our intentions (as in, for example, Haaland

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

[16] We found that this nudge accelerated vaccination but did not increase the fraction ever vaccinating. We do not analyze it further.

and Roth (2020)) by opening with questions about pharmacy shopping and include questions about non-vaccine medications, and non-flu illness. We found that few participants switched to reporting not vaccinating, and the obfuscation appeared successful. When we asked participants what they thought the survey was about, none reported a connection to our earlier work. We conclude that the failure to upload records reflects difficulty with or mistrust of the technology, rather than dishonest reporting of vaccination status.[17]

3.6 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."[18] 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 in Section 7 that beliefs and preferences are individually predictive of actual vaccinations and beliefs about personal flu risk are predictive of actual flu incidence, meaning that our belief and preference measures capture distinct, behaviorally-relevant information which corresponds to choices and experiences.

[17] Consistent with this interpretation, we observed, out of over a thousand opportunities for dishonestly, only two apparently deceitful record uploads - one an AI-generated nonsense image, the other an image that was not evidence of vaccination.

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

3.7 Sample formation

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.[19]

Appendix Table A.1 shows how the sample size changes with successive inclusion criteria. 3,447 participants started wave 1. Our main analysis sample excludes respondents with any failed attention check or inconsistent stated preference,[20] 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 including participants reporting contraindications where possible. Our qualitative and quantitative conclusions are quite similar; see Appendix D.

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 2024 BRFSS, a nationally representative survey of adult health behaviors. Like most online panels, our sample skews towards younger, higher educated, and higher income participants. However, the vaccination rate in our sample is very similar to the national average, after accounting for exclusion of participants) who report a contraindication to the flu vaccine. In our sample (excluding those participants), the vaccination rate is 43 percent. Including those participants, the vaccination rate falls to 40 percent, which is the vaccination rate in the BRFSS. Thus our sample is roughly representative on this key metric.[21] Other work measuring vaccine hesitancy

[19] 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.

[20] 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.

[21] We measure vaccinations as follows. First, we say a participant has vaccinated if they report already vaccinating in the baseline survey. Next we record their vaccination status usin the wave 3 data if they participated in wave 3, and the vaccination status using the wave 2 data if they participated in wave 2 and not wave 3. Finally we impute vaccination as 0 for contraindicated participants who do not report already vaccinating (the vast majority), as contraindicated participants were not invited to hte

has also used online panelists (e.g. Akel et al. (2021)), and we find similar hesitancy rates.[22]

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. About a quarter of participants are in each category. The BRFSS does not ask a comparable question.[23]

3.8 Pre-registration

This experiment was pre-registered in the AEA RCT registry.[24] 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 E 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 (in the non-information-provision arm) 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.[25]

follow-up surveys. This procedure sets vaccination status to missing for participants who participated in neither follow-up survey and were not already vaccinated.

[22]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.

[23]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.

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

[25]The problem stems from the nonlinearity of vaccination rates with respect to beliefs, and differences in the belief distri-

4 Model-free results on the importance of vaccine characteristics

4.1 Vaccine choices are sensitive to effectiveness and side effects

We begin with a simple result from the conjoint experiment showing the importance of effectiveness and side-effects, even among the vaccine hesitant. This simple cut was not pre-registered. Table 2 reports the fraction of participants choosing to vaccinate, in each vaccine intention group, for four different choice sets. The first row 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.[26] All the other choice sets in our experiment have a vaccine that is considerably better than the realistic vaccine. At the worst choice set we offer, take-up rises to 69 percent, with large increases among people uncertain about vaccinating or saying they do not intend to vaccinate. At the best choice set in our data, 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. The final row therefore shows take-up rates in a choice set where all vaccines have a positive side effect rate. Even here, vaccine take-up is substantially higher relative to the realistic choice set or the worst choice set.

4.2 Beliefs about flu risk and vaccination effects differ by hesitancy

Vaccine take-up is sensitive to vaccine characteristics. This sensitivity is important only if beliefs about vaccine characteristics are misaligned with actual characteristics. We therefore turn to investigating beliefs. We report mean objective beliefs in Table 3, Panel A, stratified by vaccine intention groups. The table shows two important facts. First, there is little disagreement across intention groups about national averages. Believed average flu rates are all about 35 percent and hospitalization rates vary from 12.6 to 14.2 percent.

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.[27] Beliefs about death given hospitalization are consistent

bution 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.

[26] The choice sets in this figure are 0, 1, 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.

[27] We report the CDC's range, i.e. the reported 95 percent confidence interval for risk estimates from Centers for Disease Control and Prevention (2024a), in the bottom row. The CDC reports overall rates per 100,000. We convert these to probabilities, unconditional for flu, conditional on flu for hospitalizations, and conditional on hospitalization for death.

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 beliefs about average risk may diverge from beliefs about personal risk, i.e. subjective beliefs. We therefore report the mean beliefs about personal risk and vaccine effects in Table 3, Panel B. We report subjective beliefs only for our uninformed participants, because these beliefs are likely better reflective of population beliefs. Appendix Figure A.2 shows the distribution of beliefs in the information and control provisions. In addition to the personal analogs of the beliefs in Panel A, the table reports beliefs about personal risk of severe flu (but not hospitalization) and personal side effect consequences of vaccination.[28]

In contrast to beliefs about national averages, subjective beliefs differ substantially across vaccine intention groups. Vaccine confident participants think they face roughly twice the flu risk as the most vaccine hesitant participants: the subjective flu rate, severe flu rate, hospitalization rate, and death rate are all roughly twice as large for people already vaccinated or intending to vaccinate as those not intending to vaccinate. Similarly, beliefs about vaccine effectiveness are more than 50 percent larger for vaccine confident as vaccine hesitant. (The "may or may not vaccinate" group is intermediate here.) Side effect beliefs are especially divergent. 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.

4.3 Side effect beliefs are excessively pessimistic

We have shown that choices are sensitive to vaccine side effects, even among the vaccine hesitant, and that, relative to vaccine confident participants, vaccine hesitant participants believe side effects are much worse. Although this belief is not consistent with clinical trial experience, it could be correct, because it is possible that people have heterogeneous reactions to vaccines. 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 final component of wave 1 was a series of incentivized tasks where participants chose between an unconditional follow-up payment of $3, or an unconditional of $2, plus an additional reward of $b if they showed they were vaccinated in the follow-up survey. Participants made 5 choices,

[28] These statistics are not reported by Centers for Disease Control and Prevention (2024a) so we do not elicit beliefs about them.

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.[29]

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:

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

We estimate separate specifications for the $10 switchers and $20 switchers, as well as pooled regressions.[30] 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 4 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.[31] We note, though, that the first

[29] 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.

[30] 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.

[31] Consistent with this interpretation of generally honest reporting, we found very few instances of apparently fraudulent

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. The implied TOT has a large confidence interval, so it is consistent with both very large negative effects and more modest effects. However, despite this wide range, we can reject the hypothesis that beliefs are accurate. We note the magnitude here is implausible, and we view the exact estimates with skepticism. We report this estimate because we pre-specified that we would do so. However we interpret the results conservatively as showing that vaccination does not worsen the experience of side effects (i.e. as a null rather than a very negative estimate), contrary to beliefs. Indeed this conclusion is evident in the reduced form: a negative effect of incentives on experienced side effect symptoms is inconsistent with the expected positive effect of vaccination.

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.3 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.

5 Estimating preferences over vaccines and simulating vaccine take-up

The results so far show that vaccine take-up is sensitive to characteristics and that beliefs about one of these characteristics, side effects, are biased. To simulate the consequences of correcting this bias, as well

verifications - four verifications appeared to have been manipulated, incorrect for the circumstance, or AI-generated.

as understand the consequences and desirability of other interventions, we turn to using the discrete choice experiment data to estimate vaccine preferences.

5.1 Preference specification, estimation, and identification

Preference specification: 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.

This specification allows for two explicit reasons not to vaccinate: side effects and needles. Additional preferences against vaccination load onto β_0, which captures preferences for a vaccine with no effectiveness, side effect, convenience, delivered via a nasal spray. β_{0i} therefore reflects a general preference against pharmaceuticals or dislike of foreign agents in their body. β_{0i} also reflects scientifically unwarranted beliefs, such as the view that vaccines cause autism or harm health in the long-run, because a person holding these beliefs would be opposed to even an innocuous vaccine. If these 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.

We allow for observed and unobserved heterogeneity in preferences. We account for observed heterogeneity by letting μ_i vary with participant characteristics. In our primary, pre-registered, μ_i varies with age (above/below 55), education (less than college, college or more), and score on the vaccine scale (above/below median for the sum of responses to the trust in doctors and follow government survey questions). In robustness tests we let μ_i vary directly with vaccine intentions, as well as a richer set of covariates. We account for unobserved heterogeneity via random coefficients with variance-covariate matrix Σ, assumed to be common Σ across people.

Estimation: We estimate the model using standard Bayesian methods, 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 C; the chains quickly exhibit good mixing and stability.

Identification: Our discrete choice 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 vaccine with a high reward, that could indicate a low ϵ or a low α_i. If some people consistently decline vaccines 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.

Simulating take-up: Simulating vaccination rates requires a preference distribution, which we estimate, but also beliefs about vaccine characteristics. In some contexts these beliefs are given, such as when we simulate vaccination rates in the absence of side effects. When the beliefs are not given, such as predicting take-up of vaccines outside of our experiment, we simulate demand using mean beliefs about effectiveness and side effects in each hesitancy group (measured in the uninformed condition). We use group-mean demand so that we can use the full sample but avoid contamination from information provision, which compresses cross-group belief differences. Our simulations are similar if we use a more sophisticated belief imputation approach, such as lasso (Appendix Table D.6). We assume vaccines are not convenient and delivered via needle.

Validation To not delay presenting our main results, we discuss model validation in detail in Section 7 below and summarize here. We validate the model with a number of out-of-sample prediction exercises. The model predicts a sharp in vaccination rates across intention categories. In a hypothetical choice not used in estimation, predicted take-up rates align with actual take-up on average. Estimated reward sensitivity correlates highly with sensitivity to actual in incentivized choices (not used in estimation). The model overpredicts actual vaccination rates, especially for people who say they intend to vaccinate but have not done so yet, consistent with our interpretation of the model as reflecting preferences over intentions rather than actions.

5.2 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 5. Side effect aversion is substantial but heterogeneous. While the average participant is willing to tolerate one percentage point greater side effect if it increases effectiveness by 2.15 percentage points, the median participant requires only a 0.6 percentage point increase in effectiveness to tolerate an additional 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.[32] Likewise for convenience: for the most sensitive group, a convenient vaccine is equivalent to a roughly $0.46 at the median. The coefficient on convenient is essentially zero.[33]

The estimates show a clear but varied willingness to trade side-effects against effectiveness. The median participant tolerates a 1 percentage point increase in side effects if it is accompanied by a 0.62 percentage point increase in effectiveness. However, there is a great deal of variation; some participants are almost indifferent to effectiveness, and others almost indifferent to side effects.

We similarly see a great deal of indifference in the willingness to pay to vaccinate. We define this WTP as the expected utility of vaccinating (given beliefs) scaled by reward sensitivity. The median participant modestly values vaccination, with a WTP of $22. This WTP is again highly dispersed. Many participants place a negative value on vaccination, some substantially so. Others place a very high value on vaccination. The low WTP of some participants is driven by a combination of pessimistic side effect beliefs and low "inside option" coefficient. This coefficient is literally the utility from a vaccine with no side effects or effectiveness. We interpret it as measuring like or dislike towards pharmaceuticals, i.e. a baseline preference or aversion to medicines. For most people this preference is (unsurprisingly) negative, meaning most people would not take a medicine with no effectiveness.

As our interest is in cross-group differences in preferences, we report median preferences by group in Table 6. The results show striking differences in sensitivity to vaccine characteristics by vaccine intention groups. Vaccine confident participants value effectiveness much more highly than do vaccine hesitant. While the median already vaccinated participant is willing to pay $9 for a one percentage point increase in effectiveness, for participants who do not intend to vaccinate, median willingness to pay is $1.69. Willingness to pay to avoid side effects is also higher for already vaccinated than for those not intending to

[32] 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.

[33] The low weight on convenience might be surprising in light of prior evidence. Banerjee et al. (2010); Duflo and Banerjee (2011) find in a developing country context that reliable, convenient vaccination has a large effect on vaccination take-up. Anderson et al. (ming); Worsham et al. (2020) find that pediatric flu vaccination is sensitive to availability of the flu vaccine in regular check-ups, pointing to important convenience effects. The small weight on convenience we find is likely attributable to (i) our focus on a developed country context where most people have relatively easy access to a pharmacy, and (ii) our use of stated preference data; people may not realize the importance of convenience for their choices.

vaccinate. This difference is driven in part by greater reward sensitivity among the vaccine hesitant, but also a lower value of effectiveness relative to side effects. The median already vaccinated participant values these characteristics nearly equally (with a marginal rate of substitution close to 1), while the median participant not intending to vaccinate values effectiveness at only 40 percent of side effects. Preferences over needles or convenience appear to contribute little to cross-group differences in vaccine preferences. All groups place a large, negative weight on the inside option coefficient. 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. Thus the key preference difference between vaccine hesitant and confident participants is the value they place on on vaccine effectiveness and side effects.

To put the value of effectiveness in quantitative perspective, we can use it to infer a value of a statistical life, under two additional assumptions. First we assume that participants value flu vaccine effectiveness only because of its mortality-reducing effects. This assumption is strong because preventing severe illness or hospitalization also has value, but it is plausible because death is plainly the most severe consequence of flu. This assumption implies that the value of a statistical life is

$$VSL = \frac{\beta_{VE}}{\beta_{reward}} \cdot 100 \cdot \frac{1}{\text{Pr(die from flu | unvaccinated)}}.$$

We have obtained the posterior distribution of $\beta_{VE}/\beta_{reward}$, the willingness to pay for a percentage point of effectiveness. We need an estimate of the chances of dying from flu in the absence of vaccination. We obtain this under our second assumption, which we make two ways. We first use the objective–i.e. CDC-reported (Centers for Disease Control and Prevention, 2024a)–age-specific mortality rate, adjusting for vaccination rates.[34] This implies an unconditional flu death rate of 2.7 per 100,000 for unvaccinated adults aged 18-49, 19 per 100,000 for ages 50-64, and 74 per 100,000 for ages 65 and older. As an alternative, we use subjective beliefs, setting the flu mortality rate equal to each intention group's median value (among participants not in the information arm). The last two rows of Table 6 show the implied VSL, by intention group, for these two approaches.

These approaches reveal large differences in VSL by hesitancy group. Under the objective measurement approach, vaccine confident and marginally hesitant participants have VSLs between $9 and $14 million. These are close to commonly used government benchmarks of $12-$13 million (U.S. Environmental Protection Agency, 2024; Kearsley, 2025). People not intending to vaccinate, however, have an

[34] We obtain the age-specific mortality rates unconditional on vaccination for 2024-25 from (Centers for Disease Control and Prevention, 2024a), $P(die, flu)$. The unvaccinated mortality rate, $P(die, flu|vacc = 0)$, is given by $P(die, flu) \cdot (1 - P(vacc) \cdot VE)$, where VE is vaccine effectiveness. We obtain the age specific vaccinates rates from Kahn et al. (2026), and vaccine effectiveness from Frutos et al. (2025), taking the average of two site-specific hospitalization vaccine effectiveness estimates.

implied VSL of $4 million, much lower than common benchmarks. Thus given objective flu risks in the 2024-25 season, vaccine hesitant participants appear to substantially undervalue flu vaccine effectiveness. However, subjective risk assessments suggest an alternative interpretation. Using subjective mortality to infer VSLs, we find that vaccine hesitant participants have very high VSLs, and vaccine confident participants have low VSLs. As we lack the data to validate these subjective mortality beliefs—which would require large sample sizes–we conclude that the low value of vaccine among reflects some combination of low VSL or flu risk perceptions well below objective benchmarks.

6 Counterfactual simulations

So far we have identified two clear differences between vaccine hesitant and vaccine confident participants: beliefs about side effect risks and preference over vaccine effectiveness. Vaccine hesitant participants' beliefs about side effect risks are biased relative to their own experiences. Their preferences for effectiveness imply VSLs that are lower than common benchmarks, at objective mortality risks (but not subjective mortality risk). We next show that moving beliefs and preferences to alternative levels—such as correcting side effect beliefs or setting preference for effectiveness to achieve benchmark VSLs—substantially increases vaccination rates, especially among vaccine hesitant participants. Finally, we investigate the policy implications of viewing high side effect beliefs or low VSLs as a bias.

6.1 Vaccine take-up implications of beliefs and preferences

So far we have shown substantial differences in beliefs about side effects and in preferences for effectiveness across vaccine intention groups. Vaccine hesitant participants believe side effects are worse than they are, and their preferences for effectiveness imply VSLs below normative benchmarks. We now show that these beliefs and preferences have substantial implications for vaccination rates. To do so, we simulate vaccination rates in four scenarios: (i) at baseline beliefs and preferences, (ii) eliminating side effect beliefs, (iii) changing preferences for effectiveness so each group's median VSL is equal to $12 million, and (iv) both.

The results in Table 7 show that beliefs and preferences are important for vaccination rates, individually and jointly. Changing beliefs about side effects alone would raise vaccination rates by about 4 percentage points overall, driven by a 20 percentage point increase in marginally hesitant participants and a 5 percentage point increase in the most hesitant participants (who have the most pessimistic beliefs but are also farthest from the margin of vaccinating). Setting preferences for effectiveness to align with normative benchmarks has a similar sized effect, again concentrated among the marginally hesitant. Both changes actually reduce take-up among vaccine confident participants, because their side effect views and preferences for effectiveness are above the benchmarks. The two changes have a complementary effect on

vaccination rates: changing both beliefs and preferences for effectiveness together would increase vaccination rates to 72 percent, and close much of the gap in predicted vaccination rates tween vaccine hesitant and confident participants.

Robustness: These simulation results are not especially sensitive to to modeling details and belief imputation. Appendix D reports simulation results from models that allow for preferences to depend on richer sets of covariates, include participants with reported contraindications, impute beliefs using richer specifications, and that stratify on experiment order (i.e. preference elicitation before belief elicitations). While the exact simulation results naturally vary across these approaches, our core finding is robust: side effect beliefs and preferences for effectiveness contribute substantially to vaccine take-up do not.

6.2 Policy implications of biased choices

Although changing beliefs or preferences raises vaccination rates, it may not be feasible to change beliefs. We therefore consider subsidy design to correct potentially biased vaccine choices.

Model setup: Allcott and Taubinsky (2015) provide a framework for evaluating policy decisions for individuals making potentially biased choices. In their framework, people choose an action such as vaccinating or not. The value to the individual of vaccinating is v, but choices are potentially distorted by a bias b, so that people vaccinate only if $v - b > p$, where p is the price they face. In our context, b could reflect excessively pessimistic beliefs or a VSL that is below normative benchmarks. We will sometimes refer to v as normative utility and $\tilde{v} \equiv v - b$ as decision WTP. We assume that v and b are heterogeneous with joint distribution $F(v, b)$. There is a social marginal cost of the action c.

The first best outcome is for everyone to vaccinate when $v > c$. However, assuming competitive pricing (so that $p = c$), people vaccinate when $v - b > c$, which will not typically achieve the first best outcome. Directly correcting biases, if feasible, achieves this outcome under competitive pricing (so that $p = c$). In practice ,correcting biases may be difficult because it is difficult to persistently move beliefs or increase the value of health.

In the absence of direct interventions on the bias, an alternative policy approach is to introduce a subsidy s. An optimal subsidy maximizes social welfare W, which is

$$W(s) = \int_{v,b} 1\{v - b > c - s\})(v - c)dF(v, b).$$

Social welfare aggregates up the (possibly negative) surplus from purchases. Non-purchase yields a surplus of zero. Purchase can yield a negative surplus for some individuals if their normative willingness to pay is below c but above $c - s$. Allcott and Taubinsky (2015) show that an optimal subsidy s^* satisfies

$$s^* = E[b|v - b = c - s^*]. \tag{5}$$

The optimal subsidy is equal to the average bias among people who are marginal to the subsidy, i.e. they are indifferent between purchasing or not at that subsidy level. If the bias were constant, this subsidy would equal the bias, consistent with the standard intuition for corrective taxation that a subsidy should equal the wedge between social marginal value and private marginal value.

Biases and subsidies in a mixed logit model Allcott and Taubinsky (2015) develop their framework for experiments explicitly measuring willingness to pay and dollarized bias. We show how to recover bias and decision WTP in a mixed logit model.

In such a model, people choose between a vaccine with characteristics X, or no vaccine. They have preference parameters β. They choose to vaccine if the decision utility from vaccinating, $\tilde{u} = \bar{u}(p, X, \beta) + \epsilon_1$, is greater than the decision utility from not vaccinating, ϵ_0. Mean utility $\bar{u}(p, X, \beta)$ is $\beta_p p + X'\beta$. We assume that ϵ_j are Type 1 EV distributed so that conditional on β, the purchase probability $D(p; X, \beta)$ takes the logit form. Decisions are subject to a dollarized bias b, so demand is not necessarily welfare-relevant. In our application, b will be a function of β and X, but we suppress that notation for the moment. Welfare-relevant utility is

$$u = \tilde{u} + \beta_p b. \tag{6}$$

Note that we multiply b by β_p to convert dollars to utils. We assume that the marginal cost of product is c, a subsidy s is available, and competition drives the pre-subsidy price to c before the subsidy.

Given a joint distribution of preferences and biases, we can determine the first best allocation by simulating demand with $b = 0$. We can find the optimal subsidy and second best allocation by calculating the average marginal bias at each subsidy level. The average marginal bias here is

$$AMB(s) = E[b|\epsilon_0 - \epsilon_1 = \bar{u}(p, X, \beta)].$$

Integrating over the logit error distribution, this implies,

$$AMB(s) = \frac{\int \beta_p b D(X, \beta) dF(\beta, b)}{\int \beta_p D(X, \beta) dF(\beta, b)}. \tag{7}$$

To calculate this integral, we average over many draws r from our posterior preference distribution:

$$AMB(s) \approx \Phi(s) = \frac{\sum_r \beta^r_{reward} b(X, \beta) D(X, \beta^r)}{\sum_r \beta^r_{reward} D(X, \beta^r)}. \tag{8}$$

Here we emphasize that b is potentially a function of observables (e.g. beliefs) and preference parameters (e.g. β_{VE}).

Measuring bias We consider three potential specifications of bias

1. Bias is due to misperceived side effects, where the correct perception is no side effect. Then the

dollarized bias is

$$b^{belief}(\beta, X) = \beta_{SE} SE_i / \beta_p rice. \tag{9}$$

2. Bias is due to normatively incorrect value of health, where the correct value corresponds to a hesitancy-group median VSL of $12 million. Letting $\tilde{\beta}_{VE}$ be the corrected VE, the bias is

$$b^{VSL}(\beta, X) = (\tilde{\beta}_{VE} - \beta_{VE}) VE / \beta_p rice. \tag{10}$$

   To obtain $\tilde{\beta}_{VE}$, we shift the mean of $\ln \beta_{VE}$ in each hesitancy group such that the median VSL in that hesitancy group is $12 million.

3. Bias is due to both misperceived side effects and normatively incorrect value of health. Then the dollarized bias is

$$b^{both}(\beta, X) = b^{belief}(\beta, X) + b^{VSL}(\beta, X). \tag{11}$$

Like other work on behaviorally motivated public policy, this approach requires the analyst to take a stand on sources of bias (Allcott and Taubinsky, 2015; Allcott et al., 2022). Our estimates show the first best allocation and optimal subsidy under an assumed bias but whether that assumption is appropriate is often a subjective question. We view the side effect beliefs as likely biased because we have shown they are inconsistent with actual experience. Whether VSLs are too low depends on how one views revealed preference for health and the subjective beliefs about mortality rates. By standard benchmarks (objective beliefs, government VSL targets), VSL is normatively biased for vaccine hesitant participants. However, respecting their choices about health and acknowledging their beliefs as plausible (given possible heterogeneity in flu mortality rates), we would not view VSL as normatively biased. We emphasize that we abstract from externalities such as reduction in the spread of disease or uninternalized reductions in health care costs.

Marginal cost We assume a marginal cost of $0, which corresponds to the status quo price of vaccines (as insurance subsidizes their price to zero). Thus an advantage of this marginal cost is that our specified biases are the only distortion. Accounting for other costs could raise or lower social marginal costs. For example, in a pediatric context, Anderson et al. (ming) estimate costs of procuring and distributing the flu vaccines of roughly $20 per dose, but they also estimate large reductions in health care spending from flu vaccinations. These reducitons are mostly borne by insurance companies and so represent a negative social marginal cost.

6.3 Optimal subsidies

We present the results in Table 8. Eliminating biased side effect beliefs would raise vaccination rates by about 5 percentage points, and raise welfare by about $17 per person. Eliminating biased preferences for vaccine effectiveness would have similarly sized effects, raising vaccination rates by about 5 percentage points as well, and raising welfare by $24. Eliminating both biases has a larger effect, raising vaccination rates by 17 percentage points and welfare by $80.

Attaining the first best requires directly changing preferences or beliefs (or individually targeted corrective subsidies). That is likely infeasible in practice. We report the optimal feasible subsidy in the next panel. The optimal subsidy given biased beliefs is about $12, producing a 2.5 percentage point increase in the vaccination rate. The optimal subsidy given biased VSL is $23, producing a 5 percentage point increase in vaccination rates. Despite substantial increases in vaccination rates, the welfare gains from optimal subsidies are small relative to the first best. This is because subsidies are poorly targeted: while they induce vaccination among some participants with modest bias, they also induce participants among participants with low bias and slightly negative normative WTP, undoing some of their welfare benefits.

The final column reports that the optimal subsidy is near infinite under both biases. While near infinite subsidies are infeasible, this result could also be understood as potentially justifying a mandate or efforts to achieve near-universal vaccination. It is surprising that the second-best optimal policy calls for near universal vaccination given that the first best vaccination rate is only 72 percent.

Figure 2 provides intuition for this surprising result. Each dot is a participant in our experiment, and for each participant we obtain median (dollarized) bias, which we plot against median normative willingness to pay (i.e. we obtain individual-specific posterior medians). For visual clarity, we top- and bottom-code bias and WTP. Participants with positive normative WTP vaccinate in the first best. They choose to vaccinate only if their normative WTP is greater than their bias plus any subsidy they receive. Participants below the shifted 45 degree lines would vaccinate at various subsidies. For example, participants below the line through 0 would vaccinate without any subsidy.

The figure shows that there is a diffuse mass of participants with both large normative WTP and large bias. These are mostly participants indicating they may or may not vaccinate. Inducing these participants would generate large welfare gains, and would require large subsidies (given their large biases). However, the figure also shows that there is a mass of participants with slightly negative normative WTP, and bias between $0 and $400. Most of these participants are in the “do not intend to vaccinate" or “may or may not vaccinate" groups. Large subsidies also induce these participants to vaccinate, but doing so *reduces* social welfare. The first best only induces the first group to vaccinate. In the second best, however, the large gains from the first group outweigh the modest losses from the first group, justifying large subsidies. For marginal subsidies above several hundred dollars, there are essentially no participants with negative normative WTP, meaning all remaining participants have high and positive normative WTP but also high

bias. Inducing them to vaccinate increases welfare but requires large subsidies.

While we focus on the social marginal cost $c = 0$ case, we also present welfare and optimal subsidies for $c = 20$. We believe $c = 0$ is more appropriate because it implicitly captures an externality of vaccination, reduced downstream health care costs. In the absence of such an externality, $c = 20$ would be an appropriate benchmark. At the higher marginal cost, we find that optimal take-up and subsidies are generally lower. The optimal belief-correcting subsidy is $7 instead of $12, the optimal VSL-correcting subsidy is $14 instead of $23, and the optimal subsidy to correct both b biases is $32. These optimal subsidies recover a small share of the welfare improvement of moving to the first best.

7 Validation

7.1 Validating beliefs

We show that beliefs are meaningful constructs, separate from preferences or attitudes towards vaccination, because they are predicted by risk factors, they predict risk outcomes, and they predict choices, conditional on preferences. Age and chronic conditions predict beliefs about flu risk severity. Our age variables are indicators for age 18-39, 40-54, 55-64, and 65 or older. We measure chronic conditions with an indicator for at least one of (self-reported) asthma, diabetes, heart disease, lung disease, kidney disease, and pre-diabetes. We focus on any condition rather because most conditions are rare. Appendix TableA.5 shows that age and risk factors are correlated with believed risk in an appropriate way.[35] Beliefs are also (imperfectly) predictive of outcomes. Appendix Table A.6 shows that beliefs about side effect experiences and flu risk are predictive of subsequently reported side effect experiences and flu risks. These beliefs are not necessarily well calibrated but they capture participant information about their likely risks and experiences.

Beliefs also predict vaccine take-up, distinct from preferences. To show this we regress vaccination rates on our belief measures and on individual-level predicted vaccination of a benchmark vaccine (50 percent effective, 1 percent side effect rate) given individual level preference estimates. Appendix Table A.7 shows that beliefs and preferences are independently predictive of vaccination. Together this evidence shows that subjective beliefs contain information about personal risks, are informative for vaccination choice, and are distinct from preferences.

[35] 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 flu, but much more likely to experience severe cases of the flu if they get it Centers for Disease Control and Prevention (2024a).

7.2 Model validation via out-of-sample predictions

We validate our model by simulating choices in several out-of-sample contexts. First we show that model-predicted choice rates in our baseline vaccine scenario are close to actual choices. Table 9 shows simulated take-up of the baseline vaccine, by intention group, along with the fraction of participants actually choosing the baseline vaccine in hypothetical choices. This is out of sample because we do not use this choice in estimation. As we would expect for an out of sample prediction, choice rates do not perfectly match. The overall rates are quite close. They vary less by intention group than we see in the actual data, however.

Next we compare model predicted vaccination rates, given group beliefs, to actual vaccination rates. For already vaccinated people the vaccination rate is of course 1. For the other groups, we use vaccination rates recorded in wave 2 or wave 3. As we experienced some attrition, vaccination rates are missing for some participants; we exclude such participants from the simulation and the comparison. We find a very steep gradient in predicted simulated vaccination rates across intention groups. Although expected, this provides simple evidence on the validity of the preference estimates. However, the prediction substantially over-predicts actual vaccinations, especially for people who intend to vaccinate (but have not done so by late in the season) and people who say they may or may not vaccinate. This overprediction is consistent with an action-intention gap in vaccination documented elsewhere (Milkman et al., 2011; Bronchetti et al., 2015), and is why we interpret our estimates reflecting preferences over intentions. Our model predictions are also informative of individual level vaccination rates. Appendix Figure A.4 shows a binned scatter plot of actual vaccination against model-predicted vaccination rates. Actual vaccination rates rise clearly with predicted vaccination rates, with an overall slope of 0.72.

7.3 Model validation via 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.[36]

[36] To make sure participants understood the incentives, we first asked them about bonus that hypothetical participants

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: [37]

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

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

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

where $1(\cdot)$ is the indicator function.

We report the actual and model-predicted take-up of incentives in Appendix Figure A.5. 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. The low take-up of modest incentives, but not high incentives, suggests that participants perceive a fixed cost of taking-up incentives, for example a cost of verifying their vaccination status. To accommodate this possibility we simulated choices with a cost of receiving the incentive, introducing a single additional parameter.[38] The first from this model is better; predicted take-up rises steeply with incentives. This suggests that our estimated price sensitivities are predictive of responses to actual prices.

Indeed, at an 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

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.

[37] 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.

[38] We introduce verification costs as follows. We assume that people choose incentives as if the reward from vaccinating is $b-c$, where b is the offered bonus and c is a verification. We assume the verification cost is exponentially distributed with mean λ, indepnedent of all other parameters. We require heterogeneous costs because otherwise predicted take-up would be zero (since if $c > b$ it is better not to take the incentive), and we use the exponential distribution because it is always positive and simple to implement. We find the value of λ which minimizes the mean squared prediction error in take-up rates. This is roughly \$5.

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 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 10 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.

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

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

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

Allcott, H. (2015). Site selection bias in program evaluation. *The Quarterly Journal of Economics 130*(3), 1117–1165.

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.

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

Allcott, H., D. Cohen, W. Morrison, and D. Taubinsky (2025, May). When do "nudges" increase welfare? *American Economic Review 115*(5), 1555–96.

Allcott, H., J. Kim, D. Taubinsky, and J. Zinman (2022). Are high-interest loans predatory? theory and evidence from payday lending. *The Review of Economic Studies 89*(3), 1041–1084.

Allcott, H., B. B. Lockwood, and D. Taubinsky (2019). Should we tax sugar-sweetened beverages? an overview of theory and evidence. *Journal of Economic Perspectives 33*(3), 202–227.

Allcott, H. and D. Taubinsky (2015). Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review 105*(8), 2501–2538.

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., 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*.

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.

Baicker, K., S. Mullainathan, and J. Schwartzstein (2015). Behavioral hazard in health insurance. *The Quarterly Journal of Economics 130*(4), 1623–1667.

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*.

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.

Bernheim, B. D. and D. Taubinsky (2018). Chapter 5 - behavioral public economics. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics - Foundations and Applications 1*, Volume 1 of *Handbook of Behavioral Economics: Applications and Foundations 1*, pp. 381–516. North-Holland.

Blue, C. L. and J. M. Valley (2002). Predictors of influenza vaccine: acceptance among healthy adult workers. *Aaohn Journal 50*(5), 227–233.

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, F. H. Schneider, S. Meier, D. Pope, and E. Wengström (2021). Monetary incentives increase covid-19 vaccinations. *Science 374*(6569), 879–882.

Center for Disease Control (2025a). Flu diease burden: Past seasons. `https://www.cdc.gov/flu-burden/php/data-vis/past-seasons.html`, 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. `https://www.cdc.gov/flu-burden/php/about/index.html`, Accessed: 2025-02-17.

Centers for Disease Control and Prevention (2024b). Seasonal flu vaccine effectiveness studies. `https://www.cdc.gov/flu-vaccines-work/php/effectiveness-studies/index.html`, Accessed: 2025-02-17.

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

Chandra, A., E. Flack, and Z. Obermeyer (2024). The health costs of cost sharing. *The Quarterly Journal of Economics 139*(4), 2037–2082.

Chang, T. Y., M. Jacobson, M. Shah, R. Pramanik, and S. B. Shah (2026). Impact of reminder messages, with and without financial incentives, on influenza vaccination: A randomized trial in a california health system. *Vaccine 75*, 128263.

Chapman, G. B. and E. J. Coups (1999). Predictors of influenza vaccine acceptance among healthy adults. *Preventive medicine 29*(4), 249–262.

Cheipesh, O. (2025). Just the flu? the labor market spillovers of influenza vaccination.

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.

Daysal, N. M., H. Ding, M. Rossin-Slater, and H. Schwandt (2021). Germs in the family: The short- and long-term consequences of intra-household disease spread. Technical report, National Bureau of Economic Research.

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.

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.

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

Frutos, A. M., S. Cleary, E. L. Reeves, et al. (2025). Interim estimates of 2024–2025 seasonal influenza vaccine effectiveness — four vaccine effectiveness networks, united states, october 2024–february 2025. *MMWR Morbidity and Mortality Weekly Report 74*, 83–90. URL: https://www.cdc.gov/mmwr/volumes/74/wr/mm7406a2.htm (Accessed: 2026-04-07).

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. Williams (2022). Optimal vaccine subsidies for endemic diseases. *International Journal of Industrial Organization 84*, 102840.

Goodkin-Gold, M., M. Kremer, C. M. Snyder, and H. Williams (2024). Optimal vaccine subsidies for epidemic diseases. *Review of Economics and Statistics 106*(4), 895–909.

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.

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.

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.

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

Janz, N. K. and M. H. Becker (1984). The health belief model: A decade later. *Health education quarterly 11*(1), 1–47.

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.

Kahn, K. E., A. Jain, T. Zhou, M.-C. Hung, A. Srivastav, P.-J. Lu, J. L. Kriss, C. L. Black, and J. A. Singleton (2026, January). Flu vaccination coverage, United States, 2024–25 influenza season. `https://www.cdc.gov/fluvaxview/coverage-by-season/2024-2025.html`. Accessed: 2026-03-19.

Kearsley, A. (2025, January). Hhs standard values for regulatory analysis, 2025. Technical report, Office of the Assistant Secretary for Planning and Evaluation, U.S. Department of Health and Human Services. `https://aspe.hhs.gov/sites/default/files/documents/639756a60fbe7e51786bcec176ad52f1/Standard-RIA-Values-2025.pdf`.

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.

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.

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., J. Beshears, J. J. Choi, D. Laibson, and B. C. Madrian (2011). Using implementation intentions prompts to enhance influenza vaccination rates. *Proceedings of the National Academy of Sciences 108*(26), 10415–10420.

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.

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.

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.

Shahrabani, S. and U. Benzion (2010). Workplace vaccination and other factors impacting influenza vaccination decision among employees in israel. *International journal of environmental research and public health 7*(3), 853–869.

Shahrabani, S., U. Benzion, and G. Yom Din (2009). Factors affecting nurses' decision to get the flu vaccine. *The European Journal of Health Economics 10*(2), 227–231.

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.

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.

Thirumurthy, H., K. L. Milkman, K. G. Volpp, A. M. Buttenheim, and D. G. Pope (2022, 03). Association between statewide financial incentive programs and covid-19 vaccination rates. *PLOS ONE 17*(3), 1–7.

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

Tsutsui, Y., U. Benzion, and S. Shahrabani (2012). Economic and behavioral factors in an individual's decision to take the influenza vaccination in japan. *The Journal of Socio-Economics 41*(5), 594–602.

U.S. Environmental Protection Agency (2024). Mortality risk valuation. `https://www.epa.gov/environmental-economics/mortality-risk-valuation`.

U.S. Food and Drug Administration (2026). Fluarix (influenza vaccine) injectable suspension, for intramuscular use: Highlights of prescribing information (2025–2026 formula). Package insert.

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	2024 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	0.51
High school or less	0.11	0.27
Some college	0.27	0.30
College or more	0.61	0.32
Income < $25,000	0.12	0.15
Income $25,000 - $50,000	0.20	0.23
Income $50,000 - $100,000	0.36	0.28
Income > $100,000	0.30	0.31
Already vaccinated	0.25	.
Intend to vaccinate	0.24	.
May or may not vaccinate	0.21	.
Do not intend to vaccinate	0.30	.
Flu vaccination	0.43	0.40

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: Model-free evidence on the importance of vaccine characteristics

Choice	Pr(vaccinate), by vaccine intention				
Scenario	Overall	Vaccinated	Intend to	May or may not	Don't intend to
Baseline choice set	0.43	0.76	0.64	0.30	0.06
Worst choice set with three options	0.69	0.96	0.93	0.75	0.23
Best choice set	0.82	1.00	1.00	0.94	0.45
Best choice set with side effects	0.76	0.96	0.96	0.85	0.36

Notes: Table reports by hesitancy group the share of people choosing to vaccinate in four hypothetical scenarios. In the baseline choice set, the options are no vaccine or a vaccine with no reward 50% effectiveness, 5% severe side effect rate, needle, inconvenient. In the worst choice set, the options are no vaccine, vaccine A with no reward, 75% effectiveness, 0% severe side effect rate, nasal delivery, inconvenient; or vaccine B with a $5 reward, 50% effectiveness, needle delivery, and convenient. In the best choice set, the options are no vaccine, vaccine A with a $25 reward, 90% effectiveness, 0% severe side effect rate, needle delivery, convenient; or vaccine B with a $5 reward, 50% effectiveness, 1% side effect rate, nasal delivery, and inconvenient. In the best choice set with side effects, the options are no vaccine, vaccine A with a $25 reward, 75% effectiveness, 5% side effect rate, and nasal spray, inconvenient; or vaccine B with a $0 reward, 90% effectiveness, 1% side effects, needle, and convenient.

Table 3: Beliefs about flu risk and the flu vaccine vaccines

Intention group:	Already vaccinated	Intends to	May or may not	Does not intend to
Panel A. Objective beliefs				
Flu rate	34.85	37.79	35.76	35.04
	(0.77)	(0.91)	(0.93)	(0.78)
Hospitalization rate	12.86	14.21	13.54	12.65
	(0.46)	(0.58)	(0.60)	(0.49)
Death rate	8.21	8.79	7.77	6.44
	(0.44)	(0.54)	(0.48)	(0.33)
Vacine effectiveness	71.97	73.54	71.75	67.21
	(0.64)	(0.66)	(0.78)	(0.79)
CDC range:	10-20	0.6-2.5	3-18	40-60
Panel B. Subjective beliefs				
Flu rate	43.11	39.52	29.87	20.08
	(1.39)	(1.52)	(1.42)	(1.15)
Severe flu rate	56.15	48.83	42.32	33.95
	(1.78)	(1.85)	(2.03)	(1.62)
Hospitalization rate	11.84	11.25	7.00	4.01
	(0.88)	(0.92)	(0.87)	(0.40)
Death rate	5.96	5.64	3.36	2.33
	(0.75)	(0.83)	(0.57)	(0.37)
Vacine effectiveness	67.55	68.25	60.23	38.68
	(1.09)	(1.13)	(1.44)	(1.57)
Side effect rate	-0.98	-2.10	7.84	15.27
	(1.49)	(1.36)	(1.48)	(1.58)

Notes: Table reports mean beliefs by intention group, for the indicated belief. Panel A reports beliefs about the CDC's report on the indicated aspects of flu risk or vaccines, and the final row is the range (given 95% confidence interval or annual variation) reported by Centers for Disease Control and Prevention (2024a). Hospitalization rates are conditional on flu and death rate is conditional on hospitalization. Panel B reports mean beliefs about personal flu risk or vaccine consequences. In panel B, flu risk is conditional on not vaccinating; hospitalization rate is (additionally) conditional on flu and death rate is (additionally) conditional on hospitalization. The sample is limited to people who did not receive information treatment.

Table 4: Side effect beliefs are excessively pessimistic

Group	Switch at $10 (1)	Switch at $20 (2)	Pool (3)
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.

Table 5: Preference estimates

			Percentiles				
Preference parameter	Mean	SD	2.5	25	50	75	97.5
Reward	0.13	0.26	0.01	0.03	0.06	0.14	0.71
Vaccine Effectiveness	0.45	0.63	0.03	0.12	0.25	0.52	2.14
Side effect rate	-1.33	4.68	-7.94	-1.12	-0.4	-0.14	-0.02
Needle	0.28	2.13	-3.97	-1.12	0.28	1.68	4.5
Convenient	0.04	0.79	-1.51	-0.49	0.04	0.57	1.6
Inside option	-17.36	14.12	-45.15	-26.89	-17.3	-7.69	10.17
MRS Effectiveness for side effects	2.09	7.04	0.03	0.22	0.62	1.75	13.14
WTP for vaccine ($)	528.93	6558.14	-2208.84	-135.74	22.33	357.09	6457.56

Notes: Table reports the posterior mean, standard deviation, and 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 6: Median preferences parameters by vaccine intention group

	Vaccine intention			
Parameter	Already vaccinated	Intends	May or may not	Does not intend to
Reward sensitivity	0.0471	0.0523	0.0659	0.0823
WTP for Effectiveness	8.78	6.41	3.57	1.7
WTP for Side effects	-8.96	-8.12	-6.33	-4.22
WTP for Needle	9.26	7.31	1.36	-2.26
WTP for Convenience	0.0526	0.36	0.467	0.36
WTP for Inside Option	-267	-219	-174	-198
MRS Effetiveness vs. side effects	0.941	0.76	0.533	0.391
WTP to vaccinate	322	230	-26.6	-233
VSL at objective risks ($Millions)	11.9	13.7	9.7	4.03
VSL at subjective risks ($Millions)	2.98	6.01	10.6	33.9

Notes: Table reports the posterior median of each preference parameter, by vaccine intention group. Willingness to pay is the coefficient for a characteristic divided by reward sensitivity. The "inside option" can be interpreted as the willingness-to-pay for a vaccine with no effectiveness, no side effects, no needle, and not convenient. MRS is the ratio $-\beta_{VE}/\beta_{SE}$. VSL calculations are explained in the text.

Table 7: Simulated and counterfactual vaccination rates

Intention group	Already vaccinated (1)	Intend to (2)	May or may not (3)	Do not intend (4)	Overall (5)
Baseline	0.93	0.91	0.39	0.05	0.55
No SE	0.92	0.88	0.59	0.1	0.59
Normative VSL	0.93	0.85	0.49	0.22	0.61
Both	0.92	0.81	0.73	0.46	0.72

Notes: Table reports the simulated vaccination rate for each intention group, and overall. The first row is at baseline beliefs and preferences. The second row sets side effect beliefs to zero. The third row adjusts preferences for effectiveness so that each group has a median value of a statistical life of $12 million. The final row makes both changes.

Table 8: Take-up, welfare, and subsidies under alternative biases

Source of bias:	Beliefs	VSL	Both
No Subsidy ($c = 0$)			
Welfare	830	859	839
Pr(vacc)	0.548	0.548	0.548
First Best ($c = 0$)			
Welfare	848	883	920
Pr(vacc)	0.594	0.608	0.717
Optimal Subsidy ($c = 0$)			
Welfare	830	859	863
Pr(vacc)	0.575	0.6	0.983
Subsidy	11.9	23.3	3000
No Subsidy ($c = 20$)			
Welfare	819	848	828
Pr(vacc)	0.503	0.503	0.503
First Best ($c = 20$)			
Welfare	837	872	907
Pr(vacc)	0.537	0.551	0.636
Optimal Subsidy ($c = 20$)			
Welfare	819	848	828
Pr(vacc)	0.517	0.534	0.574
Subsidy	6.76	13.9	31.5

Notes: Table reports welfare, vaccination rates, and subsidies under alternative scenarios. Each column models a different source of bias in vaccination choices: beliefs about side effects are incorrect (and should be zero), weight on vaccine effectiveness is incorrect (and should imply a median VSL of $12 million in each intention group), and both. Under each bias, we report baseline take-up and welfare, first best take-up and welfare. c is the assumed social marginal cost of a flu vaccination.

Table 9: Out-of-sample fit

Intention group	All	Already vaccinated	Intend to	May or may not	Do not intend
A. Baseline hypothetical vaccine					
Simulated probability	0.401	0.652	0.571	0.352	0.0866
Choice probability	0.427	0.761	0.642	0.298	0.0621
B. Actual vaccinations					
Simulated probability	0.549	0.93	0.905	0.38	0.05
Vaccination rate	0.431	1	0.532	0.161	0.0404

Notes: Table reports simulated and actual choice probabilities. In Panel A., we simulate take up of the baseline vaccine (characteristics described in the notes to Table 2) vs. no vaccine, a hypothetical choice we offered participants. In Panel B, we simulate take up of actual vaccination, for people with vaccination data (measured in waves 2 and 3). To simulate take up we impute beliefs at the intention group mean in the uninformed condition.

Table 10: Stated preferences predict incentivized choices

Reward	$1.50	$3	$10	$20	$50
	(1)	(2)	(3)	(4)	(5)
Pr(takeup)	0.36	0.46	0.57	0.54	0.46
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Reward beta	0.05	0.05	0.21	0.42	0.38
	(0.04)	(0.05)	(0.07)	(0.07)	(0.06)
N	1,756	1,758	1,755	1,749	1,756

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.

Figure 1: Experiment design

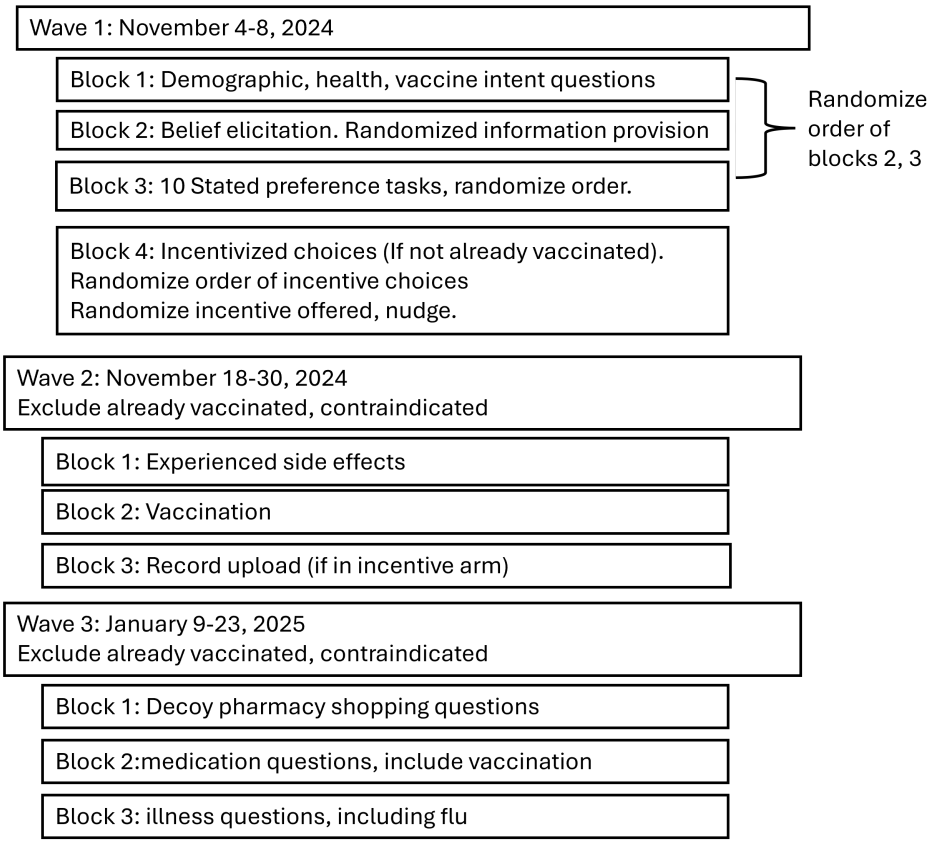
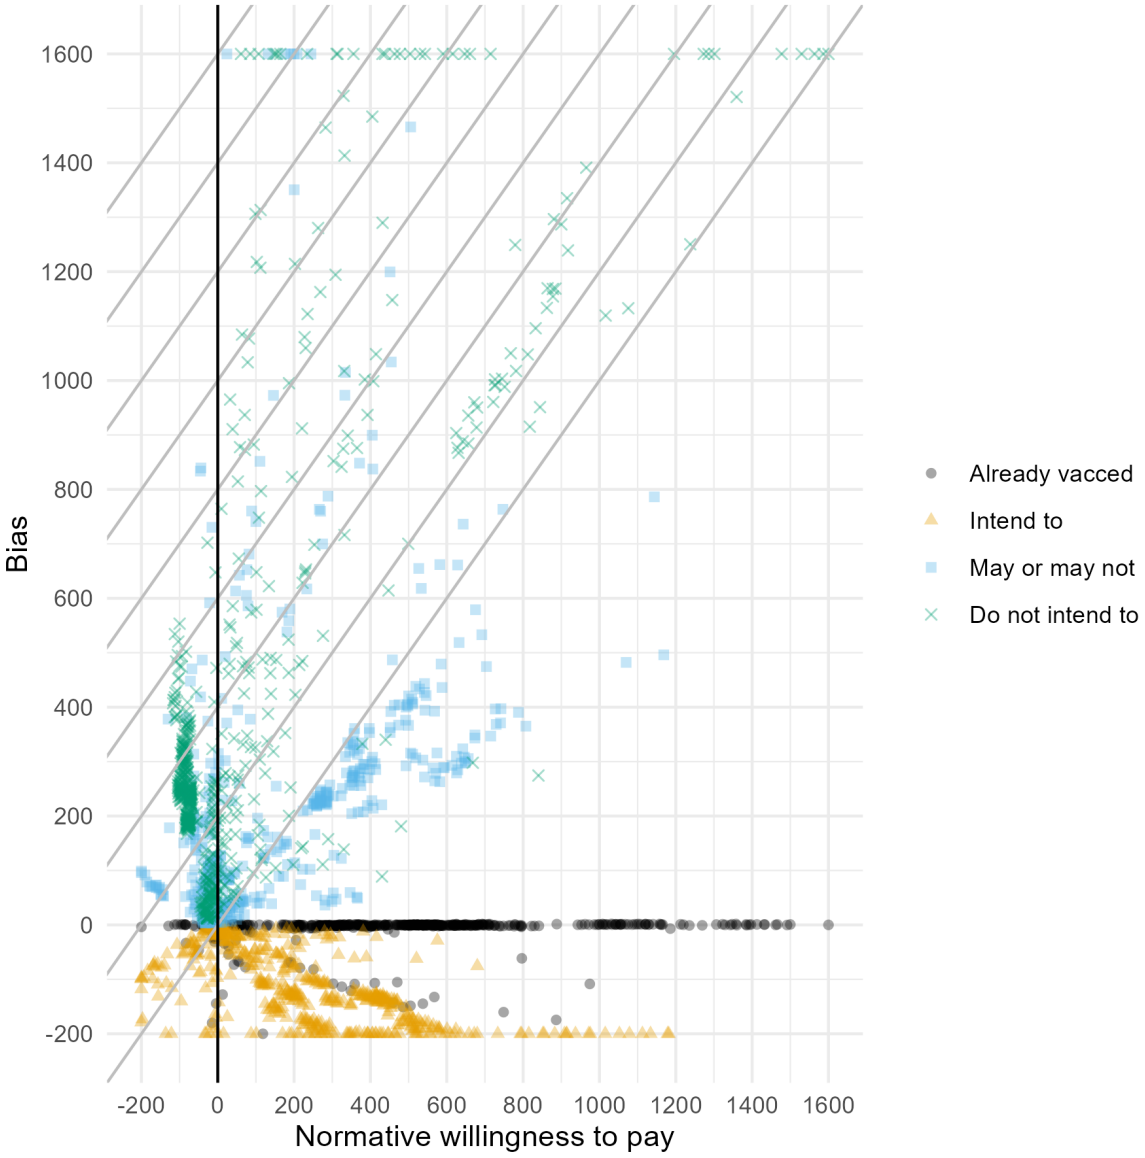


Figure 2: Bias and normative willingness-to-pay



Notes: Figure shows a scatter plot of (individual level posterior median) bias against willingness to pay. The bias assumes that a normatively appropriate VSL of $12 million and a side effect rate of zero. People with positive normative WTP vaccinate in the first best. People actually vaccinate if normative WTP minus bias is greater than the offered subsidy (as we assume a cost of zero). The solid lines illustrate different subsidies (starting at $0 and increasing by $200). People below the solid line choose to vaccinate for that subsidy.

Appendices

For online publication only

A Appendix Exhibits **50**

B Belief elicitation questions **59**

C MCMC Diagnostics **63**

D Sensitivity analysis **65**
D.1 Further details on contraindications 65
D.2 Sensitivity analysis of counterfactual simulations 66

E Pre-registration modifications and adjustments **70**

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: 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.

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: Risk factors predict flu risk beliefs

Belief	Pr(Flu) (1)	Pr(Severe |Flu) (2)	Pr(Hospitalization| Flu) (3)	Pr(Death | Hospitalization) (4)
Age 40-54	-0.959	3.235	-1.295	-0.465
	(1.718)	(2.255)	(0.857)	(0.709)
Age 55-64	-5.728	3.207	0.839	1.305
	(2.519)	(3.295)	(1.489)	(1.219)
Age 65+	-7.849	0.693	2.891	2.650
	(2.148)	(2.969)	(1.466)	(1.250)
At least one risk factor	8.908	8.185	5.485	1.403
	(1.846)	(2.271)	(1.088)	(0.858)
Joint significance	0.000	0.002	0.000	0.045
Observations	1,176	1,176	1,176	1,176

Notes: Table reports the results of a regression of the indicated subjective flu risk belief on age indicators and an indicator for having at least one risk factor in the set asthma, diabetes, heart disease ,lung disease, or kidney disease. The sample excludes participants in the information arm. The rows report the coefficients (robust standard errors in parentheses). The row labeled joint significance gives the p-value of the test that coefficients are jointly zero. Flu risks are measured here on a 0-100 scale.

Table A.6: Beliefs predict experienced side effects and reported flu

Outcome	Experienced side effect (1)	Reported flu (2)
Believed SE symptom rate	0.97	0.05
	(0.13)	(0.10)
Believed flu rate if unvacced	-0.11	0.16
	(0.11)	(0.08)
Observations	715	723

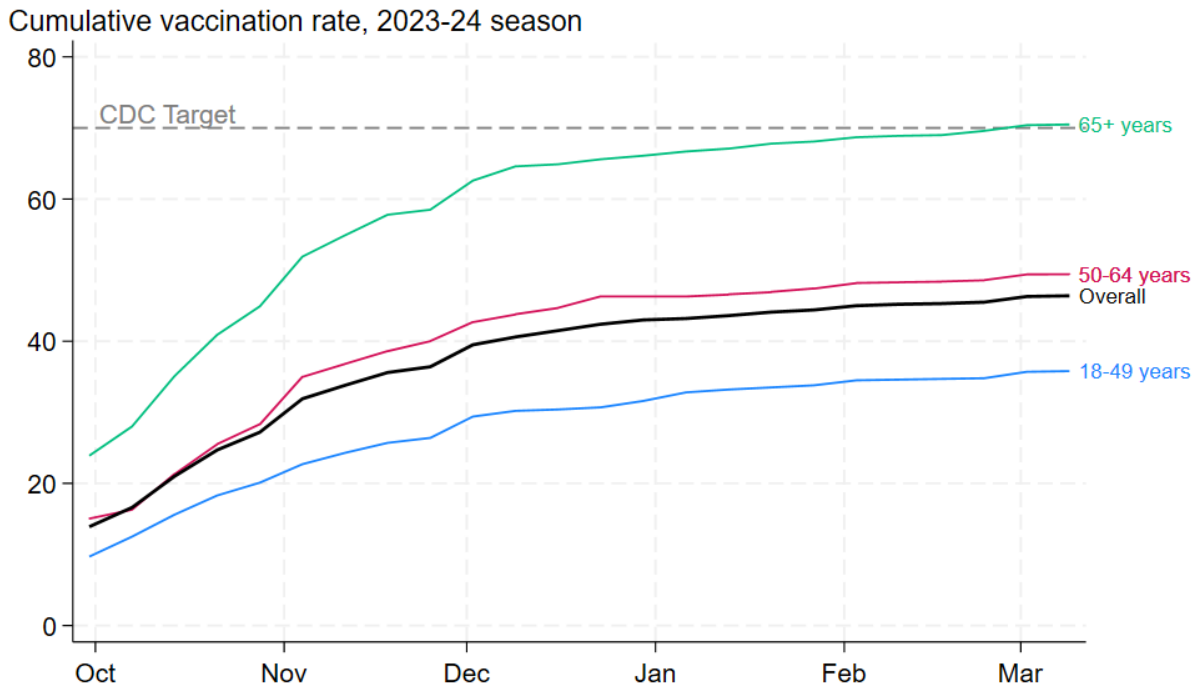
Notes: Table reports the results of a regression of experienced side effect or experienced flu on the indicated beliefs, plus indicators for age categories and an indicator for flu risk factors. We instrument for quantitative belief measures using indicators for responses to multiple choice versions of the questions, to correct for measurement error. We exclude participants in the information arm. Robust standard errors in parentheses.

Table A.7: Vaccination is predicted by beliefs and stated preferences

Beliefs:	Flu risk only	Vacc effect only	Both	+Preferences
Pr(Flu)	0.90		0.64	0.45
	(0.09)		(0.09)	(0.09)
Pr(Severe | flu)	0.15		0.11	0.07
	(0.07)		(0.07)	(0.07)
Pr(Hospitalization |flu)	0.32		0.12	0.04
	(0.24)		(0.23)	(0.22)
Pr(Death | Hospital)	0.10		0.17	0.08
	(0.30)		(0.28)	(0.26)
Side effect rate		-0.26	-0.21	-0.15
		(0.08)	(0.08)	(0.07)
Vaccine effectiveness		0.73	0.60	0.40
		(0.06)	(0.07)	(0.07)
Conjoint-implied vaccination probability				0.35
				(0.03)
Observations	1,022	1,021	1,021	1,021

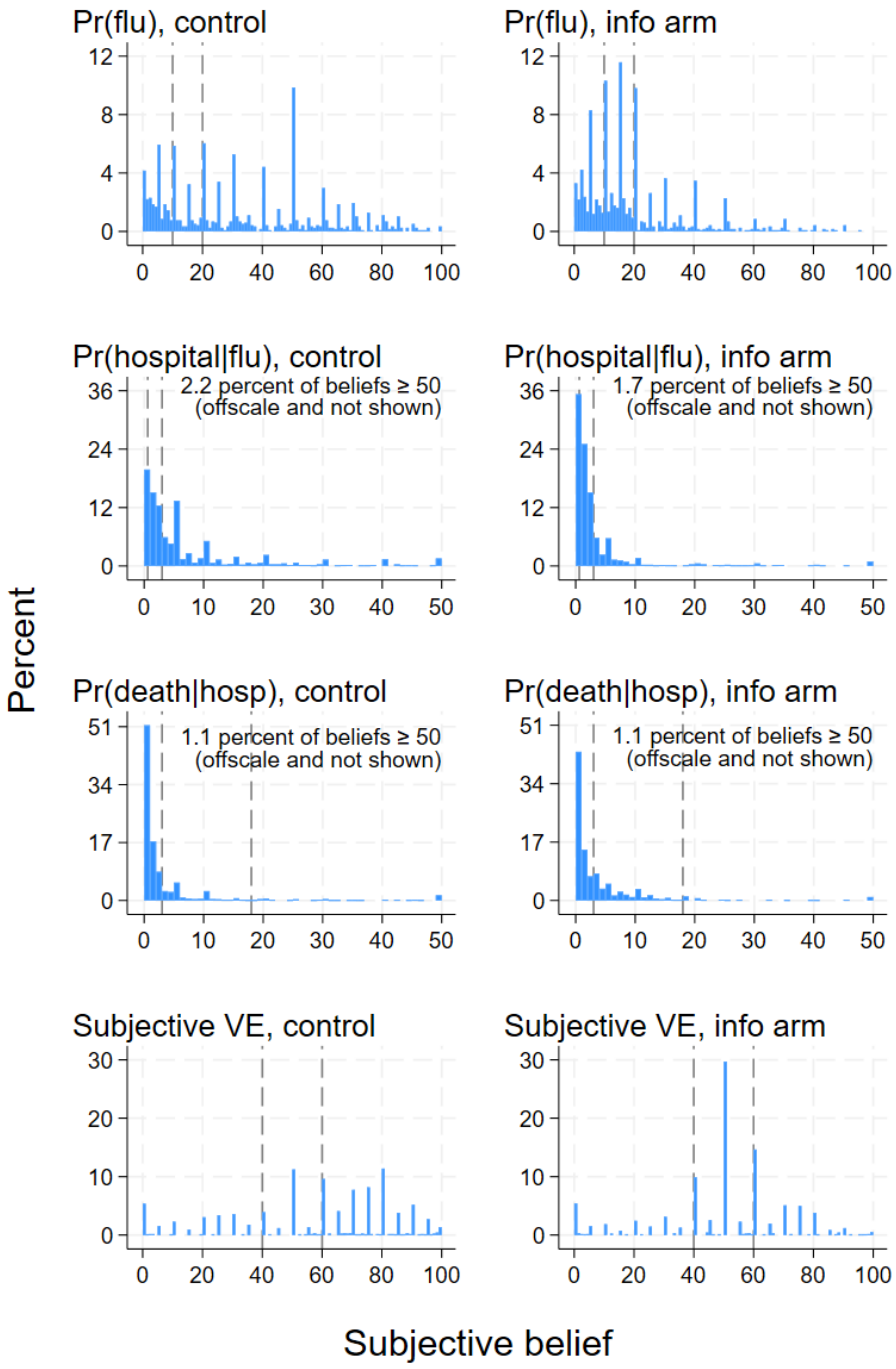
Notes: Table reports results from a regression of vaccination status on subjective beliefs. Vaccination status is an indicator for already vaccinated as of wave 1, or self-reported vaccination in wave 3. Sample excludes people with missing belief measures or missing vaccinated status (which happens for people not vaccinated in wave 1, with no follow up data in wave 3), and people in the information provision arm. 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." The conjoint implied vaccination probability is each individual's posterior mean predicted vaccination rate given her posterior mean β, assuming 5 percent side effect rate and 50 percent vaccine effectiveness (i.e. fixing beliefs at roughly their objective level). 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: Distribution of subjective beliefs



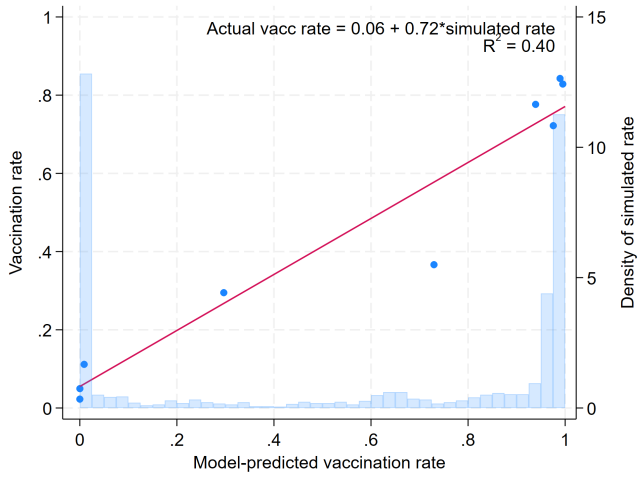
Notes: Figure plots the percent of respondents in each 1 percentage point bin of beliefs, for our subjective belief measures for which we provided information on CDC estimates. The left panels show the control group (uninformed) distribution and the right the information provision group distribution. The vertical line shows the provided ranges.

Figure A.3: Example choice set

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

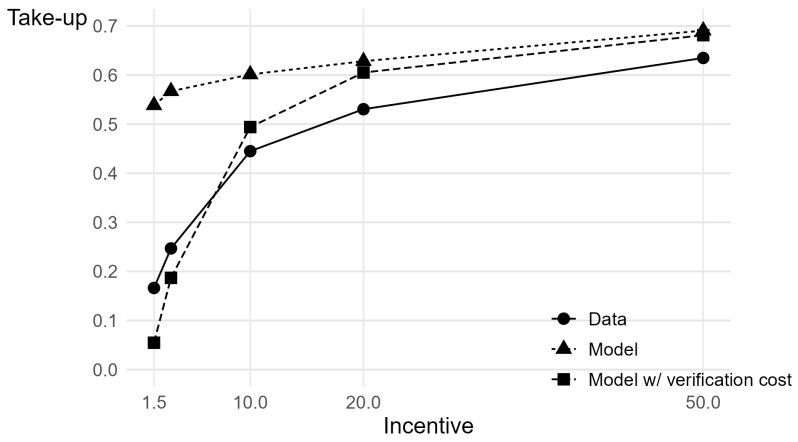
Notes: Figure shows how we displayed product characteristics in the conjoint design. Below this table we gave respondents a button to click for a reminder of the meaning of the different characteristics.

Figure A.4: 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 A.5: 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. We simulate from our baseline model, as well as from a model with a verification cost which must be paid to claim the bonus. This model assumes verification costs are exponentially distributed and we plot results for the best-fitting verification mean verification cost.

B Belief elicitation questions

This appendix describes our belief questions and side-effect 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

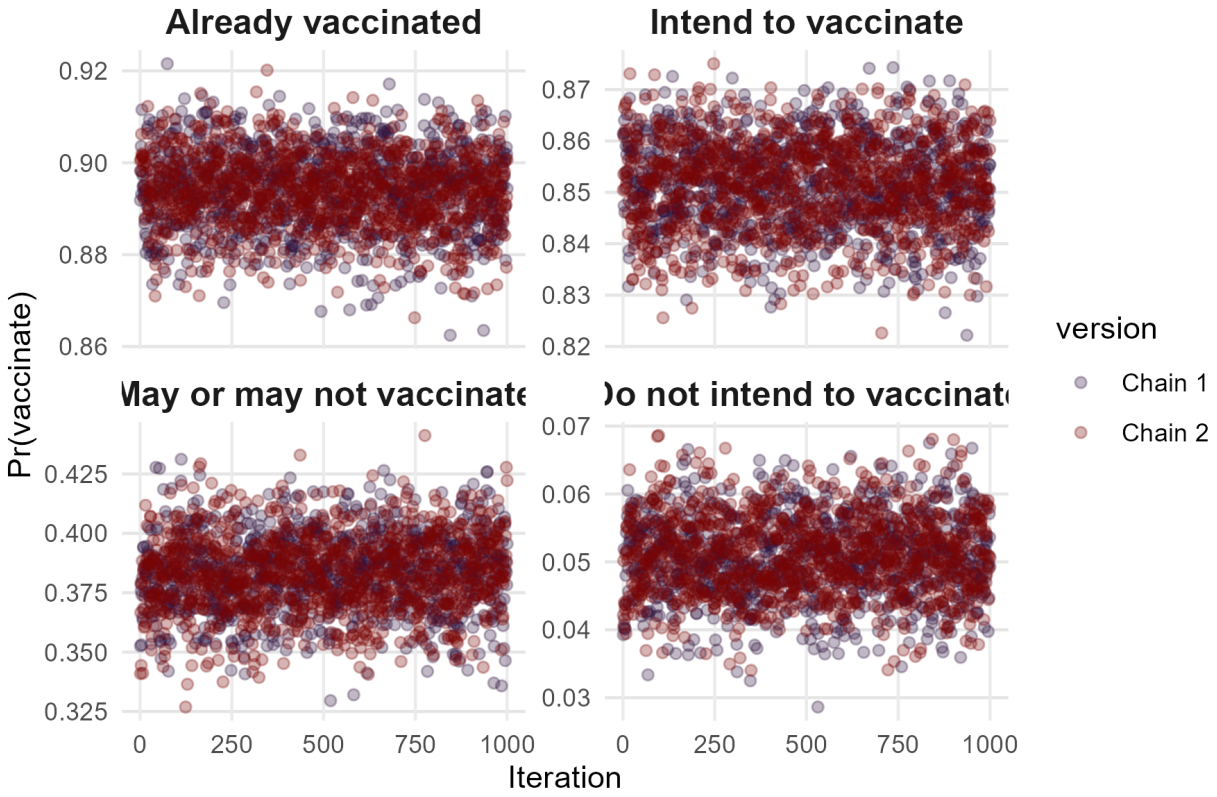
- Subjective side effect rate, vaccination, quantitative: Now we'd like you to put a percent on this likelihood. Imagining you get the flu vaccine in the next two weeks, and thinking about yourself, what do you think is the percent chance that you would experience a severe case of one or more of these 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]

Experienced side effects: In wave 2 of our survey, we asked about experience of side effect symptoms over the prior two weeks. We designed the question to closely parallel our side effect belief elicitation, so that we could compare expected and experienced side effects. Specifically, we asked,

> Sometimes people experience headaches, muscle aches, upset stomach, tiredness, joint pain, fever, or general discomfort.

Thinking about your health **over the last two weeks**, did you experience any of these symptoms **severe enough to disrupt your daily activity?**

Figure C.1: Markov chains are well-mixed

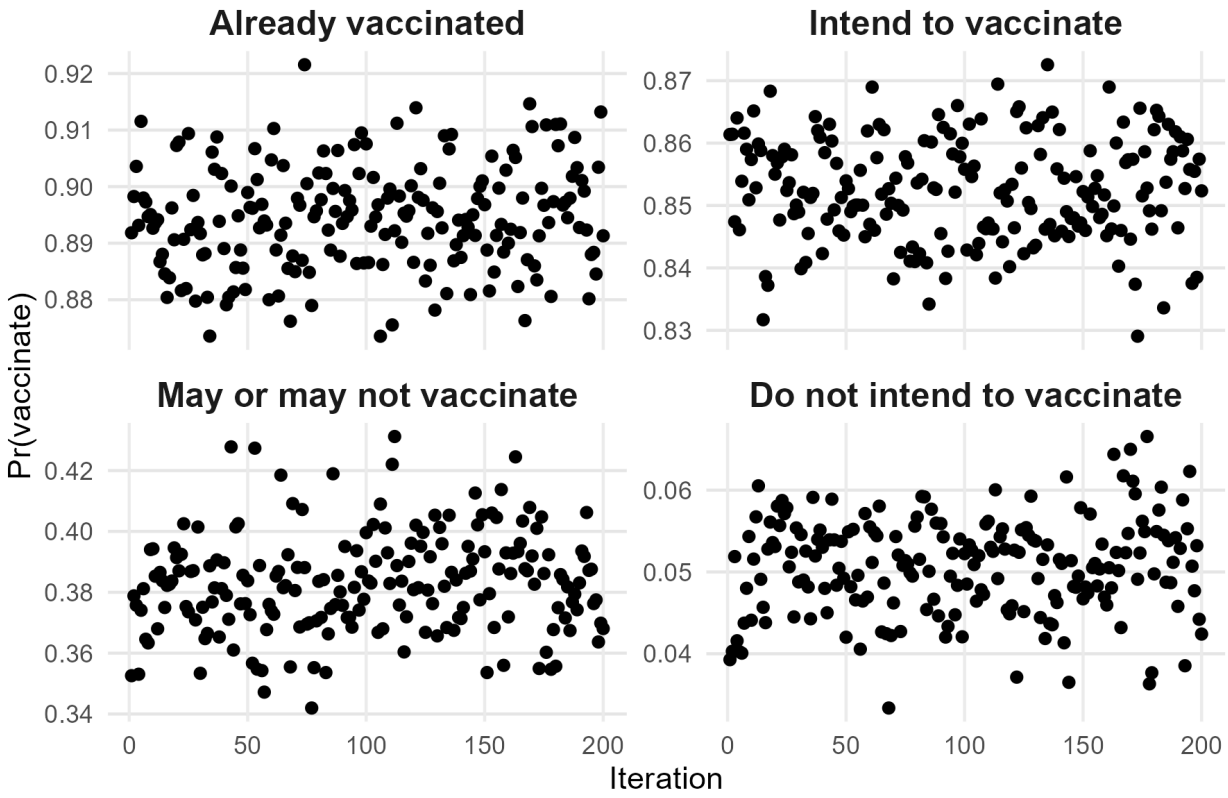


Notes: Figure plots the simulated vaccination rate, at each iterations of the Markov chain, for two separate starting values.

C 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 C.1 shows that our chains exhibit good mixing, and Appendix Figure C.2 shows that they are stationary. To construct Appendix Figure C.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 C.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 C.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 C.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.

D Sensitivity analysis

D.1 Further details on contraindications

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 D.1.[39] 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 D.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 contraindications, 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 D.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

[39] 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.

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.

D.2 Sensitivity analysis of counterfactual simulations

We show robustness of our main model-based results—the simulated effect of changing belief or preferences–in Appendix Tables D.4 - D.3. Appendix Table D.4 reports simulation results when we allow the mean of preference distributions to vary by intention group, but not our other pre-specified; Appendix Table D.4 includes the pre-specified covariates and intention group indicators. The results are quite similar to each other and to our baseline specification.

Appendix Table D.6 shows that our simulation results are not overly sensitive to our handling of beliefs. In our main specification we impute belief at intention group mean (to avoid the influence of our information intervention). In this table, we instead use a LASSO-based imputation, predicting subjective beliefs about side effects and effectiveness as a function of objective beliefs, demographics, hesitancy scale scores, and intention group indicators. The results are similar to our baseline, although the cross-group differences are slightly compressed because the greater within-group variability tends to push group mean vaccination rates towards 50 percent.

Appendix Table D.7 shows that the substantial role we find for side effects and effectiveness is not a consequence of our belief elicitation, which might have focused attention on these characteristics. To show this, we present our main simulation results (averaging over intention groups), but stratifying on the (randomly assigned) experiment order. If our belief elicitation had overly focused attention on side effects and effectiveness, we would expect larger posterior weight on these characteristics among people who went through the belief elicitation, and correspondingly larger changes to eliminating side effects or changing preferences for effectiveness. However, Appendix Table D.7 shows nearly identical simulated probabilities across the experiment ordering.

Table D.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 D.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 D.3: Simulated and counterfactual vaccination rates - included contraindicated

Intention group	Already vaccinated (1)	Intend to (2)	May or may not (3)	Do not intend (4)	Overall (5)
Baseline	0.93	0.9	0.38	0.05	0.53
No SE	0.92	0.87	0.57	0.1	0.58
Normative VSL	0.94	0.85	0.51	0.19	0.59
Both	0.93	0.8	0.75	0.43	0.7

Notes: Table is identical to Table 7 except the estimation and simulation sample is expanded to include participants reporting contraindication to the flu vaccine.

Table D.4: Simulated and counterfactual vaccination rates - preferences vary by intention

Intention group	Already vaccinated (1)	Intend to (2)	May or may not (3)	Do not intend (4)	Overall (5)
Baseline	0.94	0.92	0.33	0.031	0.54
No SE	0.93	0.9	0.57	0.073	0.59
Normative VSL	0.95	0.88	0.49	0.22	0.62
Both	0.94	0.84	0.77	0.5	0.75

Notes: Table is identical to Table 7 except the underlying preferences model allows the mean of preferences to depend on intention group indicators and not the prespecified covariates (indicators for college education or more, age above median, and above median score on vaccine hesitancy scale).

Table D.5: Simulated and counterfactual vaccination rates - richer covariates

Intention group	Already vaccinated (1)	Intend to (2)	May or may not (3)	Do not intend (4)	Overall (5)
Baseline	0.94	0.92	0.33	0.03	0.54
No SE	0.93	0.9	0.57	0.073	0.59
Normative VSL	0.94	0.85	0.49	0.21	0.61
Both	0.93	0.81	0.77	0.47	0.73

Notes: Table is identical to Table 7 except the underlying preferences model allows the mean of preferences to depend on both intention group indicators and the prespecified covariates (indicators for college education or more, age above median, and above median score on vaccine hesitancy scale).

Table D.6: Simulated and counterfactual vaccination rates - alternative belief imputation

Intention group	Already vaccinated (1)	Intend to (2)	May or may not (3)	Do not intend (4)	Overall (5)
Baseline	0.87	0.85	0.43	0.072	0.54
No SE	0.9	0.86	0.6	0.13	0.6
Normative VSL	0.88	0.79	0.54	0.3	0.62
Both	0.9	0.8	0.73	0.46	0.71

Notes: Table is identical to Table 7 except beliefs in the simulation are the LASSO-predicted subjective beliefs given objective beliefs, vaccine intention groups, and demographics.

Table D.7: No survey order effects in counterfactual simulations

	Simulated vaccination rate for...	
	Beliefs first	Conjoint first
	(1)	(2)
Baseline	0.55	0.55
No SE	0.6	0.59
Normative VSL	0.61	0.61
Both	0.71	0.72

Notes: See notes to Table 7. This table is identical to the final column but we stratify the simulated probability based on the (randomly assigned) order of completing the survey (belief elicitation first or vaccine choices first).

E Pre-registration modifications and adjustments

We pre-registered our experiment in the in the AEA RCT registry.[40] 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:

[40] 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 one way ways. We intended to simulate vaccination rates using individual-level beliefs. However, our information provision arm had a large effect on reported beliefs, and these effects appear driven by anchoring rather than true changes, because we see bunching in the belief distribution, but providing this information did not change behavior or preferences. To avoid using these anchored beliefs in our simulations, we impute all beliefs using intention-group averages in the uninformed condition in simulations which require beliefs.