Journal of Risk and Uncertainty (2024) 68:25–49
https://doi.org/10.1007/s11166-023-09424-0

COVID-19 vaccine and risk-taking

Shanike J. Smart[1] · Solomon W. Polachek[1,2]

Accepted: 18 December 2023 / Published online: 23 February 2024

Abstract

We assess whether the COVID-19 vaccine induces COVID-19 risky behavior (e.g., going to bars and restaurants) and thus reduces vaccine efficacy. A key empirical challenge is the endogeneity bias when comparing risk-taking by vaccination status since people choose whether to get vaccinated. To address this bias, we exploit rich survey panel data on individuals followed before and after vaccine availability over fourteen months in an event study fixed effects model with individual, time, sector, and county-by-time fixed effects and inverse propensity weights. We find evidence that vaccinated persons, regardless of the timing of vaccination, increase their risk-taking activities. The evidence is consistent with the "lulling effect". While vaccine availability may reduce the risk of contracting COVID-19, it also contributes to further spread of the virus by incentivizing risk-taking in the short term.

Keywords COVID-19 · Risk-taking · Vaccine · Lulling effect

JEL Classification I1 · I12 · I13 · I18

1 Introduction

The COVID-19 vaccine is estimated to have averted over a million deaths and about ten million hospitalizations in the United States (U.S.), one year following its initial distribution (Schneider et al., 2021).[1] Despite these gains, precautionary measures (e.g., social distancing)

[1] Kniesner and Sullivan (2020) estimated 47 million COVID-19 cases in the U.S., with about one million hospitalizations, resulting in an estimated $2.2 trillion in non-fatal economic losses ($46,000 per case) as of July 27, 2020.

✉ Solomon W. Polachek
polachek@binghamton.edu

Shanike J. Smart
ssmart1@binghamton.edu

1 State University of New York at Binghamton, Binghamton, New York, USA

2 IZA, Bonn, Germany

are still warranted in the short term, when vaccine efficacy is uncertain and coverage is low, and new COVID-19 variants emerge (Lovelace, 2021). But the safety the vaccine affords might incentivize the vaccinated to lower their guard and take more risks, potentially reducing the effectiveness of the vaccine. This behavior could be rational risk-taking behavior, or it might be people overestimate the efficacy of vaccines. This response can be explained by what Viscusi (1984) termed the "lulling effect", which describes situations where people may respond to the adoption of safety measures (e.g., vaccines) with a compensatory increase in risky behavior (also see Peltzman, 1975). Our study aims to investigate the relationship between vaccination status, COVID-19 risky behaviors (e.g., going to bars and restaurants), and the potential impact on the vaccine's effectiveness.

A growing literature examines the effect of health interventions on risk-taking behavior in different contexts. Findings have been mixed, with some studies reporting increased risky behavior following HIV treatment (e.g., Lakdawalla et al., 2006) and Lyme disease vaccine (e.g., Brewer et al., 2007), while others indicate a reduction in risky behavior following the HPV vaccine (e.g., Moghtaderil & Dor, 2021). The effects can be significant. Lakdawalla et al. (2006) find that the introduction of HAART—the recommended HIV treatment regimen—doubles the number of sex partners and increases the risk of HIV infection by at least 44%. Moghtaderil & Dor (2021) find that receiving the HPV vaccine increases the probability of having a Pap test, which is a diagnostic screening test to detect HPV and non-HPV-related cancers, by 22 percentage points (pp). This behavior can be due to shifts in risk perceptions. Viscusi and Cavallo (1994) found that the child-resistant features of cigarette lighters make precaution-taking less likely mainly because users think lighters are safer. The mixed evidence suggests that the relationship between health interventions and risky behavior depends on factors such as the nature of the intervention, the context being studied, and the behaviors being measured. Thus, the extent to which vaccines affect risk-taking behavior is an empirical question that requires further investigation.

The mixed evidence at least partly reflects the challenges in identifying the effect of vaccinations on risk-taking. A key issue is the endogeneity bias arising from people choosing whether to get vaccinated. The problem is unvaccinated individuals cannot serve as a suitable control group when vaccination decisions are based on factors such as vulnerability and fear, which are also usually related to one's risk-taking behavior. For example, individuals who are more vulnerable to illness or fearful of the virus may be more likely to choose vaccination, and such individuals may also be less inclined to engage in risky behavior. Therefore, comparing mean levels of risk-taking behavior between vaccinated and unvaccinated individuals may result in biased estimates, as other factors beyond vaccination status could explain any observed differences.

We make several attempts to address this endogeneity bias. First, we specify an event study fixed effects model that exploits survey panel data on individuals followed at quarterly waves before and after vaccine availability over fourteen months. This model includes individual and wave fixed effects to control for time-invariant differences in individual characteristics and broad secular trends that might impact vaccine intent and risky behavior, respectively. We also include county-by-wave fixed effects to account for the potential influence of county-specific trends, such as vaccine access and community-wide beliefs or preferences towards vaccines, which likely change over time with new information. Second, we control for difficult-to-measure time-varying individual characteristics that may influence both vaccine intent and risk-taking behavior such as willingness, vulnerability, fear, and COVID-19 exposure. Third, we use an inverse propensity weighting (IPW) technique to account for these and other covariates, allowing for potentially nonlinear effects on the likelihood of being vaccinated. IPW facilitates a comparison between vaccinated and unvaccinated persons who have similar vaccine propensities by up-weighting the unvaccinated who have a high probability of choosing to get vaccinated.[2] Finally, we check for further endogeneity issues by conducting a placebo test, using the flu vaccine as a placebo treatment for COVID-19. If lowered risk of contracting COVID-19 is the means by which the COVID-19 vaccine affects risky behavior, then one would be suspicious if the flu vaccine also affects it, given the flu vaccine does not protect against COVID-19. Using our main model, we show that the flu vaccine has a small and statistically insignificant association with risk-taking. This falsification test suggests that endogeneity concerns such as reverse causality or unobservables do not seem to be playing a role in driving our results.

We find that once vaccinated, individuals exhibit varying degrees of increased participation in some COVID-19 risky behaviors, regardless of the timing of vaccination. Specifically, we find positive and significant overall average effects of vaccination on dining indoors, attending religious services in person, and shopping, ranging from 3–11 pp. Our results also reveal heterogeneity based on the timing of vaccination with those vaccinated early displaying increased participation in more risky activities than those vaccinated later, and they do so at different times (i.e., one period post-vaccination as opposed to at vaccination) for activities common to both groups. However, at the time of vaccination, those vaccinated later (typically the general public without early vaccine access) exhibited a greater significant increase in the magnitude of the risk-taking effect (7–17 pp) than those vaccinated earlier (4–8 pp) (typically vulnerable or essential workers with early vaccine access). The evidence presented is consistent with the "lulling effect". While vaccine availability may reduce the risk of contracting COVID-19, it also contributes to the further spread of the virus by incentivizing risk-taking in the short term.

To the best of our knowledge, a study by Agrawal et al. (2022) is the only other study to examine the relationship between the COVID-19 vaccine and risk-taking. Using repeated cross-sectional data, Agrawal et al., implement an RD using discontinuity in vaccine availability at age 65; and an IV approach using variations in state-level age-group eligibility policies to instrument for vaccine take-up. They find that vaccination does not affect mask-wearing and avoiding crowds/restaurants, but obtain mixed evidence for handwashing. Both the RD and IV models estimate a local average treatment effect, as such they provide estimates for the age-65 cohort and for those influenced by vaccine eligibility policies, respectively.

Our study contributes to the literature in several respects. First, we are among the first to empirically analyze the impact of the COVID-19 vaccine on COVID-19 risk behavior. This analysis contributes to the existing literature on how vaccines

[2] See Callaway and Santana (2021) for recent discussion of the IPW approach.

may have different effects in practice than those observed in controlled settings such as clinical trials. In clinical trials, participants are less likely to engage in compensatory behavior since they do not know if they received the vaccine or a placebo and have no evidence of its effectiveness. Hence, even well-designed clinical trials may overstate the vaccine's efficacy in practice. Second, eleven kinds of COVID-19 risky behavior (e.g., visiting restaurants, friends and family, and going shopping) are used to define risk-taking. In addition to analyzing these eleven variables independently, we create composite measures that consider the overall level of risk-taking among individuals and attempt to account for the actual, as opposed to the presumed, risk of contracting COVID-19. This extends current evidence that is limited to the four behaviors examined by Agrawal et al. (2022) and might not fully represent the risky behavior of the population. Third, we are the first to use panel data to examine the effect of the COVID-19 vaccine on risk-taking. Panel data allows us to control for endogeneity concerns arising from pre-vaccine differences between the vaccinated and unvaccinated in several ways. In addition to controlling for standard demographic variables (e.g., gender and income), the richness of the panel data allows us to control for difficult-to-measure individual characteristics that typically drive vaccine intent (e.g., fear of virus). Fourth, our approach produces results that are generalizable to the U.S. adult population. Finally, the study provides the first estimates of how the timing of the vaccine decision affects the vaccine's impact on risk-taking behavior.

The rest of our paper proceeds as follows. In Section 2, we describe the data and define key variables used in our model. Section 3 describes the methodology used to estimate the effect of vaccination on risk-taking. We present and discuss the results in Section 4. Finally, we conclude in Section 5.

2 Data

We use data from the Socio-Economic Impacts of COVID-19 Survey (SEICS), a nationally representative survey of adults 18 and older residing in various counties across different states. The survey was launched by the Social Policy Institute (SPI) at Washington University in St. Louis and was conducted online using Qualtrics. It consists of five waves administered at quarterly intervals between April 2020 and June 2021. Participants were compensated for completing the survey.[3] The survey recruits new respondents in each wave while also allowing for re-contacts of prior-wave respondents. This allows us to track changes in respondents' behaviors and outcomes over a period of fourteen months after COVID-19 was declared a pandemic.

[3] Because the survey was only offered online, everyone in the sample had at least basic internet access. Respondents were also provided with monetary incentives to complete the survey which could introduce bias by only attracting individuals motivated by a reward. Hence, even though the SEICS original sample closely approximates the general population of adults in the U.S., the analysis includes population weights provided in the survey generated from the ACS to further ensure that SEICS sample closely approximates the general population.

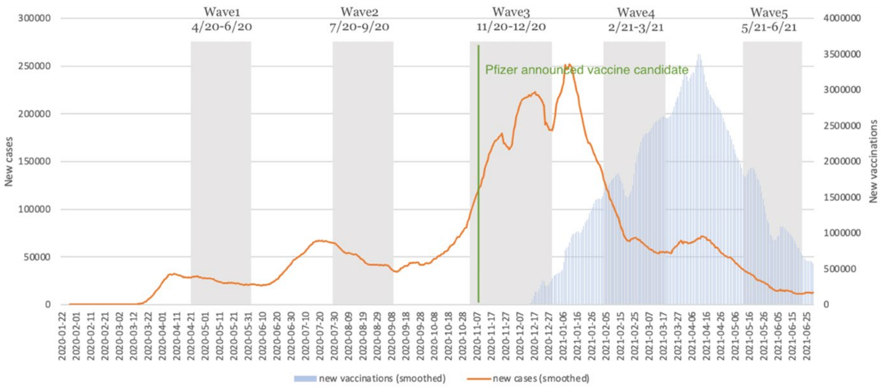

Fig. 1 Survey time periods. Source: Constructed by the author. Data for COVID-19 cases and vaccination retrieved from Global Change Lab, Oxford University

Questions about respondents' risky behavior (the main outcome of interest) were introduced beginning in wave 2 to approximately 5,000 respondents per wave.[4]

We exploit the panel nature of the survey by using only a subsample of individuals who participated in at least one survey wave before and after the vaccine was made available for distribution. Figure 1 displays a timeline of the survey waves and the corresponding dates of vaccine approval and availability. As shown in Fig. 1, mass vaccinations began on December 14, 2020, towards the end of wave 3 of the survey. Thus, our final sample included 3600 individuals who took part in either waves 2 and/or 3 and returned to participate in either waves 4 and/or 5.

2.1 Representativeness

As with any survey, there may be ways in which the sample systematically differs from the full adult population in the U.S. Systematic differences might also persist since we only use a subset of the survey respondents who were surveyed in early waves of the survey and returned at a later wave. To check that the chosen subsample does not suffer from attrition and is an overall representative sample, we compare the characteristics of the weighted SEICS baseline sample with those of the population of adults in the U.S., as measured through the Census Bureau's 2020 American Community Survey (ACS) (5-year-average estimates).[5] The results of this comparison are presented in Table 4.

As Table 4 shows, our subsample reasonably approximates the population of adults living in the U.S. on key demographic dimensions such as gender, race, household income, and age. The racial/ethnic composition of the sample closely

[4] More details about the survey methodology available at socialpolicyinstitute@wustl.edu.

[5] The baseline sample refers to respondents that were a part of wave 3 (82% of the sample) as this is the wave right before the vaccine became available and the wave before anyone ever got vaccinated.

matches that of the population, as does the distribution of income. On average, the sample of respondents was roughly five years older than adult ACS respondents. In terms of education and marital status, 8% of respondents were more likely to hold a bachelor's degree or higher level of education, and 8% were more likely to be married compared to the population of adults in the U.S. The gender distribution of the SEICS sample was also similar to the ACS adult population but with 2% more females. Given that the sample only exhibits minor deviations from the ACS and that the key demographic features are closely linked to other population characteristics (e.g., debt level, geography, political preference), the sample can be considered reasonably representative of the U.S. population on a wide range of dimensions.

2.2 COVID-19 risk-taking measures

Our measures of risk-taking come from the responses to SEICS survey questions about the degree to which respondents' social distance before and after being vaccinated. These social distancing activities such as dining out, were normally safe but had temporarily become risky due to the COVID-19 pandemic. Starting in wave 2, respondents were asked to report the number of times within the past three months—ranging from '0 = never' to '4 = more than 10 times'—they left their home to engage in the following "risky" activities: (i) shopping for food and essentials, (ii) visiting friends and family, (iii) going to work in person, (iv) using public transit, (v) dining indoors at restaurants, (vi) dining outdoors at restaurants, (vii) attending a religious service, (viii) going to a bar, (ix) traveling by airplane, (x) attending a political rally speech or campaign event, and (xi) attending an organized protest, march or demonstration of any kind. We recategorize responses for each of these eleven variables into an indicator variable: 0 if the respondent never left their home to engage in the activity; and 1 if the activity is done at least one or two times. We first consider each of these indicator variables as separate measures of risk-taking, representing whether the respondents engaged in the respective risky activity or not.

Social distancing is a multi-dimensional construct given it entails several activities, from going to a bar to visiting friends. Considering each of these acts as a separate measure of risk-taking behavior does not account for the inherent variability in individual behavior. For example, if we solely consider visits to a bar, we might erroneously conclude that someone is not engaging in risky behavior if they avoid bars, despite their participation in other risky activities such as increased shopping visits and dining indoors at restaurants. Our analysis therefore includes ways of aggregating these activities to get a composite measure of risk-taking that combines the eleven risky activities. We consider four different aggregation approaches.

The first is constructed as a count of all eleven activities. We sum across the eleven indicator variables for each respondent to create a count (ranging from 0 to 11) increasing in risk-taking, which represents the number of these risky activities done by any respondent (mean = 3.8, SD = 2.5).[6] While a count of risky activities

[6] A Cronbach alpha check of internal consistency equals 0.80 for the count of risky activities showing that this variable is internally reliable. Factor analysis revealed one underlying factor with an eigenvalue greater than 1. Taken together, this statistic is considered reliable and internally consistent.

provides insights into the extent to which individuals engage in multiple COVID-19 risky activities, it is important to acknowledge the inherent limitations of simply summing the number of individual activities. This method assumes each activity is equally important in increasing the spread of COVID-19, which may not be the case in this setting. For instance, the impact of an increase in indoor dining on COVID-19 transmission is likely higher than a comparable increase in outdoor dining. Additionally, the nature and frequency of events vary, with some (e.g., shopping) being commonplace and less amenable to social distancing than others (e.g., protests and campaign events) (see Dave et al., 2021).

To account for the relative importance of each risky activity, we develop a second risk measure using weights derived from factor loadings in a principal component analysis of the eleven activities. Computed by summing the products of each variable's value and its corresponding factor loading, this method accounts for the significance of each activity in relation to an assumed latent construct, representing COVID-19 risk. While it is reasonable to assume the primary latent factor adequately represents risk-taking in our observed setting given the nature of the activities and the survey's explicit design for gauging social distancing during the COVID-19 pandemic, we acknowledge its limitations. Thus, we explore alternative approaches to assign weights based on the actual, rather than presumed, risk of contracting COVID-19.

It is challenging to obtain precise calibrations of the actual risks of non-compliance with different parts of the social distancing guidelines, whether in terms of life expectancy or even infection probability (Heffetz & Rabin, 2023). In this context, we aim to provide suggestive correlations, which we believe are useful. For our third aggregate measure, we account for actual COVID-19 risk by estimating a probit regression to model the probability of a household member testing positive for COVID-19 based on each of the eleven activities as predictors.[7] The resulting coefficients, from which we obtain the marginal probabilities of a household member getting COVID-19, serve as weights for combining the eleven activities into a third aggregate measure by summing the products of each activity's value and its respective weight. To ensure positive predicted values, aligning with the expectation that risky activities contribute positively to COVID-19 risk in our context, we apply the exponential function to constrain the coefficients. For robustness, we construct a fourth measure using coefficients as weights done previously, but now from an Ordinary Least Squares (OLS) model. To address potential endogeneity between household COVID-19 positive status and engagement in risky activities, we model the estimation of COVID-19 exposure (household testing positive or not) before vaccine availability, based on lagged values of the eleven risky activities.[8] This approach

[7] We use the household measure as a proxy of actual COVID-19 risk since a direct measure of individual COVID-19 status is unavailable.

[8] More formally, we estimate the following OLS and probit models: $COVIDPOS_{it} = \beta_0 + \boldsymbol{R}'_{i,t-1}\delta + v_{it}$; and $\Pr(COVIDPOS_{it} = 1) = \beta_0 + \Phi\left(\boldsymbol{R}'_{i,t-1}\delta\right)$, respectively. $COVIDPOS$ is an indicator variable where 1 = someone in the household tests positive for COVID-19. $\boldsymbol{R}'_{i,t-1}$ is a vector of the lag of the eleven COVID-19 risky activities. $\Phi()$ is a probit function. The details of these regression estimations and output are available upon request.

mitigates the possibility of reverse causality, in which contracting COVID-19 affects subsequent behavior, because in our main analysis COVID-19 follows rather than precedes the risky behavior.

All demographic characteristics necessary to categorize respondents based on vaccination status are taken directly from the responses provided in the survey. Table 1 provides summary statistics of these and all other model variables, which we discuss next.

2.3 Vaccination decision covariates

In this section, we summarize the literature on the vaccine hesitancy determinants drawing mainly from Aw et al.'s (2021) meta-analysis to better understand possible differences by vaccine status, which might act as confounders in our model. We first discuss the three key categories they identified—contextual, group/individual, and vaccine-specific factors—followed by how we account for these factors (listed in Table 1).

2.3.1 Contextual factors

Of the contextual factors most examined in the literature, socio-demographics, namely, young age, female, non-white, and lower education, were most commonly linked to higher vaccine hesitancy (Aw et al., 2021). Aw et al. (2021) and Dhanani and Franz's (2022) meta-analysis show systematic differences in vaccine intent among different socio-demographic groups due in part to vulnerability to severe illness or death from COVID-19 based on factors such as occupation, age, race/ethnicity, and health status.[9] Political affiliation and inclinations also play a significant role in vaccine hesitancy (Aw et al., 2021; Dhanani & Franz, 2022) and predict reported policy preferences, mask use, and social distancing (e.g., Barrios & Hochberg, 2020; Bruine de Bruin et al, 2020). Studies find that Republicans report significantly higher vaccine hesitancy than Democrats (Dhanani & Franz, 2022) and that Democrats are more likely than Republicans to take the threat of the virus seriously and support efforts to control it (Allcott et al., 2020).

2.3.2 Group/individual factors

The primary group/individual factors influencing vaccine intent encompass beliefs/attitudes about health and prevention (e.g., fear of contracting the virus or beliefs about disease severity), past experiences with vaccinations, and trust and experience with the health system and providers (Aw et al., 2021). Akesson et al. (2022) underscores the significance of beliefs, revealing through their online experiment, that individuals dramatically overestimate the dangerousness and infectiousness of

[9] More than 81% of COVID-19 deaths occur in people over age 65 (See Center for Disease Control and Prevention, 2023). Those in close contact with these persons also tend to have a greater incentive to vaccinate and take precautions (Dabla-Norris et al., 2021).

Table 1 Weighted average of baseline characteristics

	Overall	Vaccination Status			
		Early	Late	Never	
Controls:					
18 to 29 years	0.08	0.05	0.05	0.14	***
30 to 44 years	0.23	0.18	0.20	0.29	***
45 to 64 years	0.40	0.27	0.46	0.41	***
65 and older	0.29	0.50	0.29	0.15	***
Female	0.53	0.47	0.49	0.61	***
Income	95,238	108,778	97,932	82,174	***
High school or less	0.29	0.20	0.28	0.36	***
College but no Bachelor	0.30	0.28	0.29	0.34	
Bachelor or higher	0.41	0.52	0.43	0.31	***
White	0.62	0.62	0.62	0.63	
Married	0.58	0.60	0.62	0.52	***
Have kid(s) < 6 yrs. in HH	0.06	0.03	0.05	0.08	***
Enrolled in school	0.11	0.14	0.09	0.12	*
Elderly (65+) in household	0.40	0.60	0.39	0.27	***
Any chronic illness in HH	0.39	0.50	0.41	0.29	***
Have disability	0.17	0.16	0.16	0.18	
Work in person	0.35	0.35	0.34	0.36	
HH member test COVID positive	0.04	0.06	0.03	0.04	
Fear of COVID	0.58	0.70	0.63	0.44	***
Perceived chance of COVID illness/ death (0–100)	33.3	41.5	33.5	27.1	***
Friend/family died from COVID	0.09	0.13	0.08	0.06	***
Liberal	0.29	0.34	0.31	0.24	***
Moderate	0.35	0.33	0.35	0.36	
Conservative	0.36	0.33	0.33	0.41	**
Republican	0.32	0.29	0.28	0.40	***
Democrat	0.39	0.47	0.44	0.29	***
Independent	0.25	0.22	0.25	0.28	
Other Political Party	0.03	0.02	0.03	0.04	
Willing to get vaccinated	0.64	0.84	0.72	0.40	***
Observations	2722	720	1100	902	

Note: All values reported in the table represent average proportions unless otherwise stated. The baseline is wave 3 which is immediately prior to vaccine availability. *p < 0.10; **p < 0.05; ***p < 0.01 indicates a significant difference between the three groups.

COVID-19 relative to expert opinion. Providing expert information partially corrects these beliefs. Within the broader literature, heightened vaccine hesitancy is commonly associated with a low perceived risk of contracting COVID-19, absence of chronic medical conditions, and the belief that COVID-19 is not severe (Aw et al., 2021). Adding to the complexity, one's risk preference is also a significant factor. Studies illustrate a negative correlation between risk-taking and COVID-19 vaccine uptake, suggesting that individuals with a higher inclination towards precautionary measures are more likely to get vaccinated (Dabla-Norris et al., 2021; Latkin et al., 2021).

2.3.3 Vaccine-specific factors

Vaccine-specific considerations significantly shape COVID-19 vaccine intent, with a key focus on perceptions regarding the safety and efficacy of the COVID-19 vaccine (Aw et al., 2021). These concerns mirror issues observed with established vaccines, such as the flu vaccine, where safety and efficacy remain central considerations (Dibonaventura & Chapman, 2008). Additionally, structural barriers, including the location for vaccination and the time required for transport, further contribute to vaccine-specific factors influencing intent (Aw et al., 2021).

2.3.4 Proxies for vaccine decision influences

We account for contextual determinants/socio-demographic differences in our model by including several control variables derived from SEICS data collected on each respondent. These variables include respondents' age, gender, income, race, education, marital status, sector employed, and child (<6 years) living in the household. We also observe whether an elderly (65+) lives in the respondent's household, whether they or anyone in the household have one or more health conditions,[10] whether anyone in the household ever tested positive for the coronavirus, and whether they are required to work in person regardless of the sector in which they work. We also include respondents' political beliefs, liberal, moderate, or conservative as well as their reported political party affiliation as control variables. The survey data also includes self-reported measures of usually difficult-to-measure time-varying characteristics, that we use to account for group/individual influences. These are the self-reported fear of COVID-19 (reported on a scale that ranges from '0 = no fear' to '100 = very afraid'), the probability of becoming ill and requiring hospitalization or dying (reported on a scale from '0 = no probability' to '100 = certainty'), and whether the respondent has a close friend or relative die from COVID-19.[11]

To account for unobserved individual differences, we first use SEICS data on individual's self-reported measure of willingness to receive the vaccine as a proxy for vaccine intent. This variable serves as a catchall for other individual drivers

[10] The chronic medical illnesses reported are chronic lung diseases, moderate to severe asthma, serious heart conditions, being immunocompromised, diabetes, kidney disease (undergoing dialysis), liver disease, hypertension, severe obesity (body mass index of 30 or higher) and sickle cell disease.

[11] We recoded self-reported fear as an indicator variable equal to 1 for reported fear exceeding the mean value and 0 otherwise.

of vaccine intent that are not directly observed in our data. Secondly, as we show later, we exploit the panel nature of our data to control for other unobserved differences, such as personality traits, through individual and wave fixed effects. We also account for vaccine-specific influences such as accessibility through county-by-wave fixed effects.[12]

2.3.5 Balance of covariates

We check for balance between vaccinated and unvaccinated groups using all observed covariates that might explain the decision to vaccinate. For each covariate, we compare the average value for the early and late vaccinated group to the average value among the never vaccinated group using an F-test. Covariates are considered balanced if the resulting test statistic is not significantly different from zero at the 10% level of significance or lower. The result of the balance tests is presented in Table 1. We find a significant difference for the majority of the covariates (see Table 1). Therefore, simply comparing the mean risk-taking behavior by vaccine status could result in biased estimates, as any observed differences could be due to factors other than vaccination. In the next section, we describe how our methodology addresses this problem.

3 Empirical methodology

Table 1 suggests that simple comparisons of risk-taking by vaccination status can be misleading, given the number of observable confounders. As a result, we specify a model that attempts to address the endogeneity of one's COVID-19 vaccine decision.

In our setting, persons receive their initial vaccination at different points in time: in wave 4—which we refer to as early vaccination, and wave 5—which we refer to as late vaccination. In this staggered treatment set-up, recent literature warns that the most popular model used—two-way fixed effects models—produces estimates that are not reliable in the presence of effect heterogeneity, and that they potentially even have the wrong sign.[13] To avoid this, we estimate separate event study fixed-effects models for each vaccinated group. Formally, we estimate the effect of vaccination for the early and late vaccinated by employing the following event-study fixed effects model:

$$Y_{ict} = \sum\nolimits_{e\neq-1} \beta_e \cdot 1(E = e) + \eta_i + \lambda_t + \delta_{ct} + v_{ict}, \tag{1}$$

[12] One other factor that might be related to vaccine intent is religiosity. While the literature reviewed does not highlight religiosity as a key determinant in vaccine intent, it is usually referred to tangentially in partisan beliefs (e.g., Albrecht, 2022). While we do not observe one's religiosity directly, we observe several variables that the literature finds are highly correlated with religiosity such as willingness to get vaccinated, political affiliation, as well as political party, and we later adjust for county-by-wave fixed effects which proxies any community effect of religious influences.

[13] See Goodman-Bacon (2021), Callaway and Sant'Anna (2021), De Chaisemartin and D'Haultfoeuille (2020), and Borusyak et al. (forthcoming) for a recent discussion on the failure of two-way fixed effects regressions for time-varying treatment effects.

where Eq. (1) is estimated as separate regressions for early and late vaccinated. The control group consists of individuals that did not yet get the vaccine by wave t. Y_{ict} is an indicator variable representing whether individual i in county c at wave t participates in a given type of risky behavior. For the composite measures of risk-taking, we use Eq. (1) to model the expected value of a Poisson random variable given the positively skewed nature of all measures and the fact that one of them follows a count distribution. E is the length of time before and after vaccination. Hence, if vaccinated early, E takes on values -2, -1, 0, and +1. If vaccinated late, E takes on values -3, -2, -1, and 0. The period immediately prior to vaccination (i.e., -1), is the reference exposure period for both vaccinated groups and is therefore excluded from the analysis. β_e is the key parameter that we are interested in estimating. It represents the change in the average level of risk-taking behavior among the vaccinated group relative to their pre-vaccination period, compared to what would have occurred if they had not been vaccinated. In addition to the event study estimates, we report the overall average of the estimated effects for the early and the late vaccinated, weighted by the share of individuals in each group.

We account for broad secular trends and time-invariant differences in individual characteristics (observable and unobservable) that might affect both vaccine intent and risky behavior, using wave and individual fixed effects (λ_t and η_i), respectively.[14] We also account for county-specific trends, such as vaccine access, infection and hospitalization rates and county-wide popular beliefs or preferences towards vaccines, which likely change over time due to new information and resources. We do so by including county-by-wave fixed effects (δ_{ct}).

It is also important that we control for individual time-varying observables which likely drive vaccine intent and risk-taking behavior. We adjust Eq. (1) for the observed covariates using inverse propensity weights. We estimate the propensity weight as follows:

$$\omega_i = \sqrt{D_i + \frac{\left(1 - D_i\right)\widehat{p(x_i)}}{1 - \widehat{p(x_i)}}}, \tag{2}$$

where $D_i = 1$ when individual i is vaccinated (and 0 otherwise), and $p(x_i)$ is the estimated propensity score for individual i. The propensity score is computed by estimating a separate logit model, which predicts the probability of selecting vaccination as a function of all covariates listed in Table 1 along with dummies for an individual's sector of employment. Because there are two separate vaccination dates, there is a unique propensity score for every group. The analysis uses covariates from only the wave prior to vaccination in the propensity score calculation. Using a logit model to predict the propensities allows the effects of the covariates to have a non-linear impact on the probability of being vaccinated as opposed to just a linear one if we were to simply include these covariates in the model directly.

[14] One such difference that would impact vaccine intent and risky behavior is ease of access to a vaccine site.

IPW assigns higher weights to unvaccinated individuals that have higher odds of being vaccinated. This reduces the differences between the vaccinated and unvaccinated individuals in terms of their propensity to take the vaccine. To incorporate the uncertainty from estimating the propensity score in a separate regression, we use bootstrapped standard errors clustered at the individual level based on 1,000 replications. We also perform two standard checks on the IPW. The first assesses the common support requirement that there should be both vaccinated and unvaccinated individuals who have a similar propensity to get vaccinated. The second, checks for covariate balance, which determines whether the IPW adequately adjusts for differences in covariates between the vaccinated and unvaccinated, allowing us to compare individuals of similar covariate distributions. We discuss the results of these checks in Section 4.

The validity of our approach rests on a conditional parallel trend assumption. This implies that average risk-taking behaviors for the vaccinated and unvaccinated groups would have followed parallel paths in the absence of vaccination, conditional on adjusting for covariates.[15] So far, we are assuming that applying inverse propensity weights to Eq. (1) and including individual, wave, county, and county-by-wave fixed effects accounts for any pre-vaccine differences that could cause a violation of the parallel trend assumption. In Section 4, we test empirically whether we have sufficiently controlled for differences between the vaccinated and unvaccinated prior to vaccine availability by estimating analogous vaccine effects for the pre-vaccine periods. If the vaccinated and unvaccinated groups follow parallel paths in the absence of the vaccination, the pre-vaccination estimates should be insignificant and close to zero. These results, reported as a part of Tables 2 and 3, are discussed with our main results in Section 4. We also consider the potential role of other endogeneity issues, such as reverse causality and unobservables on our findings by implementing a placebo test in Section 4.2, using the flu vaccine as a placebo treatment for COVID-19. We show using the outcomes from the placebo that these issues do not appear to be driving our results.

4 Empirical results

4.1 Diagnostic results

In this section we conduct checks on the IPW estimation and present evidence of the parallel trends. These findings set the stage to measure how vaccination affects risk-taking behavior of individuals.

[15] Note that it would be unreasonable for parallel trends to hold without conditioning on covariates. One reason is the rapid and progressive policy changes implemented by states that might affect risky behaviors and vaccine intent. These changes include limits on social gatherings, adding states to travel quarantine lists, mandating face masks and encouraging residents to stay home. For example, in late March and early April 2020, a majority of states issued orders directing residents to stay at home and for schools and nonessential businesses to close. By late April (and extending up to the beginning of June) states then began reopening.

Table 2 Estimated effect of vaccine status on risk taking measured by type of social-distancing activity done

	Visit friends/ family	Dining indoor	Dining outdoor	Attend religious service	Visit bar	Travel by airplane	Attend campaign	Attend protest	Use public transit	In-person work	Shopping
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Overall average effect	0.038	0.106***	0.002	0.108***	0.046*	-0.014	0.027*	0.022	0.028	0.030	0.045*
	(0.029)	(0.025)	(0.029)	(0.035)	(0.026)	(0.024)	(0.015)	(0.014)	(0.018)	(0.025)	(0.025)
Panel A: Early vaccinated											
Exposure											
1 period before	0.005	0.024	-0.019	-0.038	-0.059**	-0.009	-0.018	-0.043*	-0.021	-0.010	-0.012
	(0.039)	(0.035)	(0.044)	(0.031)	(0.025)	(0.027)	(0.023)	(0.026)	(0.025)	(0.027)	(0.024)
Period of vaccination	0.020	0.049	-0.015	0.037	0.014	-0.022	0.038*	0.027	0.023	0.079***	0.017
	(0.034)	(0.032)	(0.046)	(0.045)	(0.038)	(0.030)	(0.021)	(0.019)	(0.024)	(0.030)	(0.025)
1 period later	-0.019	0.200**	-0.047	0.150*	0.089	0.064	-0.002	0.020	0.092	0.024	0.049
	(0.061)	(0.079)	(0.075)	(0.080)	(0.069)	(0.074)	(0.039)	(0.036)	(0.059)	(0.044)	(0.061)
Observations	7741	7732	7727	7740	7734	7725	7726	7737	7731	7729	7747
Panel B: Late vaccinated											
Exposure											
2 periods before	-0.110**	-0.010	-0.041	0.024	0.044	-0.021	0.033	-0.012	-0.005	0.003	-0.039

Table 2 (continued)

	Visit friends/ family	Dining indoor	Dining outdoor	Attend religious service	Visit bar	Travel by airplane	Attend campaign	Attend protest	Use public transit	In-person work	Shopping
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	(0.051)	(0.043)	(0.053)	(0.045)	(0.040)	(0.037)	(0.022)	(0.025)	(0.028)	(0.034)	(0.029)
1 period before	-0.092**	-0.010	-0.032	0.015	0.046	0.012	0.009	-0.007	-0.006	-0.008	0.009
	(0.041)	(0.038)	(0.053)	(0.047)	(0.035)	(0.033)	(0.018)	(0.019)	(0.021)	(0.027)	(0.030)
Period of vaccination	0.018	0.156***	0.001	0.168***	0.070*	-0.021	0.020	-0.004	0.013	-0.024	0.062
	(0.048)	(0.046)	(0.045)	(0.060)	(0.037)	(0.037)	(0.026)	(0.022)	(0.029)	(0.044)	(0.046)
Observations	7099	7093	7080	7089	7095	7086	7081	7097	7085	7093	7098

Note: The table reports estimated coefficients and standard errors derived by running separate fixed effects linear regressions with IPW and county-by-wave fixed effects, for each vaccinated group and each type of COVID-19 risky activity. The dependent variable is each risky measure shown in columns (1) to (11), defined as an indicator variable where 1 = individual participated in the respective behavior at least one or two times and 0 = otherwise. Controls used for the IPW adjustment are those listed in Table 1 as well as dummies for an individual's sector of employment. Results are all weighted with ACS population weights provided as a part of the survey. Bootstrapped standard errors clustered at the individual level are in parentheses. The number of bootstrap replications is 1,000. *, **, *** indicate statistical significance at 10%, 5% and 1% levels.

Table 3 Estimated effect of vaccine status on risky activities

	Count of COVID Risky Activities		COVID Risk	
	(1)	(2)	(3)	(4)
	Sum	PCF	OLS	Probit
Overall average effect	0.120***	0.158***	0.166**	0.162**
	(0.036)	(0.053)	(0.079)	(0.070)
Panel A: Early vaccinated				
Exposure				
1 period before	-0.037	-0.002	0.019	0.050
	(0.048)	(0.078)	(0.119)	(0.101)
Period of vaccination	0.074**	0.156**	0.134	0.130
	(0.035)	(0.074)	(0.107)	(0.091)
1 period later	0.173**	0.417**	0.513*	0.481**
	(0.077)	(0.172)	(0.264)	(0.235)
Observations	7758	7383	7383	7383
Panel B: Late vaccinated				
Exposure				
2 periods before	-0.023	0.004	0.024	0.003
	(0.056)	(0.082)	(0.145)	(0.123)
1 period before	-0.011	-0.023	0.096	0.059
	(0.044)	(0.070)	(0.103)	(0.091)
Period of vaccination	0.137**	0.108	0.164	0.166
	(0.067)	(0.083)	(0.111)	(0.104)
Observations	7110	6810	6810	6810

Note: Estimated coefficients and standard errors obtained from individual fixed effects Poisson regressions as specified in Eq. (1). Dependent variables: number of COVID-19 risky activities (column 1); risk-weighted average of COVID-19 risky activities (columns 2–4), where the weights are generated using principal component factor analysis, OLS, and probit models, respectively. Each regression adjusts for inverse propensity weights and individual, wave, and county-by-wave fixed effects. Controls used for IPW are those listed in Table 1 as well as dummies for an individual's sector of employment. Results are all weighted with ACS population weights provided as a part of the survey. Bootstrapped standard errors clustered at the individual level are in parentheses. The number of bootstrap replications is 1,000. *, **, *** indicate statistical significance at 10%, 5% and 1% levels.

4.1.1 Common support and covariate balance

To generate the weights for the IPW we described in Section 3, we estimate the probability of selecting to be vaccinated in the form of a propensity score. Once a propensity score has been calculated for each observation, one must ensure that there is overlap (typically referred to as common support) in the range of propensity scores across vaccinated and unvaccinated groups. This ensures the availability of both vaccinated and unvaccinated individuals that have a similar propensity to get vaccinated. To detect the extent of overlap, we follow Imbens' (2004) recommendation and plot the density of estimated probabilities for selecting vaccination separately

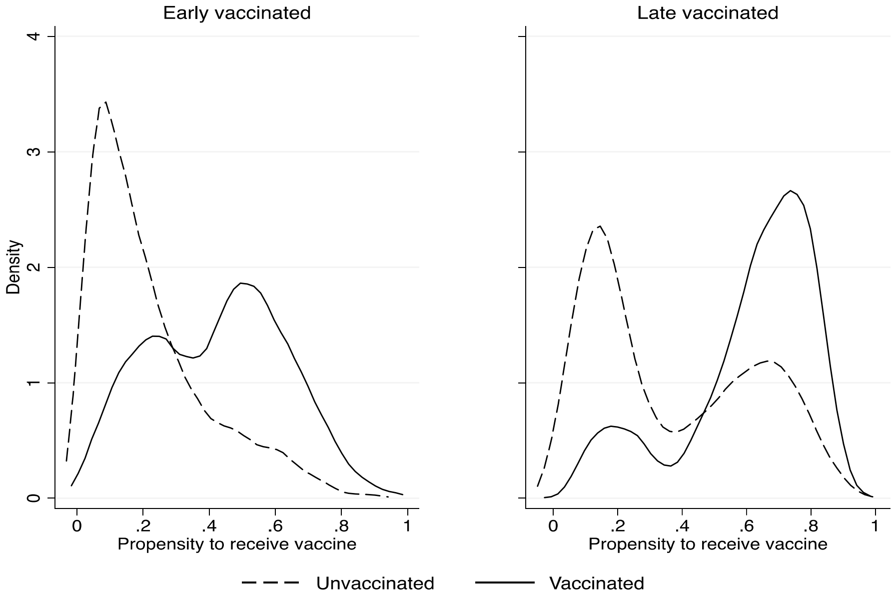

Fig. 2 Density plots of the distribution of the probability of vaccination among vaccinated and unvaccinated groups for early and late vaccinated groups

for the vaccinated and unvaccinated groups. From Fig. 2 we see that both plots overlap (i.e., extend across the same estimated probabilities of being vaccinated) quite well. In only about 2% of the cases (in the tails of the distributions), there are no estimated probabilities in common for vaccinated and unvaccinated individuals. There is therefore a sufficient common support for the distribution of propensity scores.

We then check that reweighting Eq. (1) with IPW in fact allows us to compare individuals of similar covariate distributions. Once units are reweighted, the characteristics of the constructed vaccinated and unvaccinated groups should not be significantly different. As is typical in the literature (see e.g., Morgan & Todd, 2008), we use the absolute standardized mean difference (ASMD) to compare the means of a given covariate between vaccinated and unvaccinated individuals before and after implementing IPW.[16] ASMD is used to assess the degree to which the characteristics are different on average between the vaccinated and unvaccinated individuals. Standardizing allows us to compare the difference in means across covariates. Figure 3 displays the resulting ASMDs for all covariates with and without propensity score adjustment to show the selection bias. We follow Imbens and Rubin's (2015)

[16] Specifically, $ASMD = \left(\overline{X}_T - \overline{X}_C\right) / \sqrt{\left(S_T^2 + S_C^2\right)/2}$, where $\overline{X}_T$ and $\overline{X}_C$ denote the covariate means of the treatment and comparison groups, respectively, and S_T^2 and S_C^2 represents the corresponding variances.

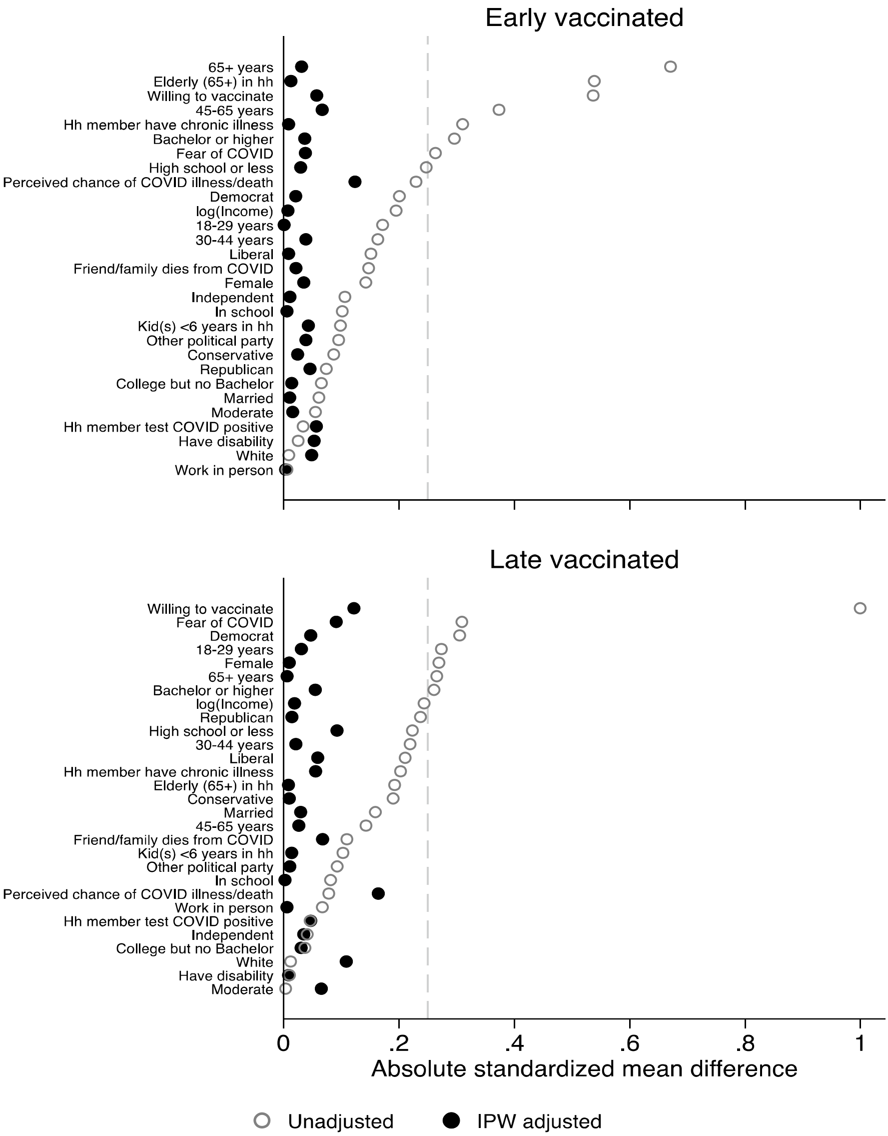

Fig. 3 A check for balance in covariates between vaccinated and unvaccinated individuals using the absolute standardized mean difference with and without IPW adjustment for early and late vaccinated

recommended rule of thumb that an ASMD not exceeding one quarter is considered balanced. The generated output shows that when setting the threshold for the mean difference to one quarter, all covariates were balanced after the IPW adjustment, indicating a significant reduction in the baseline bias due to IPW.

4.1.2 Parallel trend assumption

The empirical analysis also relies on the conditional parallel trends assumption, which asserts that the vaccinated and unvaccinated with similar characteristics would follow a similar trend in risk-taking behavior in the absence of vaccination. To test this assumption, we estimate analogous treatment effects for the pre-vaccination periods of the early and late vaccinated. Specifically, we estimate one pre-treatment effect for early vaccinated individuals, corresponding to the period just prior to vaccination (Tables 2 and 3, Panel A); and two pre-treatment effects for late vaccinated individuals, corresponding to one and two periods before vaccination (Tables 2 and 3, Panel B). If parallel trends are supported, we expect these coefficients to be statistically insignificant and small (i.e., close to zero), indicating that there are no significant differences in risk-taking behaviors of vaccinated and unvaccinated individuals in the absence of vaccination.

When we consider the risk-taking activities separately (Table 2), pre-treatment estimates are close to zero and insignificant for all risky activities except attending bars and participating in protests for those vaccinated early (Panel A, columns 5 and 8) and visiting friends and family for the late vaccinated (Panel B, column 1). In Table 3, we report results for four different aggregations of COVID-19 risky activities. Using these composite measures, all pre-trend estimates are insignificant and close to zero. In short, our pre-vaccination results provide strong evidence in support of the conditional parallel trends assumption for both the early and late vaccinated.

4.2 Main results

Thus far, we have shown all the evidence supports the assumption that our model, as outlined in Section 3, makes the vaccinated and unvaccinated individuals sufficiently comparable. Thus, we leverage vaccine status to identify the parameters of Eq. (1).

Table 2 displays estimates for each of the eleven COVID-19 risky activities using the event study linear probability version of our model (columns 1–11). The first row presents the overall average estimates (i.e., the average of the estimated effects for early and late vaccinated, weighted by the share of individuals in each group) followed by event study estimates for early and late vaccinated individuals in Panels A and B, respectively. Parallel trends hold for all cases except three outcomes—attending bars and participating in protests for the early vaccinated (Panel A, columns 5 and 8) and visiting friends and family for the late vaccinated (Panel B, column 1)—preventing us from drawing any definitive conclusions about their impact of vaccination on these behaviors. Thus, we focus our discussion on the remaining risky activities.

We observe an overall 3–11 pp statistically significant increase in risky activities—specifically, participation in indoor dining (11 pp), religious services (11 pp), attendance at campaigns (3 pp), and shopping (5 pp)—among vaccinated individuals, regardless of the vaccination timing (row 1 of Table 2). These effects appear to vary based on the timing of one's vaccination. At the time of vaccination, the later vaccinated individuals showed a greater increase in the magnitude of the risk-taking effect (7–17 pp) (Panel B, columns 2, 4, and 5), than those who received the vaccine earlier (4–8 pp) (Panel A, columns 7 and 10). Furthermore, although both the early and late vaccinated groups demonstrated heightened risk-taking behavior by increasing engagement in indoor dining and attending religious services, the later vaccinated group did so at the time of vaccination, while the early vaccinated group showed an increase one period after vaccination, at which point the general public had vaccine access. Furthermore, the impact of vaccination on attending campaigns and in-person work seems to be driven by the early-vaccinated group, being the only group for which we observe a significant and positive increase in these activities. Notably, the discernable effect is observed only at the time of vaccination for this group and diminishes in subsequent periods (Panel A of Table 2). The later vaccinated cohort also showed a significant positive increase in risky behavior by attending bars (7 pp). One period after vaccination we also see the early vaccinated group shows a 5–9 pp increase in the probability of risk-taking behavior by traveling by plane, using public transport, and shopping, and a 5 pp decrease in dining outdoors, despite a loss in precision. Conversely, visiting friends and family display only small changes (average 2 pp) in the coefficients relative to the period prior to vaccination, suggesting that vaccination does not appear to have a significant effect on risk-taking behavior associated with this activity.

Table 3 shows the estimates of the effect of vaccination on the aggregate risk measures using the Poisson model specified in Section 3. The first row shows the overall average effects across both groups, and subsequent rows show event study estimates for the early and late vaccinated in Panel A and B respectively. We find that being vaccinated, regardless of differences in the time of vaccination, resulted in an overall average increase in risk-taking (row 1, columns 1–4). These amount to a 12% increase in the number of risk-taking activities (column 1) and a 16 pp increase in the probability the respondent or a family member would contract COVID-19 (column 4).

The effects vary based on the timing of vaccine availability. For the early vaccinated, we find a sizable increase in risk-taking at the time of vaccination, albeit insignificant in models 3 and 4 (Panel A). This group further increased their risk-taking behavior one period later. These results suggest that individuals took on relatively more risk, the longer ago they were vaccinated. Those vaccinated later also increased their risk-taking following vaccination but by a relatively smaller magnitude compared to the early vaccinated one period later (Panel B). The early vaccinated group increased the mean number of COVID-19 risky activities by 7–17% (Panel A), while those who got vaccinated later increased their mean number of activities by about 14% (Panel B). While there is some loss in precision for models 2–4, overall the results remain qualitatively the same across all specifications—once vaccinated we observe a sizable increase in risky activities (columns 1 and 2) and risk (columns 3 and 4).

4.3 Placebo test

Much effort has been devoted to addressing endogeneity issues arising from people choosing whether to get vaccinated. Nevertheless, there might still be unobservables, such as impulsivity or impatience, influencing one's vaccination decision and risk-taking behavior. We therefore check for further endogeneity issues by conducting a placebo test, using the flu vaccine as a placebo treatment for COVID-19.

Two features of the flu vaccine make it a credible placebo treatment in our setting. First, we expect both vaccines to share similar influences, since the decision to take either vaccine occurred during an unusual period of heightened vaccine awareness and hesitancy.[17] However, a potential bias arises from healthcare providers recommending both vaccines simultaneously. Individuals opting for the flu vaccine first may be influenced by factors such as age, comorbidities, occupation, and location. Assuming these factors similarly influence the choice of both vaccines, our model effectively controls for these influences, as previously outlined. Second, unlike the COVID-19 vaccine, the flu vaccine does not offer protection against coronavirus. Thus, if a lower risk of contracting COVID-19 is the means by which the COVID-19 vaccine affects risky behaviors, then one would be suspicious if the flu vaccine also affects it. A non-zero effect would suggest that there are still unobservables at play.

While flu vaccines do not protect against COVID-19, they reduce the risk of a simultaneous double infection. Thus, we use respondents' flu vaccine status in wave 3 (November to December 2020), prior to the peak of the flu season and the availability of the COVID-19 vaccine when the COVID-19 vaccine could not affect behavior. Using the same Poisson regression as in Section 3, we estimate the effect of the flu vaccine (in lieu of the COVID-19 vaccine) on the number of risky activities. We find that the flu vaccine has a small and statistically insignificant association with risk-taking in wave 3 (0.014, s.e. = 0.036).[18] In short, the evidence supports our claim that unobservables have no discernible effect on our results.

4.4 Reverse causality

One might still be concerned that the results are being driven by reverse causality whereby being vaccinated did not cause individuals to take more risk, but instead, how risky individuals are, drives their decision to be vaccinated. In this regard, we recall the evidence presented from the literature that showed that individuals who are risk-takers, when defined as less likely to comply with mask-wearing and social distancing, are less likely to get vaccinated (Dabla-Norris et al., 2021). This therefore implies that the presence of reverse causality would, if anything, introduce a downward bias on the coefficients of interest, making the observed effect smaller than it should be. The results from the aggregate measures for the early vaccinated (Table 3, Panel A) show that the lagged effects, which cannot be affected by reverse causality, are more than twice the magnitude of the contemporaneous effects. The relatively smaller contemporaneous effects support the claim that reverse causality

[17] The new flu vaccine released for 2020 should ideally be taken in September/October which coincided with the start of the pandemic.

[18] This finding is unchanged even if we consider the effect in wave 4 when COVID-19 vaccine becomes available and could possibly be affected by potential challenges introduced by the availability of the COVID-19 vaccine (0.042, s.e. = 0.050).

causes a downward bias in the estimates. With a downward bias, if one believes we are unable to fully control for all potential bias from risk-taking on vaccination, the results presented should be accepted as at least conservative estimates of the vaccination on risk-taking.

5 Conclusion

This research is one of the first to estimate the impact of vaccines on COVID-19 risk-taking behavior. Although previous research examines the relationship between health interventions and risky behavior, the findings vary by the type of intervention, the context in which it is implemented, and the specific behaviors under consideration. We evaluate eleven COVID-19 risky behaviors that are believed to pose a high risk of COVID-19 transmission. Using the first individual-level panel data available, we address endogeneity biases resulting from pre-vaccine differences between vaccinated and unvaccinated groups through a variety of panel methods and tests. This creates an environment where both groups are sufficiently comparable allowing us to identify the effect of vaccination on COVID-19 risk-taking. Several tests validate our approach.

We provide two basic findings. First, being vaccinated induces risk-taking in some risky activities and not others. Specifically, the effects seem to be more pronounced for dining indoors, attending religious services in person, attending campaigns, and shopping. Second, these results are heterogeneous in the timing of vaccination. Early vaccinated groups (typically vulnerable or essential workers with early vaccine access) participate in more risky behaviors than those vaccinated later (typically the general public without early vaccine access) and they do so at different periods relative to the time of vaccination. However, at the time of vaccination, the later vaccinated individuals showed a greater increase in the magnitude of the risk-taking effect (7–17 pp) than those who received the vaccine earlier (4–8 pp).

This behavioral response could be rational risk-taking behavior arising from the lower likelihood of contracting COVID-19 after vaccination. Alternatively, it might be that people overestimate the efficacy of the vaccine. The resulting increased risky behavior in the presence of safety devices is what Viscusi (1984) termed the “lulling effect”. Thus, while the vaccine availability may reduce the risk of contracting COVID-19, it also contributes to the further spread of the virus by incentivizing risk-taking in the short term. While this study does not suggest that the vaccine was a bad policy choice, our findings show the importance of accounting for the unintended effects that might result from policy implementation of a similar nature. This is especially important when dealing with a highly contagious and life-threatening virus like coronavirus, which involves unique issues for individual behavior.

Appendix

Table 4 National representativeness of the sample (Adults 18+)

	Baseline Sample Average	ACS 2020 (5-Yr average)
Demographic Variables		
Age (yrs.)	53	48
18 to 29 years	0.08	0.21
30 to 44 years	0.23	0.25
45 to 64 years	0.40	0.33
65 and older	0.29	0.21
Female	0.53	0.51
HH Income ($)	95238	95395
High school or less (25+)	0.29	0.38
College but no Bachelor (25+)	0.30	0.29
Bachelor or higher (25+)	0.41	0.33
White	0.62	0.60
Black	0.11	0.12
Hispanic	0.15	0.18
Other race	0.12	0.10
Married	0.58	0.50
Liberal	0.29	0.26[a]
Moderate	0.35	0.365[a]
Conservative	0.36	0.375[a]
Observations	2720	

Note: The table compares demographic characteristics of the sample with those of the population of adults 18+ in the U.S. as measured by the Census Bureau's 2020 (and 2019 when 2020 not available) ACS (5-year average estimates) unless otherwise indicated. All values reported in the table represent proportions unless otherwise stated. Baseline sample averages calculated using wave 3 which is the wave immediately prior to vaccine availability. All averages reported in the table are weighted using individual weights from respective surveys. [a]Retrieved from Gallup poll 2020 publication. Normalized to add to 100% for comparability with the baseline sample, which originally added to 96% since 4% indicated no opinion.

Acknowledgements The authors acknowledge the Social Policy Institute at Washington University in St. Louis for access to the Socio-Economic Impacts of COVID-19 Survey data, which supported the analysis in this manuscript. The collection of the data used in this study was funded by Mastercard Center for Inclusive Growth, the JPMorgan Chase Foundation, the Annie E. Casey Foundation, and Centene Corporation. The funders had no role in the study design, the analysis, or the preparation of this manuscript. In addition, the authors would like to thank the Editor and an anonymous referee, as well as David Slichter, Marlon Tracey, and Binghamton University's Labor Group for their valuable input and feedback.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

References

Agrawal, V., Sood, N., & Whaley, C. M. (2022). The ex-ante moral hazard effects of COVID-19 vaccines. *NBER Working Paper, 30602.*

Akesson, J., Ashworth-Hayes, S., Hahn, R., Metcalfe, R., & Rasooly, I. (2022). Fatalism, beliefs, and behaviors during the COVID-19 pandemic. *Journal of Risk and Uncertainty, 64*, 147–190.

Albrecht, D. (2022). Vaccination, politics and COVID-19 impacts. *BMC Public Health, 22*, 1–12.

Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., & Yang, D. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics, 191*, 104254.

Aw, J., Seng, J. J. B., Seah, S. S. Y., & Low, L. L. (2021). COVID-19 vaccine hesitancy—A scoping review of literature in high-income countries. *Vaccines, 9*, 900.

Barrios, J. M., & Hochberg, Y. (2020). Risk perception through the lens of politics in the time of the COVID-19 pandemic. *NBER Working Paper, 27008.*

Borusyak, K., Jaravel, X., & Spiess, J. (Forthcoming) Revisiting event study designs: Robust and efficient estimation. *Review of Economic Studies*.

Brewer, N. T., Cuite, C. L., Herrington, J. E., & Weinstein, N. D. (2007). Risk compensation and vaccination: Can getting vaccinated cause people to engage in risky behaviors? *Annals of Behavioral Medicine, 34*, 95–99.

Bruine de Bruin, W., Saw, H. W., & Goldman, D. P. (2020). Political polarization in US residents' COVID-19 risk perceptions, policy preferences, and protective behaviors. *Journal of Risk and Uncertainty, 61*, 177–194.

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics, 225*, 200–230.

Center for Disease Control and Prevention. (2023). *People with certain medical conditions.* Retrieved from https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/people-with-medical-conditions.html.

Dabla-Norris, E., Khan, H., Lima, F., & Sollaci, A. (2021).Who doesn't want to be vaccinated? Determinants of vaccine hesitancy during COVID-19. *IMF Working Paper, 2021/130*.

Dave, D., Friedson, A., Matsuzawa, K., McNichols, D., Redpath, C., & Sabia, J. J. (2021). Risk avoidance, offsetting community effects, and COVID-19: Evidence from an indoor political rally. *Journal of Risk and Uncertainty, 63*, 133–167.

De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review, 110*, 2964–2996.

Dhanani, L. Y., & Franz, B. (2022). A meta-analysis of COVID-19 vaccine attitudes and demographic characteristics in the United States. *Public Health, 207*, 31–38.

Dibonaventura, M. D., & Chapman, G. B. (2008). Do decision biases predict bad decisions? Omission bias, naturalness bias, and influenza vaccination. *Medical Decision Making, 28*, 532–539.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics, 225*, 254–277.
Heffetz, O., & Rabin, M. (2023). Estimating perceptions of the relative COVID risk of different social-distancing behaviors from respondents' pairwise assessments. *Proceedings of the National Academy of Sciences, 120*, e2219599120.
Imbens, G. W. (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics, 86*, 4–29.
Imbens, G. W., & Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge: Cambridge University Press.
Kniesner, T. J., & Sullivan, R. (2020). The forgotten numbers: A closer look at COVID-19 non-fatal valuations. *Journal of Risk and Uncertainty, 61*, 155–176.
Lakdawalla, D., Sood, N., & Goldman, D. (2006). HIV breakthroughs and risky sexual behavior. *The Quarterly Journal of Economics, 121*, 1063–1102.
Latkin, C. A., Dayton, L., Yi, G., Colon, B., & Kong, X. (2021). Mask usage, social distancing, racial, and gender correlates of COVID-19 vaccine intentions among adults in the US. *PLoS ONE, 16*, e0246970.
Lovelace, B. J. (2021). *WHO urges fully vaccinated people to continue to wear masks as delta Covid variant spreads*. CNBC. Retrieved from https://www.cnbc.com/2021/06/25/delta-who-urges-fully-vaccinated-people-to-continue-to-wear-masks-as-variant-spreads.html.
Moghtaderi, A., & Dor, A. (2021). Immunization and moral hazard: The HPV vaccine and uptake of cancer screening. *Medical Care Research and Review, 78*, 125–137.
Morgan, S. L., & Todd, J. J. (2008). A diagnostic routine for the detection of consequential heterogeneity of causal effects. *Sociological Methodology, 38*, 231–281.
Peltzman, S. (1975). The effects of automobile safety regulation. *Journal of Political Economy, 83*, 677–725.
Schneider, E. C., Shah, A., Sah, P., Moghadas, S. M., Vilches, T., & Galvani, A. P. (2021). *The U.S. COVID-19 vaccination program at one year: How many deaths and hospitalizations were averted?* Commonwealth Fund. Retrieved from https://doi.org/10.26099/3542-5n54.
Viscusi, W. K. (1984). The lulling effect: The impact of child-resistant packaging on aspirin and analgesic ingestions. *The American Economic Review, 74*, 324–327.
Viscusi, W. K., & Cavallo, G. O. (1994). The effect of product safety regulation on safety precautions. *Risk Analysis, 14*, 917–930.