Influencing COVID-19 Vaccination Opinions via Engaging Media

Erik Zinn, Lester Yang, Sundeep Mahensaria
W241 Section 005 - Final Project - April 2021

Abstract

In this experiment, we study the effects of exposing people to facts about COVID to see if their willingness to receive a COVID vaccination changes. Our treatment utilizes engaging media to supplement the facts we expose participants to, to see if the engaging media we provide has any additional persuasive effects. We ran our experiment on Berkeley Xlab students using a survey with an ROXO experimental design, in hopes of running a difference in differences analysis. However, the data we got back was heavily top-coded with a 93.6% of participants *Pre-Treatment Decision* to vaccinate being *Definitely Yes* or that they had already been *Vaccinated*. So we were unable to perform the difference in difference analysis we had originally hoped to perform. But we were able to analyze the participants *inclination* to vaccinate, and we did observe a slight increase in participants' *inclination* to get vaccinated, though it was not statistically significant.
Our code and survey are available on Github[1].

1 Introduction

The COVID-19 pandemic has affected lives to an extent that shifted the paradigms of normalcy in our everyday lives. From virtual learning and working from home, to wearing masks wherever necessary. The COVID-19 vaccine has become the only hope in bringing life back to normalcy. Vaccines are able to create sterilizing communities in which vaccinated individuals will not transmit the virus to unvaccinated individuals. However, few vaccines, even those most effective, guarantee this sterilizing immunity. The inactivated poliovirus vaccine did not block transmission or infection, but after 6 years of being introduced, in 1961, paralytic polio cases had dropped by more than 90%, when only 54% of the US population was vaccinated [1].

It is still unclear what percentage of the population needs to be vaccinated or recovered from COVID-19 before herd immunity kicks in, as the head of the National Institute of Allergy and Infectious Diseases (NIAID), Dr. Anthony Fauci acknowledged on CNN, "I think we have to be honest and humble. Nobody really knows for sure." Nonetheless, it should be made clear that the purpose of the COVID-19 vaccine is to prevent symptomatic infection, hospitalization, and death, but not the viral spread. For example, efficacy studies on available vaccines have only produced results that suggest, but not prove, AstraZeneca vaccine stymies viral spread [2]. In combination with the discovery of viral mutants that dramatically reduced the rate of symptomatic protection of existing vaccines, it is not hard to imagine that COVID vaccines will need to be constantly updated, and a large portion of a community revaccinated in order to prevent rapidly rising cases. In fact, the CEO of Pfizer, the manufacturer for the Pfizer-BioNTech vaccine recently stated that annual COVID-19 booster shots will be required to stay ahead of new variants [3]. Due to these conditions, promoting vaccination and achieving a high percentage of vaccination in the population is more important than ever.

However, skepticism has only grown. Skepticism about vaccinations is not new - and it is no different, if not more, for the Covid-19 vaccines. Despite the continued emphasis of the importance of getting

[1] `https://github.com/smahensaria/W241-final-project`

vaccinated by governments, global health organizations, and news media, the vaccination rate remains relatively low (compared to what is required for achieving herd immunity). As of April 2021, the vaccine has been out for over 4 months, nearly 900M people worldwide and 200M in the US have been vaccinated [4] - yet in some US states over 30% doses remain unused [5]. Evidently, convincing the public to get vaccinated is a difficult task, but this distrusting dynamic between health experts and the public is nothing new. During the 1918 influenza pandemic, the Anti-Mask League of San Francisco was established to protest the ordinance that required mask wearing in San Francisco, California. Violators of the ordinance were fined or jailed for not wearing masks.

2 Hypothesis

Countless studies have shown the effectiveness of visualization in aiding information consumption and persuasion. A clinical trial in 2005 (n=87) investigating the effect of pictograms on medicine labels for understanding and adherence has shown that a high adherence of greater than 90% was found for 54% of the treatment group compared to just 2% of the control group [6]. A review of the literature for the effects of text illustrations showed that people following directions with text and illustrations do approximately three times better than those following directions without illustrations [7]. A more related study on the pervasiveness of medical information by McCabe and Castel [8] has shown that adding pictures of brain scans and mentioning cognitive neuroscience make people more inclined to believe what they are reading. Even though the supporting literature is plentiful, it is unclear whether these techniques are as effective for a topic as polarizing as vaccines. Our hypothesis is that by sharing COVID related facts with an engaging media, people's opinions of COVID vaccination are more likely to be swayed by the facts. In hopes to leverage these research findings, we plan to share facts with engaging media to increase the subject's willingness to get vaccinated, when compared to sharing the same facts in text alone.

3 Treatment and Control

The treatment and control are designed to provide the same core facts regarding pandemics to the subjects. The difference arises in the way the survey questions engage the subjects. For example, in one of the treatment questions, a photo taken during the 1918 pandemic of an open-air barbershop with people wearing masks is shown. The treatment subject is then asked to think about where they think the photo was taken, they are then shown the answer on the following page. Whereas the same question for the control group is simply a sentence describing mask wearing.

Example Treatment question:
To the right is a photograph of a barbershop taken in 1918, during the Spanish Flu. Notice two features of this photograph. All of the barbers are wearing masks; and, The shop is outside as a method of reducing the spread of the Spanish Flu virus.

If you had to guess, in what city or state do you think this photograph was taken? Keep your answer in mind while moving on. We will give you the correct answer on the next page.

[on the following page] This photo was taken at the University of California, Berkeley!

Example Control question:
Mask wearing was common during the 1918 Spanish Flu at the University of California, Berkeley. **[Photo Excluded]**

The methods of engagement we incorporated into our survey was not limited to visualizations. We also include multiple choice questions, questions that require text answers, and rephrasing of a statement into questions. The complete Qualtrics survey is available for reference on the Github repository.

4 Experiment Design

We used Qualtrics to survey UC Berkeley XLab participants for our experiment. We achieved randomization by applying the Qualtrics survey block randomizer to our Treatment and Control survey blocks to generate equal exposure to each condition. To test our hypothesis we designed a difference-in-differences experiment with an ROXO design, taking measurements before and after the exposure condition as illustrated in Figure 1.

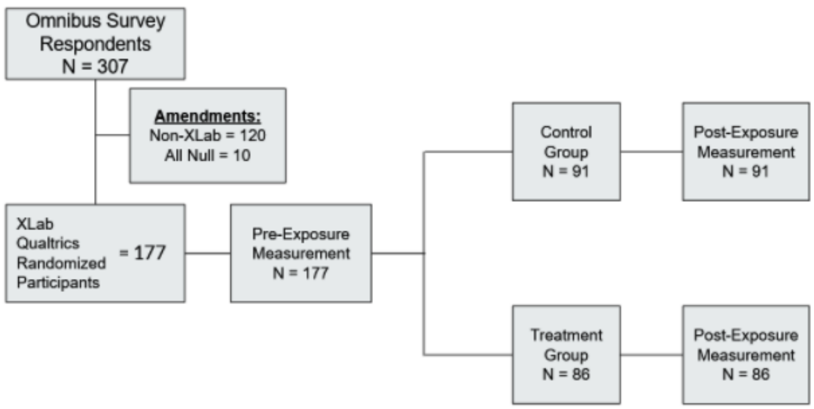

Figure 1: Experimental flow diagram. 177 XLab participants, 91 control participants, 86 treatment participants.

5 Outcome Measures

To operationalize the subjects willingness to receive the COVID19 vaccine we used the following survey question pre- and post-exposure (this outcome will be referred to as the participants vaccination *Decision*).

1. Given that the COVID-19 vaccine became available, would you accept the vaccine if it meant protecting friends, family, or at-risk groups?
   (a) No, definitely not
   (b) Unsure, but leaning towards no
   (c) Unsure, but leaning towards yes
   (d) Yes, definitely
   (e) I have already been vaccinated

By utilizing an ordinal outcome variable we are able to normalize the concept of the participants willingness to receive the vaccine. As opposed to using a slider from 0-100 (low implying declining the vaccination & high values being inclined to vaccinate) where the participants would have to map their willingness to vaccinate to a numeric value.

While we were worried about the possibility of many participants being supportive of the vaccination, we believe that using ordinal measures of the participants willingness to vaccinate or *Decision* would lead to more interpretable responses. We plan to leverage a difference-in-differences experimental design to capture any changes in participants *Decision* to vaccinate and account for the participants pre-exposure preferences to more accurately capture our treatment effect.

The second outcome measure we were interested in is the participants' *Inclination* to vaccinate - this is a post-exposure measurement where we ask participants:

2. Overall, the information provided above makes me
   (a) Much less inclined to be vaccinated
   (b) A little less inclined to be vaccinated
   (c) No less or more inclined to be vaccinated
   (d) A little more inclined to be vaccinated
   (e) Much more inclined to be vaccinated

This outcome measure will help us capture the participants' perception of the information provided [control or treatment] and how they believe that information has influenced their *Inclination* to vaccinate.

6 Data Analysis

Given we were using Berkeley Xlab to run our experiment, we were quite aware that the survey respondents would primarily be undergraduate students, with a liberal bent of mind. Figure 2 shows the distribution across gender, race, student role at Berkeley, party affiliation and religion. We observed some skews in our units - key amongst these include:

- Majority (115 of 177) were Cisgender Women
- Majority (110 of 177) were Asian
- More than 50% (93 of 177) were Democrats
- Almost 50% were Atheists

However, we do notice that the randomization worked well and the units were evenly distributed between treatment and control demonstrating good covariate balance. A full covariate balance check is available in the Appendix [§13.1].

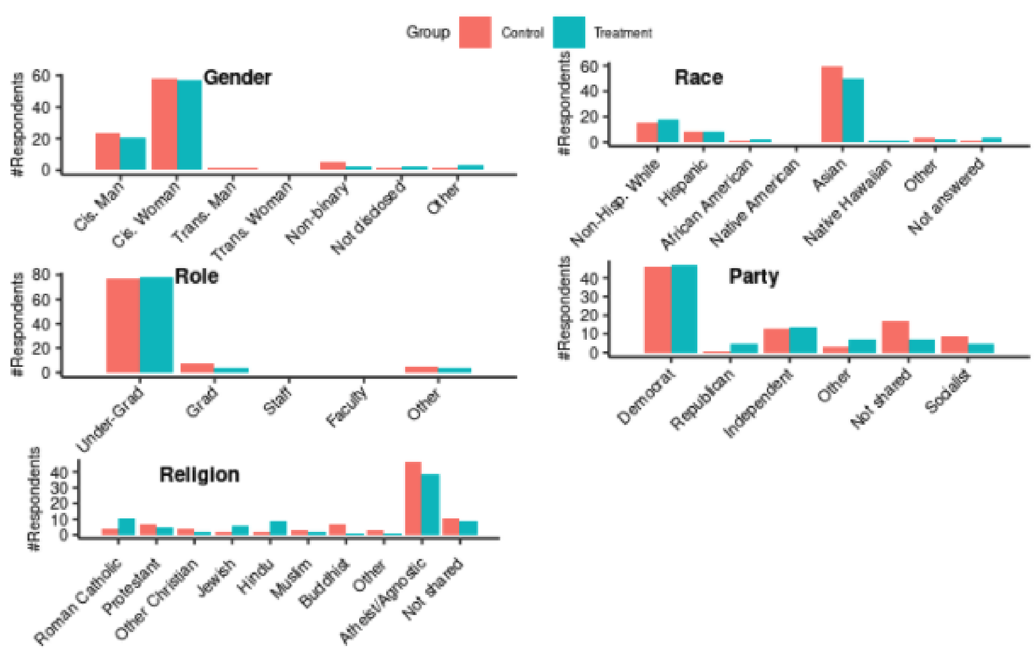

Figure 2: Demographic distributions of units.

As mentioned in the Treatment section above, in the treatment part of the survey we presented the same set of facts, as in control, to the respondents in an interactive manner. We noticed that respondents in treatment spent a mean time of 186 secs (median of 172 secs) in this section compared to mean of 60 secs (median of 47 secs) by those in the control group. Figure 3 shows a distribution of the time spent by respondents in treatment and control.

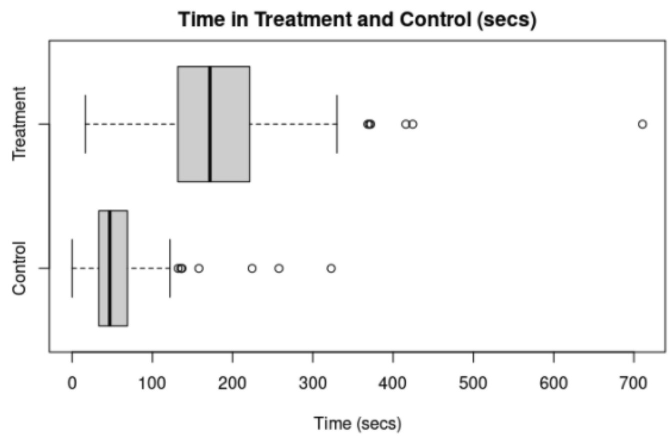

Figure 3: Time spent by participants in treatment and control sections of the survey, respectively.

Pre-exposure *Decision* to vaccinate (Figure 4): The liberal predisposition of our units became evident, as we observed that 58% respondents (102 of 177) were already vaccinated and another 34% (61 of 177) were a 'Definitely Yes' to get vaccinated in the pre-exposure section of our survey. Of the remaining, 9 were 'Unsure, but leaning towards yes', with 1 each in 'Unsure, but leaning towards no' and 'Definitely No'.

Post-exposure *Decision* to vaccinate (Figure 4): The post-exposure distribution was no different, though we noticed 5 units changed their decision - 2 additional ones claimed to be vaccinated, 2 moved from being 'Unsure, but leaning towards yes' to 'Definitely Yes' while 1 moved down from 'Definitely Yes' to 'Unsure but leaning towards yes'.

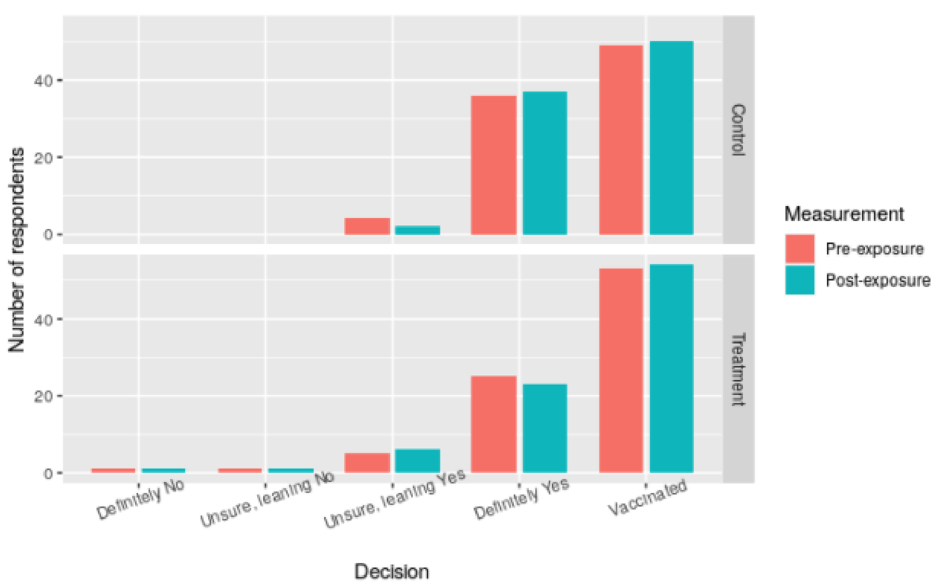

Figure 4: Distribution of pre- and post-exposure *Decision* to vaccinate.

7 Results

As mentioned above, given the survey units distribution and the duration of the experiment, we were expecting limited, if any, effect in our outcome measure 1 (i.e., Change in unit's *Decision* to get vaccinated). This was further corroborated by the data analysis (Figure 4 above). Though not very insightful, for reasons of completeness of analysis, we created two models - Linear Regression and Ordinal Logistic Regression for this outcome measure (see Models section below for details).

Next, we turned our focus on outcome measure 2 (i.e., Change in unit's *Inclination* to get vaccinated, after reading the facts on COVID). We observed that all respondents were either 'neither less nor more' or more inclined after reading the facts, regardless of whether they were in treatment or control (see Figure 5). Further, given the categorical nature of the *Inclination* dependent variable, which requires a more 'complex to interpret' Ordinal Logistic Regression, we decided to bin the responses into a binary variable *More Inclined* (see Figure 6). This enabled us to define a Logistic Regression model (see Models section below for details).

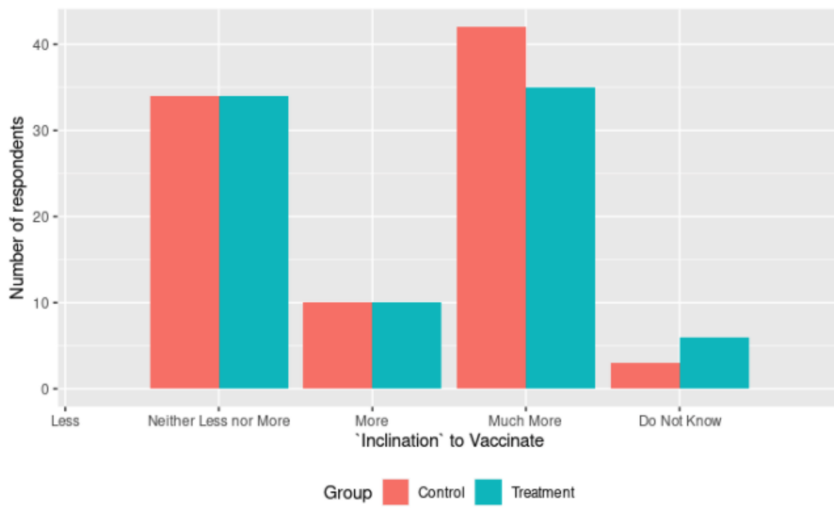

Figure 5: Units change in *Inclination* to get vaccinated after reading the facts on Covid - a categorical dependent variable.

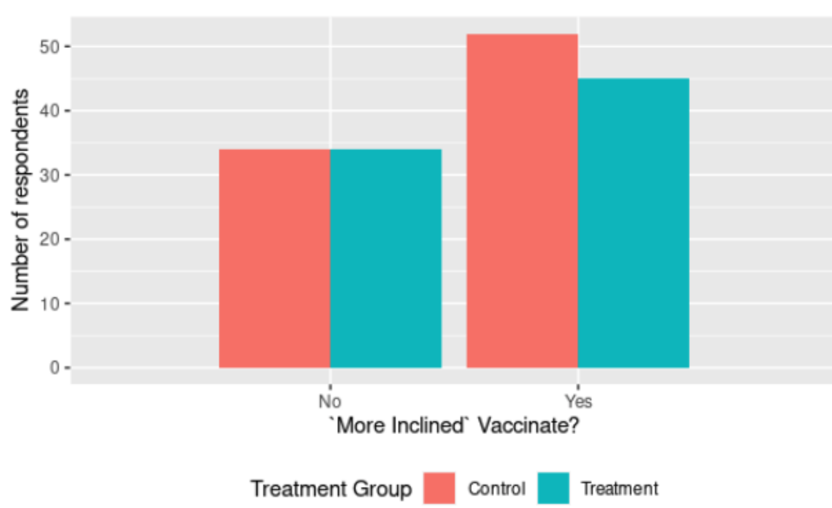

Figure 6: Units inclination to vaccinate after reading the facts on Covid - a binary outcome variable.

8 Models

Given the categorical nature of our outcome measures, we compared results from Linear Regression (OLS) and Ordinal Logistic Regression (OLR). Additionally, having transformed our second outcome measure (Change in *Inclination*) to a binary outcome variable, we used Logistic Regression to model our results. We chose linear regression in our models due to the ordinal nature of our outcome variables.

Outcome measure 1 - Change in *Decision*: These models are for academic purposes only, because as mentioned in the data analysis section above, we did not observe a material change in decision of our units. Table 1 captures the results of these models. From Column (1), the linear regression model, we observe a strong, statistically significant positive effect (0.87) of pre-treatment decision of the units and a negative effect (-0.623) of the treatment at 10% significance level - this seems to be a modeling anomaly as we only observed one unit change her decision in the negative direction. Additionally, from Column (2), the ordinal logistic regression model, we observe that the odds ratio of *Treated*, $e^{-24.665}$, is not different from zero[2].

Table 1: Change in *Decision* to get vaccinated
Observed a negative effect of the treatment in the model using linear regression (Model 1), while the OLR model shows an effect not different from zero (Model 2). Pre-treatment Decision is statistically significant in both models.

	Dependent variable:	
	Post_Decision	Post_Decision
	OLS	*ordered logistic*
	Model 1	Model 2
	(1)	(2)
Treated	-0.623^{*}	-24.665
	(0.359)	(29.335)
Pre-treatment Decision	0.870^{***}	6.829^{***}
	(0.071)	(1.190)
Treated:Pre-treatment Decision	0.130^{*}	5.686
	(0.075)	(7.271)
Constant	0.623^{*}	
	(0.340)	
Robust Std Errors	Yes	No
Observations	165	165
R^2	0.927	
Adjusted R^2	0.926	
Residual Std. Error	0.166 (df = 161)	
F Statistic	684.013^{***} (df = 3; 161)	

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

[2] In R polr the ordinal logistic regression model is parameterized as:

$$logit(P(Y \leq j)) = \beta_{j0} - \eta_1 x_1 - \cdots - \eta_p x_p$$

Outcome measure 2 - Change in *Inclination*: To model the change in *Inclination* of units to get vaccinated, we used two sets of models, namely:

1. Linear Regression (OLS) and Ordinal Logistic Regression (OLR) on categorical outcome variable *Inclination*, given the ordinal nature. [Table 2]
2. Logistic Regression on the binarized outcome variable, *More Inclined*. [Table 3]

Table 2 (below) captures the results of the first set of models. From Column (1) we observe that using OLS on a short model, the treatment effect leads to units being 1.322 times 'more or much more inclined' to get vaccinated, though not statistically significant. Using OLR, from Column (2) we observe that the odds ratio of *Treated*, $e^{2.816}$, is 16.7 resulting in the interpretation that for units that were treated, the odds of being 'more or much more inclined' to get vaccinated is 16.7 than units in control - though this is not statistically significant. However, the pre-treatment decision is statistically significant. Column (3) shows the results of OLR when controlled for demographic covariates of gender, race, student role, party affiliation and religion. In this model the odds ratio of *Treated*, $e^{2.107}$, is 8.22 resulting in the interpretation that for units that were treated, the odds of being *more or much more inclined* to get vaccinated is 8.22 than units in control. Though here too, this is not statistically significant. Additionally, we observe that some genders, races and religions are significant at 10% level. We believe that these are modeling anomalies as the number of units with these demographics were very small.

Table 2: Change in *Inclination* to get vaccinated.
From the OLS model we observe that units in treatment are 1.322 times 'More or Much More Inclined' to get vaccinated (Model 1). Using OLR we observe the units in treatment to have an odds ratio of 16.7 of being 'More or Much More Inclined' to get vaccinated when we don't control for demographic covariates (Model 2). Controlling for demographic covariates we observe the units in treatment to have an odds ratio of 8.22 of being 'More or Much More Inclined' to get vaccinated (Model 3). Pre-treatment Decision is statistically significant in all models. Please refer to §13.2.1 for the unabridged table.

	Dependent variable:		
	Inclination	Inclination	
	OLS	*ordered logistic*	
	Model 1	Model 2	Model 3
	(1)	(2)	(3)
Treated	1.322	2.816	2.107
	(1.131)	(2.260)	(2.915)
Pre-treatment Decision	0.576^{***}	1.146^{***}	1.258^{***}
	(0.195)	(0.385)	(0.456)
Treated:Pre-treatment Decision	−0.311	−0.654	−0.532
	(0.246)	(0.496)	(0.628)
Constant	1.488		
	(0.905)		
Gender, Race, Role, Party & Religion categories	No	No	Yes
Robust Std Errors	Yes	No	No
Observations	165	165	165

*p<0.1; **p<0.05; ***p<0.01

Notes: Base Gender: Cisgender Man. | Base Race: Non-Hispanic White.
Base Role: Undergraduate Student. | Base Party: Democrat. | Base Religion: Roman Catholic.
We treat Pre-treatment Decision as an ordinal variable for sake of simplicity in analysis.

Additionally, we applied Logistic Regression on the binarized outcome variable, *More Inclined*, using restricted and unrestricted covariates. Table 3 captures the results of the second set of models. Column (1), the result of a restricted model, shows a negative effect, $e^{-0.145}$ (i.e., odds ratio of treatment effect is 0.865 to be *More Inclined* to vaccinate). Column (2) shows, when interacted with pre-treatment decision, we notice a negative effect $e^{-1.006}$ at 10% significance level (i.e., odds ratio of 0.366 to be *More Inclined* to vaccinate). The effect is not statistically significant per these models.

We next controlled for pre-treatment decision and demographic variables gender, race, student role at Berkeley, party affiliation, and religion - see Column (3). We find that the treatment effect interacting with pre-treatment decision has an odds ratio $e^{-0.581}$, (i.e., odds ratio of treatment effect is 0.559 to be *More Inclined* to vaccinate). In this model also the effect is not statistically significant. However, other covariates, pre-treatment decision and some races are statistically significant at 5% or lower. While Asian was the predominant race in the data, Hispanic or Latino and Black or African American constituted 9% (16 of 177) and 2% (3 of 177) of the sample size. Thus, we believe these to be modeling anomalies, considering the 'shoe-leather wisdom'.

Table 3: Change in *More Inclined* to get vaccinated.
In the restricted model (Model 1) we observe the treatment has an odds ratio of 0.865 to be 'More Inclined' to vaccinate. When interacted with Pre-Treatment Decision (Model 2) the odds ratio is 0.366 to be 'More Inclined'. When we control for demographic covariates we observe an odds ratio of 0.559 to be 'More Inclined' (Model 3). None of these are statistically significant. The statistical significance of certain demographic covariates seems a modeling anomaly given their small proportion in the sample. Pre-treatment Decision is statistically significant in all models. Please refer to §13.2.2 for the unabridged table.

	Dependent variable:		
	MoreInclined		
	Model 1	Model 2	Model 3
	(1)	(2)	(3)
Treated	-0.145	4.371	2.193
	(0.321)	(3.074)	(3.189)
Pre-treatment Decision		1.242^{**}	1.304^{***}
		(0.550)	(0.497)
Treated:Pre-treatment Decision		-1.006	-0.581
		(0.668)	(0.691)
Constant	0.425^{*}	-5.159^{**}	-5.436^{**}
	(0.223)	(2.527)	(2.329)
Gender, Race, Role, Party & Religion categories	No	No	Yes
Robust Std Errors	Yes	Yes	No
Observations	165	165	165
Log Likelihood	-111.703	-106.765	-87.301
Akaike Inf. Crit.	227.406	221.529	236.602

*p<0.1; **p<0.05; ***p<0.01

Notes: Base Gender: Cisgender Man. | Base Race: Non-Hispanic White.
Base Role: Undergraduate Student. | Base Party: Democrat. | Base Religion: Roman Catholic.
We treat Pre-treatment Decision as an ordinal variable for sake of simplicity in analysis.

9 Power Analysis

We ran a power analysis of our short model that regressed *More Inclined* with our treatment using logistic regression. We simulated sample sizes varying from 50 to 500, having a similar proportion of *More Inclined* outcomes as in our sample of 177, across both treatment and control. We observe a power of 0.8 for our sample size. The power progressively increases and touches 0.99 at a sample size of 400. However, using 'shoe-leather wisdom' this seems to be a modeling anomaly because given our observed mix of 'liberal' units, we do not believe that sample size can improve our modeling power.

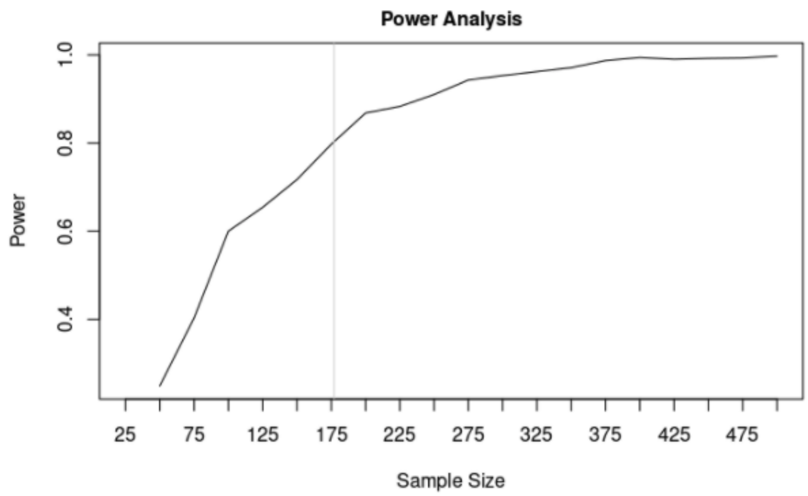

Figure 7: Power analysis.

10 Observations

While the experiment may not have turned out as we hoped we still learned a lot from our first field experiment, with a few key insights being:

- A majority of participants were *More Inclined* to vaccinate after we shared facts about COVID for both Treatment Control control groups.
- It is very important to have a representative sample of the population.
- One should apply 'shoe-leather wisdom' and not rely solely on regression models.

Another observation we had that was slightly puzzling was that 42% of respondents who were **already vaccinated** became **more inclined to vaccinate**, which shouldn't be possible since they've already been vaccinated. We believe this to be due to a *response bias*, where participants believe they should respond in a certain way after being exposed to our survey. Another caveat to our experiment was that after getting our results back and running our analysis, we were told that UC Berkeley had a campus wide initiative to vaccinate which started several weeks before our experiment, which may have also contributed to the low variance we observed in our sample.

11 Conclusion

To summarize our experimental results, we observed very little variance in our sample, with a 93.6% of participants *Pre-Treatment Decision* being *Definitely Yes* or that they had already been *Vaccinated*, this is clearly illustrated in Figure 4. We also noticed that participants *Decision* was so heavily top-coded, only 5 participants changed their *Decision* after exposure, thus we did not pursue the difference in differences model we had originally planned.

We found *Pre-Treatment Decision* to be statistically significant in all models where it was included which indicated that the participants *Pre-Treatment Decision* was the most indicative of what their *Post-Treatment Decision* or *Inclination* to vaccinate was after exposure. We did observe a slight increase in participants' *Inclination* to get vaccinated, though it was not statistically significant.

12 Future Research

Given the importance of vaccines to control COVID, it is imperative that we reach herd immunity. Considering that we observed a material increase in inclination of the units to vaccinate after being exposed to the facts, more research is needed to understand if using an interactive method to share the facts is effective or was our observation merely due to the liberal nature of the units in our sample. We need a more representative sample of the population to ascertain this. As of this writing there are still many US states where greater than 30% vaccines remain unused - thus, recruiting from these states may result in real insights.

References

[1] Sophie Ochmann and Max Roser. Polio. *Our World in Data*, 2017. `https://ourworldindata.org/polio`.

[2] Merryn Voysey et al. Single Dose Administration, And The Influence Of The Timing Of The Booster Dose On Immunogenicity and Efficacy Of ChAdOx1 nCoV-19 (AZD1222) Vaccine. 2021.

[3] Jared S. Hopkins. Annual covid-19 vaccine booster shots likely needed, pfizer ceo says. 2021. `https://on.wsj.com/3sDWahw`.

[4] Hannah Ritchie et al. Statistics and research coronavirus (covid-19) vaccinations. *Our World in Data*, 2021. `https://ourworldindata.org/covid-vaccinations`.

[5] Anna Edney and Drew Armstrong. Unused vaccines are piling up across u.s. as some regions resist. 2021. `https://bloom.bg/3v8ZR0i`.

[6] Ros Dowse and Martina Ehlers. Medicine labels incorporating pictograms: do they influence understanding and adherence? *Patient Education and Counseling*, 58:63–70, 2005.

[7] W. Howard Levie and Richard Lentz. Effects of text illustrations: A review of research. *Educational Communication and Technology*, 30:195–232, 1982.

[8] David McCabe and Alan Castel. Seeing is believing: the effect of brain images on judgments of scientific reasoning. *Cognition*, 107:343–352, 2008.

13 Appendix

13.1 Covariate Balance Check

Since randomization of the experiment is conducted through Qualtrics, it is unlikely that randomization has produced unbalanced treatment assignments. Nonetheless, we show in Figure 8 that with adjustment, there are no systematic differences in the covariates between treatment and control groups.

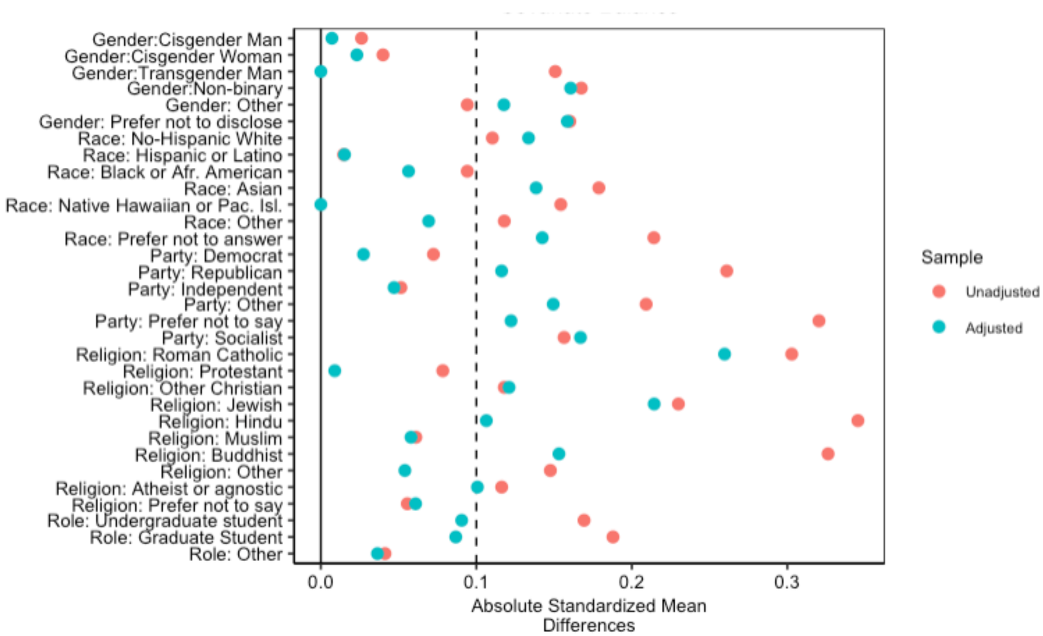

Figure 8: Absolute standardized mean differences in the measured covariates between treatment and control groups.

13.2 Unabridged Tables

13.2.1 Unabridged version of Table 2

	Dependent variable: Inclination		
	OLS	*ordered logistic*	
	Model 1	Model 2	Model 3
	(1)	(2)	(3)
Treated	1.322	2.816	2.107
	(1.131)	(2.260)	(2.915)
Pre-treatment Decision	0.576***	1.146***	1.258***
	(0.195)	(0.385)	(0.456)
Gender category: Cisgender Woman			−0.174
			(0.427)
Gender category: Transgender Man			16.274***
			(0.00000)
Gender category: Non-binary			−0.045
			(0.995)
Gender category: Other			15.721***
			(0.00000)
Gender category: Prefer not to disclose			−36.496***
			(0.00000)
Race category: Hispanic or Latino			2.485***
			(0.832)
Race category: Black or African American			3.043**
			(1.429)
Race category: Asian			1.563***
			(0.566)
Race category: Native Hawaiian or Pac. Isl.			15.083***
			(0.00000)
Race category: Other			1.606
			(1.061)
Race category: Prefer not to answer			19.349***
			(0.00000)
Role category: Graduate Student			−0.684
			(0.760)
Role category: Other			1.034
			(0.887)
Party category: Republican			1.074
			(1.180)
Party category: Independent			−0.154
			(0.486)

Party category: Other			−0.919
			(0.999)
Party category: Prefer not to say			−0.063
			(0.588)
Party category: Socialist			0.180
			(0.726)
Religion category: Protestant			−0.974
			(1.087)
Religion category: Other Christian			-2.774^{**}
			(1.250)
Religion category: Jewish			0.222
			(1.167)
Religion category: Hindu			-2.015^{*}
			(1.040)
Religion category: Muslim			−1.494
			(1.339)
Religion category: Buddhist			−1.614
			(1.179)
Religion category: Other			−0.017
			(1.655)
Religion category: Atheist or agnostic			-1.633^{**}
			(0.818)
Religion category: Prefer not to say			−1.380
			(1.024)
Treated:Pre-treatment Decision	−0.311	−0.654	−0.532
	(0.246)	(0.496)	(0.628)
Constant	1.488		
	(0.905)		
Robust Std Errors	Yes	No	No
Observations	165	165	165
R^2	0.082		
Adjusted R^2	0.065		
Residual Std. Error	0.908 (df = 161)		
F Statistic	4.811^{***} (df = 3; 161)		
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 4: Unabridged version of Table 2. Base gender: Cisgender Man. Base race: Non-Hispanic White. Base role: Undergraduate Student. Base Party: Democrat. Base Religion: Roman Catholic. We treat Pre-treatment Decision as an ordinal variable for sake of simplicity in analysis.

13.2.2 Unabridged version of Table 3

	Dependent variable:		
	MoreInclined		
	Model 1	Model 2	Model 3
	(1)	(2)	(3)
Treated	−0.145	4.371	2.193
	(0.321)	(3.074)	(3.189)
Pre-treatment Decision		1.242**	1.304***
		(0.550)	(0.497)
Gender category: Cisgender Woman			−0.346
			(0.472)
Gender category: Transgender Man			16.274
			(3,956.181)
Gender category: Non-binary			0.154
			(1.118)
Gender category: Other			15.721
			(2,182.086)
Gender category: Prefer not to disclose			−36.496
			(4,663.801)
Race category: Hispanic or Latino			2.575***
			(0.997)
Race category: Black or Afr. American			19.392
			(1,915.503)
Race category: Asian			1.451**
			(0.595)
Race category: Native Hawaiian or Pac. Isl.			15.083
			(3,956.181)
Race category: Other			1.277
			(1.191)
Race category: Prefer not to answer			19.349
			(3,956.180)
Role category: Graduate			−0.659
			(0.827)
Role category: Other			1.252
			(1.144)
Party category: Republican			2.114
			(1.583)
Party category: Independent			0.227
			(0.534)
Party category: Other			−0.015

			(1.059)
Party category: Prefer not to say			−0.071
			(0.630)
Party category: Socialist			0.209
			(0.802)
Religion category: Protestant			0.154
			(1.223)
Religion category: Other Christian			−2.982**
			(1.485)
Religion category: Jewish			0.653
			(1.220)
Religion category: Hindu			−1.708
			(1.075)
Religion category: Muslim			−0.694
			(1.574)
Religion category: Buddhist			−1.494
			(1.214)
Religion category: Other			−0.026
			(1.635)
Religion category: Atheist or agnostic			−1.243
			(0.844)
Religion category: Prefer not to say			−1.013
			(1.053)
Treated:Pre-treatment Decision		−1.006	−0.581
		(0.668)	(0.691)
Constant	0.425*	−5.159**	−5.436**
	(0.223)	(2.527)	(2.329)
Robust Std Errors	Yes	Yes	No
Observations	165	165	165
Log Likelihood	−111.703	−106.765	−87.301
Akaike Inf. Crit.	227.406	221.529	236.602
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 5: Unabridged version of Table 3. Base gender: cisgender man. Base race: no-Hispanic white. Base role: undergraduate student. Base party: Democrat. Base religion: Roman Catholic. We treat Pre-treatment Decision as an ordinal variable for sake of simplicity in analysis.