Humanities & Social Sciences Communications

ARTICLE

https://doi.org/10.1057/s41599-023-01692-0 OPEN

The heterogeneous effects of social support on the adoption of Facebook's vaccine profile frames feature

Nadav Rakocz[1✉], Sindhu Ernala[2], Israel Nir[2], Udi Weinsberg[2] & Amit Bahl[2✉]

Achieving widespread COVID-19 vaccine acceptance is a key step to global recovery from the pandemic, but hesitancy towards vaccination remains a major challenge. Social proof, where a person's attitude towards vaccination is influenced by their belief in the attitudes of their social network, has been shown to be effective for making in-roads upon hesitancy. However, it is not easy to know the attitudes of one's network, nor reliably signal one's own feelings towards COVID-19 vaccines, minimizing the impact of the social influence channel. To address this issue, Facebook launched a feature that enables users to overlay a message indicating that they support vaccination upon their profile picture. To raise awareness of these vaccine profile frames (VPFs), users received a variety of promotional messages from Facebook, a subset of which contained the social context of friends who had already adopted the frame. Leveraging this variation in promotional messaging, we analyzed the adoption pattern of VPFs in the US to determine the most effective strategies to drive VPF usage. We found that adoption is driven by a pattern of complex diffusion, where multiple exposures to the adoption decisions of others increased VPF uptake, and that there is substantial heterogeneity in the adoption response associated with prior vaccine beliefs, whether the promotion had a social component and closeness of the tie included in social promotions. Specifically, we observed resistance to adoption correlated with an aversion to follow authoritative health pages and stronger adoption effects from social promotions containing close friends. We also confirmed this finding of the value of strong ties through a randomized field experiment and heterogeneous treatment effects modeling. In contrast to studies that have shown the importance of less close relationships in vaccine decision-making, we found little effect from awareness of VPF adoption by weak ties. Finally, we detected no significant backfire effect for expressing support for COVID-19 vaccines via VPFs. Together, these results suggest that social proof provided by close friends may be a key driver for messaging campaigns intended to drive prosocial behavior such as publicly promoting vaccination and that these campaigns do not necessarily come with adverse experiences for adopters, even for a polarizing issue such as vaccines.

[1] University of California Los Angeles (UCLA), Los Angeles, CA, USA. [2] Core Data Science, Meta Platforms, Menlo Park, CA, USA. ✉email: nrakocz@ucla.edu; abahl@meta.com

Introduction

Widespread acceptance of COVID-19 vaccines is essential for achieving the coverage required for herd immunity, but in many countries, a sufficiently large proportion of people are still hesitant about receiving available vaccines (Lazarus et al., 2021; Solís Arce et al., 2021). Reaching sufficient vaccine coverage has been challenging due to barriers at multiple levels. One commonly used classification system describing these barriers is the 4Cs model which segments people based on the main driver for hesitancy: confidence (lack of trust in health institutions and pharmaceutical interventions), convenience (structural barriers preventing vaccination conversion despite intent), complacency (low perception of disease risk), and calculation (significant information searching) (Betsch et al., 2015). Furthermore, as people make vaccine decisions, Social Contagion Theory suggests that social influence also plays a role as these considerations are influenced by belief in the decisions of others (Christakis and Fowler, 2013). Recent studies have shown that such social influence can have a substantial effect on eventual vaccine decision-making, with positive associations found between acceptance and beliefs about the intentions of others to vaccinate (Brunson, 2013; Brewer et al., 2017; Agranov et al., 2021; Konstantinou et al., 2021; Moehring et al., 2021). These associations are amplified when the others in question are close, trusted ties from a person's social network (Goldberg et al., 2020; Lau et al., 2022; Rabb et al., 2022) In order to activate this social influence channel, people need to have accurate information about the beliefs of their social network, yet it's currently unclear to what extent people are aware of the vaccine decisions of others in their social network and may misestimate the degree of acceptance/resistance based on the amplification of a relatively small number of voices. For example, being exposed online to amplified messages of concerns regarding vaccine safety could decrease confidence and move calculations toward hesitancy (Loomba et al., 2021). On the other hand, positive indications that trusted ties have chosen to vaccinate can combat this phenomenon and result in increased confidence and adjust factors such as calculation and complacency towards intent (Konstantinou et al., 2021). In general, we do know that people underestimate others' adherence to a range of COVID-19 preventative behaviors (Graupensperger et al., 2021), biasing their perception of social norms towards non-compliance.

In order to make people more aware of the vaccine perceptions of their network connections, Facebook, in partnership with public health agencies, recently launched vaccination profile frames (VPFs) to enable users to surround their profile picture with a supportive message with respect to vaccination (Meta, 2021). This form of advertising one's support for vaccination is the raw material that may allow social influence to make progress on the 4Cs. Previous work has established the impact of social proof-driven behavior change on Facebook, in non-health-related areas such as voting (Bond et al., 2012), friending (Sun and Taylor, 2020), and activism for social issues (State and Adamic, 2015). However, little is known about the factors that drive social signaling of vaccination support on social media, their relative importance, their overlap with factors that drive vaccine decision-making more broadly, and whether there are any negative downstream effects of sharing one's support. In this study, we explore these issues in the context of VPF usage on Facebook.

Our first research objective (RQ1) seeks the factors that promote VPF adoption, with a particular focus on determinants related to exposure to the adoption decisions of a user's friend network. The VPF feature rollout, coupled with our knowledge of the overall Facebook social graph and user demographics, provides the variation and controls which enables us to address RQ1 quantitatively along a number of key dimensions. Specifically, (1) promotions for VPFs for the frames took on several forms, including those with/without the social context of friend adoptions, allowing us to observe the effects of social proof, (2) among promotions with social context, friends were selected at random to produce a mix of relationships, enabling a study of tie strength effect, and (3) VPF promotions were held back from a set of random users, giving us an interventional setting to validate our main findings. Together, these factors allow us to determine the effects and heterogeneity of social context on VPF adoption using both observational and experimental data.

Our second research question (RQ2) addresses the potential that while promoting vaccine beliefs may lead to improved public health outcomes in aggregate via social influence, expressing support for a polarizing issue such as COVID-19 vaccines may also come with social risks and unwanted negative social interactions (Oz, 2018; Schmidt et al., 2018). To address RQ2, we searched for a detectable backfire effect against those who adopted VPFs, where we operationalized this effect to be negative actions received on Facebook that limit social ties with a VPF adopter (unfriending, unfollowing, blocking).

To our knowledge, this is the first large-scale quantitative look at how social context applies to people's decision to socially demonstrate their choice to vaccinate, a distinct and less studied behavior compared to vaccine acceptance about which much more is known, giving attention to both positive (RQ1) and negative (RQ2) outcomes. Our results have implications for understanding the determinants of vaccine-related social signaling which is crucial for maximizing the impact of the social influence channel and for the design of messaging campaigns aiming to drive health-related behavior change via social media.

Methods

VPF adoption and promotional exposure data (RQ1). This study was conducted using de-identified data logged by Facebook in the normal usage and launch of VPFs in accordance with Facebook's data use policy. The full dataset contained ~1 million users in the US who adopted a VPF (among a set of ~40 official VPFs available for the initial feature launch) within the analysis window starting 2021-04-25 and ending 2021-05-08. We also selected a large randomly sampled set of ~10.5 million US non-adopters who were active on FB and at least 18 years of age. All analyses to measure diffusion effects and modeling to uncover adoption drivers were conducted using samples drawn from these parent sets of users while preserving the adoption rate we observed in purely random samples from 2021-04-25 (0.64% adopters). A flowchart illustrating user selection is shown in Supplementary Fig. 1.

To serve all downstream analyses with these user samples, we also collected standard demographic and Facebook activity controls, exposure counts for different VPF promotions (2 with social context, 2 without), VPF adoption dates, the number of accredited health pages they followed at the start of the analysis window (a proxy for prior vaccine beliefs), how many of their friends adopted VPFs prior to the user's adoption, and available tie strengths for any social context displayed in promotions (see subsection "Additional data section" in the "Methods" section for details on health pages and tie strengths). Table 1 lists descriptive statistics for the data.

Among the promotions considered in this study are 2 non-social variants (a message in the user's feed or profile page to adopt VPFs), and 2 social versions (messages in the feed that show that a single friend or a set of 3 friends have adopted VPFs). Examples of these promotional messages are shown in Supplementary Fig. 2.

Table 1 Descriptive statistics of features used in the regression analysis.

N **= 10,822,480**					
Adopted the VPF:					
Yes = 69,508					
No = 10,752,972					
Numerical variables	**Mean**	**Std**	**Min**	**Median**	**Max**
Age (years)	43.24	16.69	18	41	99
FB age (days)	2988.96	1700.62	1	3582	6333
User's friend count	467.23	693.47	0	243	4987
Number of days that contained Facebook activity within the last four weeks	21.51	9.49	1	28	28
Number of friends that adopted the VPFs	4.82	9.69	0	2	797
VPF posts seen	0.50	1.23	0	0	104
User's "State vaccination rate" (%)	53.04	8.01	35.72	52.85	73.34
Binary variables	**Prevalence (%)**				
Newsfeed promotion seen	19.85				
Profile promotion seen	14.22				
Friend aggregation promotion seen	17.14				
VPF post seen	25.66				
VPF post from close friend seen	0.74				
VPF post from influencer seen	0.06				
VPF post from user with high friend count seen	2.23				
VPF adoptions	0.64				
Categorical variables	**Prevalence (%)**				
Highest education level					
Graduate school	4.79				
College	38.76				
High school	20.69				
Unknown	35.76				
Sex					
Male	53.73				
Female	46.27				

Influencers' matching and difference in differences (RQ1). To examine the effects of influencer adoption on the adoption changes of followers, we chose the most followed 105 Facebook pages in the US that had adopted the VPF between 04/01/2021 and 06/24/2021. For each influencer, we measured the percentage of followers that adopted the frame one week before and after the influencer's adoption date and then calculated the difference between the two measurements. As a comparison counterfactual value of this difference, we matched each VPF-adopting influencer to 10 similar non-adopting ones, based on follower count and pre-trained Facebook graph embeddings (minimum cosine distance to neighboring pages with similar follower count; see subsection "Additional data section" in the "Methods" section for details on embeddings), and looked at the difference in differences (DID) of their follower adoption patterns. For each VPF-adopting influencer, we measured the average DID with their 10 matches as a measurement of influence on the adoption rate of their followers.

To measure the statistical significance of the DID values, we generated *P*-values using a permutation method to approximate the null distribution. Specifically, we randomly permuted each influencer's 2-week follower adoption rate vector, breaking up any temporal effect that was driven by the influencer adopting a VPF. Therefore, DID calculations based on 10,000 iterations of such permuted vectors captured the null distribution, which we used to assign *P*-values to our observed DID values.

Logistic regression model (RQ1). We implemented logistic regression where VPF adoption was the dependent variable and exposures to the different promotional formats and whether one of these included a strong tie were the independent variables of interest. We also utilized a set of confounders as described in the text. A full list of variables can be found in Table 1.

To estimate model parameters in a robust manner, we used a bootstrapping procedure where we: (1) randomly sampled 1 million users from among VPF adopters and non-adopters, maintaining the adopter ratio observed in purely random samples from 2021-04-25 (0.65% adopters; Supplementary Fig. 1); (2) fitted a logistic regression model using the statsmodel package in Python (Seabold and Perktold, 2010) to produce maximum likelihood estimates of model coefficients; (3) measured the model's performance using the area under the receiver operating characteristic curve (AUROC) on a randomly held-out sample; (4) repeated steps 1, 2, and 3 for 1000 iterations to produce a mean estimate along with a 95% confidence interval of model parameters and AUROCs.

Randomized field experiment (RQ1). To determine the causal effects of social promotions, interventional data between 06/18/2021 to 07/18/2021 were collected. The experimental design was a simple A/B test where treatment was defined to be delivery of the friend aggregation post (Supplementary Fig. 2), and the control condition was not receiving this promotion. Eligibility for inclusion in the experiment was based on being a non-adopter at the start of the experiment, age (≥18 years old), location (US-based user), not having received a friend aggregation promotion within 2 weeks of the start date, and having at least three friends who had already adopted the VPF. Approximately 645K users met this eligibility condition, with roughly 323K randomly chosen for treatment and 321K as controls. The experiment outcome was the adoption of a VPF before the experiment's end date.

Table 2 lists descriptive statistics for the experimental data, showing the covariate balance between treatment conditions.

Causal Forest for heterogeneous treatment effects (RQ1). To search for heterogeneous treatment effects in the field experiment, we used causal forests which leverage the random forest algorithm to find sub-groups on which the conditional average treatment effect is maximized. We utilized the causalforest

Table 2 Descriptive statistics of features used in the randomized field experiment analysis.

Experimental data $N = 644{,}231$	**Cases** $N = 321{,}438$		**Controls** $N = 322{,}793$	
Numerical variables	**Mean**	**Std**	**Mean**	**Std**
Age (years)	40.72	14.74	40.73	14.71
FB age (days)	3642.04	1425.84	3640.85	1428.11
User's friend count	810.91	956.64	810.32	955.72
Number of days that contained Facebook activity within the last four weeks	26.81	3.55	26.82	3.54
Number of friends that adopted the VPFs	9.51	13.69	9.47	13.68
AH pages followed	2.69	8.82	2.68	9.52
User's "State vaccination rate" (%)	52.37	8.03	52.41	8.03
Binary variables	**Prevalence (%)**			
Non-social promotion seen	41.4		42.2	
Friend aggregation promotion seen	31.2		0	
VPF post seen	24.9		24.8	
VPF adoptions	0.37		0.22	
Categorical variables	**Prevalence (%)**			
Highest education level				
Graduate school	5.3		5.2	
College	50.2		50.1	
High school	24.3		24.5	
Unknown	20.2		20.2	
Sex				
Male	59.81		59.68	
Female	40.19		40.32	
Ties with prior friend adopters				
Weak	48.1		48.3	
Medium	32.6		32.5	
Strong	19.3		19.2	

package in R for model fitting and feature importance metrics from the trained model to score the contribution to an effect size of each included covariate, which encompassed the same set as the predictive model described above. In general, covariates that were used more often and in earlier stages of tree building are provided with higher importance scores.

Backfire effects (RQ2). To examine the association between VPF adoption and backfire effects, we collected data on select negative actions taken against adopters by other users in the two weeks preceding and succeeding VPF adoption. Specifically, we selected actions that sever the relationship or inhibit information flow between a user pair: unfriending, blocking, and unfollowing. We included ~475K users in our study who adopted a vaccine profile frame between 2021-04-14 and 2021-04-18 and received at least one negative action against them in the 2 weeks surrounding VPF adoption. As controls, we sampled ~507K non-adopters (from the same time period) identified by the criteria described above for RQ1.

We conducted a stratified propensity score analysis between adopters and control users who did not adopt the VPF. As covariates, we included demographic variables (age, gender), friend count, account tenure, number of vaccine profile frames seen, and accredited health (AH) pages followed by the user. To identify treatment (VPF adopters) and control users who are statistically similar to one another along the covariates, we match individuals with similar propensity scores into strata. Each stratum, then, consists of matched treatment and control users and lets us estimate the effect of VPF adoption on backfire effects within each stratum. To compute the propensity scores, we built a logistic regression model (accuracy = 0.87, F1 score = 0.84) with the above covariates to predict one's likelihood to receive the treatment (VPF adoption). Then, based on the empirical distribution of propensity scores, our stratified matching approach groups treatment and control users with similar propensity scores into 5 strata. Table 3 shows the balance in covariates per strata. Lastly, we compute the average treatment effect per stratum with the outcome as a relative difference in negative actions targeted before and after adoption. Weighing the average treatment effect per stratum with the number of treated users in that strata gives the final average treatment effect on the treated.

Additional data

Tie strength. We used an internal scoring of edges in the Facebook friend graph which gives higher weights to pairs of users who interact more often and more directly. These scores were clustered into five non-overlapping intervals representing tie strength buckets. Scores in the highest bucket were annotated as strong ties, those in the next two were medium, and the ties in the lowest two buckets were considered weak. In general, this annotation scheme tends to place close friends and family into the strong tie bucket.

Page embeddings. We utilized internal pre-trained Facebook page-dense embeddings which embed nodes in the page–page graph where edges are determined by followers/fans, links posted, topics discussed, and other features (Slide 37 from (Facebook, no date). These vectors place nodes such that a low cosine distance from a query node gives pages that are most similar. In terms of our application to influencers, they return similar celebrities.

Accredited Health pages. Facebook has annotated pages from global and US health organizations such as the WHO, UNICEF, and CDC, as well as local/regional sources of trusted health information to be disseminated on various platform surfaces (Meta, 2020). We leveraged this list of pages and utilized their follows as a proxy for pre-existing vaccine beliefs, the assumption being that users with strong negative views towards vaccination will likely not follow such pages.

Table 3 Descriptive statistics of features used in the backfire effect analysis.

Covariates	Adopters		Controls	
	Strata 1: 29,307		**Strata 1:** 344,205	
	Strata 2: 37,854		**Strata 2:** 93,479	
	Strata 3: 37,191		**Strata 3:** 33,172	
	Strata 4: 65,034		**Strata 4:** 18,761	
	Strata 5: 305,264		**Strata 5:** 18,015	
	Mean (per strata)	**Std (per strata)**	**Mean (per strata)**	**Std (per strata)**
Age (years)	35.66	8.16	30.76	7.90
	53.97	9.41	54.12	8.02
	38.09	17.40	48.49	21.36
	45.67	8.06	47.37	12.29
	53.52	13.86	51.17	14.26
FB-age (log)	3.29	0.55	3.29	0.51
	3.27	0.59	3.25	0.61
	3.34	0.51	3.33	0.54
	3.33	0.55	3.39	0.48
	3.43	0.45	3.45	0.42
User's friend count	1422.46	1406.76	1618.50	1408.17
	1169.54	1296.80	1052.00	1228.65
	1471.66	1427.36	1201.73	1286.96
	1438.88	1417.13	1312.49	1268.47
	1229.19	1312.69	1429.05	1341.87
Number of VPFs seen	0	0	0	0
	0.03	0.18	0.01	0.09
	0.77	0.42	0.52	0.50
	0.98	0.14	0.90	0.29
	3.25	2.96	1.92	1.30
AH pages followed	4.68	6.99	2.64	4.79
	7.91	12.02	6.26	11.45
	6.52	12.60	7.01	16.04
	7.95	13.34	9.38	19.97
	13.47	28.94	16.18	50.37

Vaccine data. We downloaded COVID-19 vaccination rate data at county granularity from the CDC website (https://data.cdc.gov/Vaccinations/COVID-19-Vaccinations-in-the-United-States-County/8xkx-amqh) on 06-01-21, which represents aggregate county vaccination coverage up until that date. For Fig. 2, we used the county-level metrics for the percent covered with 1+ dose among the 18+ population.

Results

VPF adoption exhibits a pattern of complex diffusion. Signaling one's support for a polarizing issue such as vaccination comes with social risks, and may require the psychological support of first seeing friends adopt this behavior for many users to do so themselves. If so, there are downstream questions as to whether this form of influence follows a pattern of complex (increasing dose-like effect of multiple exposures) versus simple (no increasing dose-like effect) diffusion. To determine this association, we used bivariate analysis and measured the empirical adoption probability of users in our sampled data conditioned on the number of (1) their friends that had previously adopted a VPF and (2) their friend's VPFs adoption posts they saw in their News Feed (a post alerting a user's adoption is automatically generated and shared in their friend's Feeds).

Figure 1a shows that the probability of adopting a VPF rapidly increases as the number of friends that have previously adopted grows, starting at a baseline of very close to 0% when no friends have adopted and saturating at about 2% when ~40 friends have done so. However, users do not necessarily see the adoption posts for all of their friends, as they may not scroll far enough on their Feed, scroll past these posts without viewing them, or simply not be on Facebook at that time. Therefore, we also looked at this same adoption probability conditioned on adoption post impressions in Feed that were reliably seen by the user. These results in Fig. 1b show a quicker saturation effect, at ~7 exposures leading to a ~4% adoption rate. Together, these results are supportive of social proof playing a role in VCP adoption, via a complex diffusion with saturation after 7 exposures on average, and requiring upwards of 40 friend adoptions for exposures to reach this level.

Pre-existing openness to vaccines requires significantly less social proof for adoption. Not all recipients of social proof are alike, and while Fig. 1 established a population-level association between social proof and VPF adoption, further segmentation of the data reveals significant adoption heterogeneity related to existing vaccination attitudes. Specifically, we divided users based on different noisy proxies for overall vaccination attitudes to evaluate the differential impact of social exposure as a function of rising openness towards vaccination. The two proxies we utilized were (1) profile county location thresholded to divide users between high/low COVID-19 vaccination rate counties (top and bottom 25 percent quartiles) and (2) the binned number of high-quality health pages followed by the user (see the "Methods" section for details on these pages) (Frey, 1986; Yom-Tov and Fernandez-Luque, 2014). In both cases, we observed that substantially less social proof is required to reach comparable adoption rates as we move up in the levels of these proxy variables.

For example, Fig. 2a shows that as users follow more high-quality health pages from trusted health authorities, the effect of six exposures to VPF social proof increases the adoption rates by 26% (95% confidence intervals of ±19%) when comparing users who follow 10+ high-quality health pages versus those who follow none. Figure 2b shows a 47% increase (95% confidence interval of ±56%) at 3 exposures when we utilize a location-based attitudinal proxy based on the specified home county of the user (note that at higher exposures the statistical significance of the difference disappears). Overall, these different vaccine attitude proxies highlight that substantially more social proof is required to drive comparable VPF adoption when there is existing resistance toward vaccination.

Social proof from stronger ties has a greater effect on adoption. Having shown heterogeneity in adoption response from the point of view of the recipient of the social proof, we next examined differential response when the friend providing the proof are close/far ties. Facebook users with frequent interactions on the platform, such as close friends and family, generally have a higher interpersonal influence on each other than user pairs who rarely interact or where the interactions are only one way (e.g. following a celebrity) (Aral and Walker, 2014). Figure 3 shows that social proof from strong ties indeed leads to a stronger likelihood to adopt the VPF compared to weak ties, saturating at 5% versus 2% when ~40 friends have adopted, and at 6% versus 3% after being exposed to ~7 VPF posts.

Influencers showed limited effect on adoption. While social proof from close ties was overall more influential, we also looked at a subset of weak ties that are of particular interest as campaign messengers, the social influencer. To do so, we examined the most followed 105 influencers in our data set and compared the adoption rates of VPFs among their followers

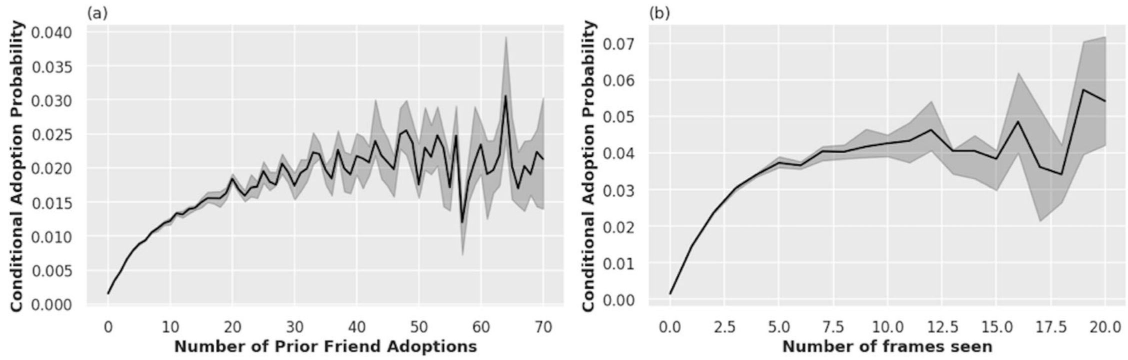

Fig. 1 VPF adoption probability conditioned on social support. The probability of adopting a VPF is conditioned on **a** number of friends who have adopted and **b** the number of friend's adoption posts seen. A pattern of complex diffusion is evident, in which as the number of social proof exposures increases, so does the likelihood of the user adopting the frame.

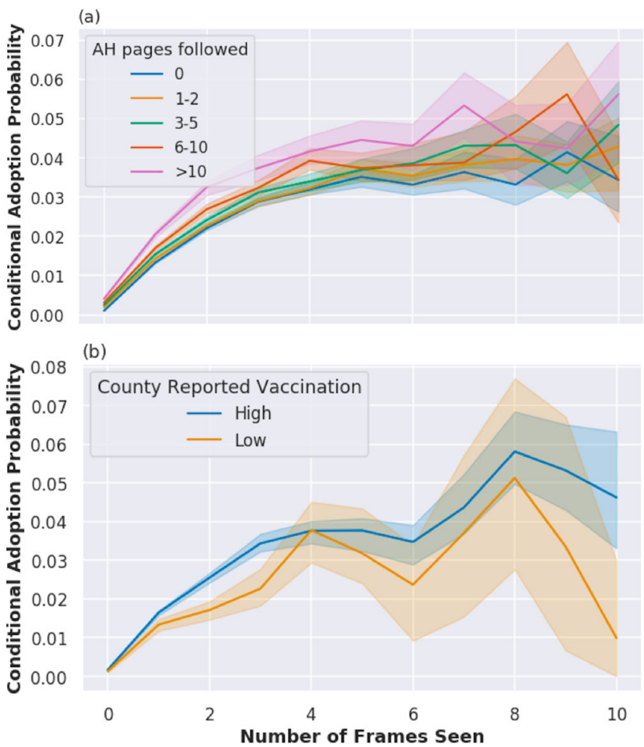

Fig. 2 VPF adoption probability conditioned on social support and segmented by pre-existing vaccine attitudes. The probability of adopting a VPF, conditioned on the number of friend's adoption posts seen, and segmented by **a** Authoritative health (AH) pages followed by the users, and **b** the COVID-19 vaccination rate in the user's home county (binned by top and bottom 25 percent quartiles). These cuts provide proxies for pre-existing vaccine attitudes and show that significantly less social proof is required to reach comparable adoption rates as we move up in the levels (representing more openness to vaccination, in aggregate).

before and after the influencers adopted the frame themselves. For a comparison control value of this difference, we matched each VPF-adopting influencer to 10 similar non-adopting ones, based on follower count and graph embeddings (see the "Methods" section for details), and looked at the changes in their follower's adoption patterns across the same time period. Figure 4a shows the difference in differences (DID) set up for our analysis for a particular adopting influencer and a matched control from our data.

Figure 4b shows the average DID between the influencers and their matches, as well as the permutation test-based *P*-values (see the "Methods" section for details). Only 5% of the 105 influencers produced a significant DID effect on the VPF adoption rates of their followers, suggesting that feature adoption was not primarily driven by social influencers, and sharpening the importance of strong ties.

Modeling the effects of social proof on adoption. Having demonstrated significant effects and heterogeneity of social proof in isolated bivariate comparisons, we next moved to model VPF adoption as a function of these and other confounding variables in order to estimate the contributions of the different promotional formats. To do so, we implemented a logistic regression where VPF adoption is the dependent variable, exposures to the different promotional formats and whether one of these included a strong tie are the independent variables of interest, and the set of confounders included prior beliefs, general Facebook activity, number of friend VPF adoptions, and user age/gender/location.

The performance of the model as measured by the area under the receiver operator characteristic curve (AUROC) using cross-validation was 0.87 [95% CI 0.87,0.87], indicating that the logistic model is a good choice to describe the relationship between the outcome and the dependent variables (Table 4). Figure 5 shows that discovery by social means has a significantly stronger effect when compared with the non-social promotion that appears on a user's profile page (OR = 6.18 and 95% CI of [5.46,6.88] for profile frame post; OR = 2.18[1.83,2.54] for friend aggregation post). The effect is even greater when comparing social discovery to non-social promotions appearing on a user's Newsfeed (OR = 10.50[9.10,11.77] for profile frame post; OR = 3.71[3.07,4.32] for friend aggregation post). The significant coefficient scores (log odds scores) for social discovery and the high OR values compared to non-social discovery highlight and quantify the value of providing social proof to drive VPF adoption.

In addition, when social proof from a strong tie (close friend or family) is provided in the social discovery, we estimate a further increase in the adoption odds score of 1.9 (holding everything else constant). Consistent with our findings of influencer effects in the previous section, we also found no significant effect when an influencer was included in the social proof. These results confirm the value of close ties in driving VPF adoption in a controlled model setting.

A randomized field experiment provides causal support for social proof and tie-strength effects on the adoption. While our modeling results point to strong social proof effects that are

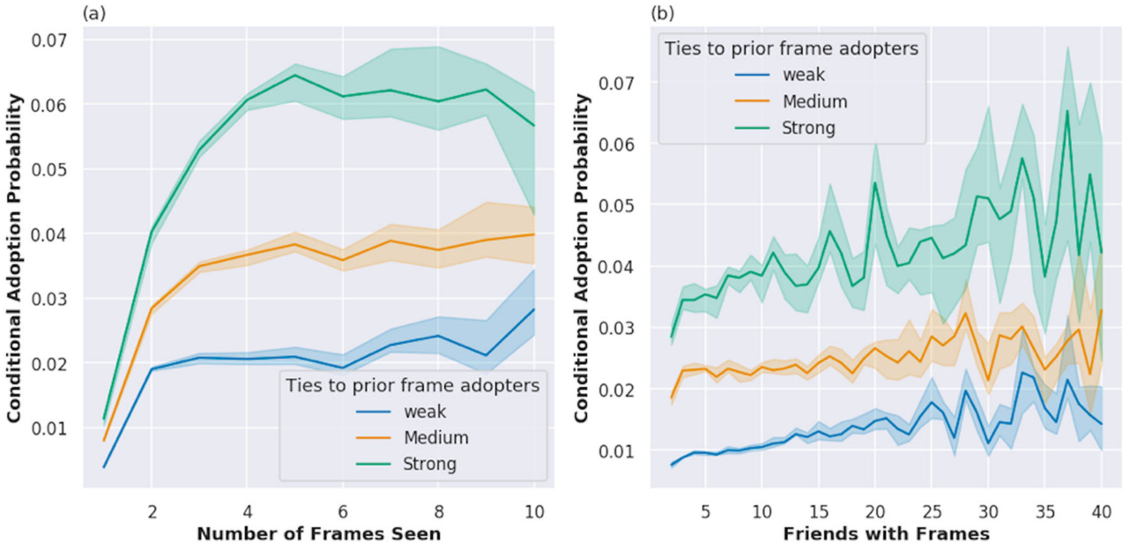

Fig. 3 VPF adoption probability conditioned on social support and segmented by tie strength. The probability of adopting a VPF segmented by the levels of tie strength with prior adopters and conditioned on the **a** number of friends who have adopted and **b** the number of friend's adoption posts seen. Users with strong ties to prior adopters seem to be more likely to adopt the VPF when social proof exposure increases compared with users that have weaker ties with prior adopters.

amplified when the source is a close tie, these estimates are observational without causal interpretation given that we cannot rule out the existence of uncontrolled/unobserved confounders. To provide some causal support for these main conclusions, we utilized an experiment that held out a random set of eligible US Facebook users from the social aggregation promotion. This control group did not receive the Newsfeed promotion shown in Fig. S2B, while the test group did so. Since eligible users (active 18+ Facebook users from the US with at least three friends that had previously adopted the VPF) were assigned to the test and control group completely at random, the conditional ignorability assumption holds and the causal effect of the VPF on eligible users can be estimated (Hernán and Robins, no date). The covariate balance across control and test groups is shown in Table 2, and the average treatment effect (ATE) on the treated for this promotion was a 0.15% (95% CI = [0.12%,0.18%]) increase in VPF adoption rates (Fig. 6, left column) which amount to a relative increase of 75% (95% CI = [56%,79%]) in VPF adoption. In cases where the adoption rate among the control is very small, the absolute effect size can also appear quite small, and so it is important to also consider the relative effect size which indicates the proportional increase caused by the treatment. Here, it is estimated that among eligible users the VPF would increase adoption of the frame by 75%. In addition, when considering the millions of users active on the Facebook platform in the US, an absolute effect size of 0.15% would result in a large increase in people who adopt the VPF.

In addition to ATE, we also attempted to estimate the conditional ATE (CATE) on the treated given an idiosyncrasy of this experiment: friends who adopted were selected for the social aggregation post at random, but the friend's identities were not retained in our logging. While the random selection created the variation to estimate a CATE of tie strength on adoption, the lack of friend identities led us to use each exposed users' maximum tie strength across all their friends who previously adopted for conditioning levels of tie strength (weak, moderate, strong). The rationale for choosing this conditioning scheme is that users with higher tie strengths on average will tend to see more social promotions from closer friends (in the aggregate). Figure 6 shows an increasing trend in CATE correlated with increasing levels of tie strength, with the strongest level significantly higher than the weakest one (CATE strong = 0.28%, CATE weak = 0.05%, difference = 0.23% (two-sided P-value = 8.4e−9).

Causal machine learning reveals additional heterogeneous treatment effects. Estimated exposure to different tie strength levels of social proof showed significant CATE differences in our pre-planned experiment. To search for other potential heterogeneous treatment effects (HTE), we applied the causal forest algorithm (Athey et al., 2019) to the experimental data and our full set of covariates from the modeling section, allowing us to rank covariates by their contribution to heterogeneity in adoption upon treatment with the social aggregation promotion (Fig. 7). This analysis confirmed that tie strength is a major HTE contributor, showing up second in the ordered feature importance scores. Lower in the list, we also observe adopter friend count, health pages followed, and state vaccination level as drivers of heterogeneity, showing that the data-driven HTE discovery approach confirms the observations in Figs. 1 and 2, where we see that these features drive heterogeneity, but to a lesser extent than tie strength.

The strongest HTE-driving feature was the age of the user, with older users (>50 years old) showing higher CATE (Fig. 8a). One factor that may be contributing to this age effect is tie strength. When looking at the proportion of friends who have adopted the VPF, segmented by levels of tie strength, we see that older users tend to have proportionally stronger tie friends (Fig. 8b). In particular, we see that once an age cohort has a proportion of close ties around 0.25 (Fig. 8c), the CATE effect becomes significantly and consistently different from 0. This is approximately the point at which it becomes probabilistically more likely than not to select at least one strong tie when choosing three friends at random for the aggregation post. There may be other explanations for the association between age and adoption, including differential risk perception and incentives to promote vaccination among older users, but these are outside the scope of our study.

Backfire effects of VPF adoption. Having established a positive relationship between VPF adoption and exposure to a friend's VPF, we next moved to examine the potential negative side of exposing one's vaccine beliefs openly. Specifically, we examined

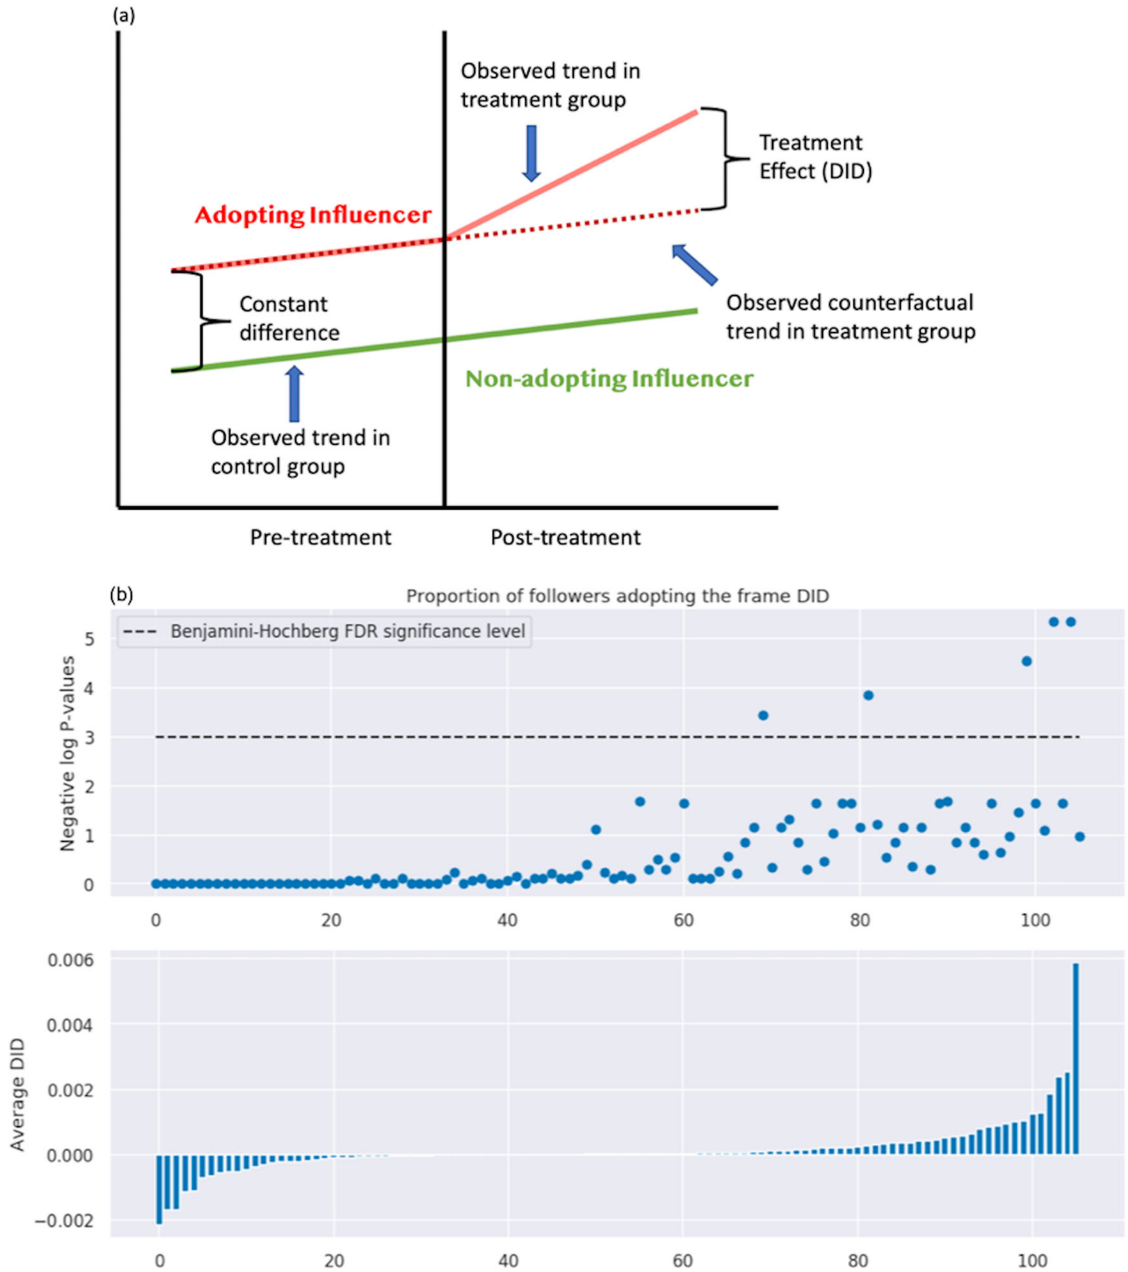

Fig. 4 The impact of influencers on VPF the adoption. a A difference in differences (DID) approach was taken to estimate the effect of an influencer's adoption decision on the decisions of their followers. In this illustrative example, we show an adopting influencer and a matched non-adopting control (in practice, we use 10 matched controls per influencer). We estimate the effect of the adopting influencer's decision on her followers by looking at the departure from the counterfactual provided by the non-adopting influencer's followers' behavior. **b** A permutation method enables deriving an empirical null distribution of DID values per influencer, allowing determination of *P*-values (adjusted for multiple hypotheses testing), and revealing that only ~5% of influencers show a significant effect at alpha = 0.05.

whether there are significant differences in targeted negative actions (unfollowing, unfriending, and blocking) upon VPF adoption for individual adopters in the 2 weeks before/after adoption.

Our findings revealed extremely low effects of VPF adoption on the number of negative actions received by adopters, suggesting that adoption did not have any significant backfire effects on individual adopters. The average treatment effect (ATE) on the treated was a 0.86 units increase in the relative difference in negative actions (Fig. 9a). This finding was robust to stratification of users by a propensity to adopt (see the "Methods" section for stratification details), with an average effect size across the strata having a Cohen's $d = 0.06$ (<0.2 suggests small differences between the two distributions; Fig. 9b–f).

In addition, across all strata, we do not observe any temporal variation in negative action trends after VPF adoption. Specifically, we observe a statistically insignificant peak (95% CI [1.176, 1.233]) in negative actions on the day of adoption, however, this flattens to the baseline value prior to adoption immediately.

Discussion

In this study, we show that social influence plays a significant role in increasing VPF adoption, an example of a health behavior change

Table 4 Regression.

Performance				
Model	**AUROC [95% Confidence Interval]**			
Model with a State Vaccination Rate feature	0.866 [0.866,0.867]			
Model with a State indicator	0.870 [0.870,0.870]			
Model with Demographics only	0.760 [0.759,0.762]			
Regression coefficients	**Coef**	**Std err**	***Z***	***P*** > | ***Z*** |
High friend count	−0.15	0.05	−2.92	3.50E−03
Influencer's post seen	−0.22	0.34	−0.65	5.16E−01
Strong ties post seen	0.65	0.06	11.5	1.25E−30
friend_agg only	1.47	0.08	18.23	2.94E−74
Frame post seen only	2.51	0.05	46.27	7.14E−293
Frame post and friend agg	2.56	0.06	40.68	7.14E−293
Profile Prompt QP	0.69	0.03	21.77	4.53E−105
Newsfeed QP	0.16	0.03	4.96	6.92E−07
Intercept	−5.17	0.13	−40.22	7.14E−293

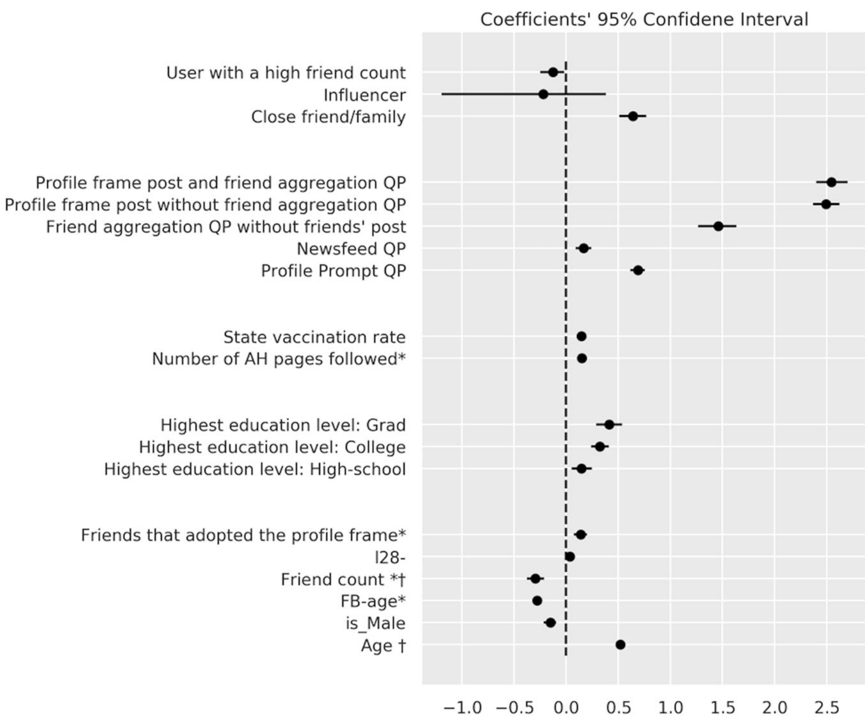

Fig. 5 Modeling VPF adoption with logistic regression. Logistic regression coefficients for adoption conditioned on the discovery channel (different QPs = "quick promotions"), tie strength of any included social proof, and a broad set of confounders, including those where heterogeneity in adoption response was observed. The coefficients (log odds scores) imply that VPF adoption is strongly affected by social aspects such as seeing a vaccination post from a close friend or seeing a promotion informing users that their friends have adopted the frame. FB-age is defined as the number of days since the user has signed up to the Facebook platform. "l28-" is an inactivity variable defined as the number of days within the last four weeks at which the user has not been active on the FB platform. (*) marks a log transformation. (+) marks a standardization transformation for zero mean and unit variance.

where users choose to advertise their support for vaccination to their social network (RQ1). We found evidence that adoption follows a complex diffusion process, where multiple instances of social proof increase the probability of adoption (Fig. 1), and that there is significant heterogeneity in this response associated with factors such as prior vaccine beliefs (Fig. 2), whether users become aware of the feature in a social context (Fig. 6), and tie strength when social context is provided (Fig. 3). In short, significantly more exposures are needed to achieve comparable adoption levels when the user holds more resistant views to vaccination, or when the exposures lack social context from strong ties. We also jointly modeled adoption using these factors, controlling for a variety of confounders, to arrive at estimates of relative contribution among the factors, revealing that social support from strong ties is the most influential factor in driving adoption (Fig. 5).

This observational result was validated using a field experiment where a promotional message presenting multiple friends who had already adopted was held back from a control group, confirming the value of strong ties in an interventional setting, and therefore giving this relationship a stronger interpretation than simple association (Figs. 6 and 7). This randomization also provides evidence that this effect is not due exclusively to homophily, which we expect to be comparable between treatment and control groups, but representative of social influence.

These results, for the distinct and much less understood behavior of advertising one's vaccine support, are consistent with the literature on factors that contribute to vaccine acceptance and add to the research in a number of ways. As in previous studies (Brunson, 2013; Brewer et al., 2017; Bruine de Bruin et al., 2019; Agranov et al., 2021; Konstantinou et al., 2021; Moehring et al., 2021), we find that social influence is a strong determinant that drives a vaccine-related decision. While some studies have shown demographics to be more important (Bruine de Bruin et al., 2019), we found that social influence is the strongest determinant. Also consistent with previous studies (Goldberg et al., 2020; Lau et al., 2022; Rabb et al., 2022), we found strong ties to be the most influential form of social influence and age to be the strongest demographic determinant.

With respect to weak ties, many studies have reinforced the notion of the "strength of weak ties" (Granovetter, 1973) in

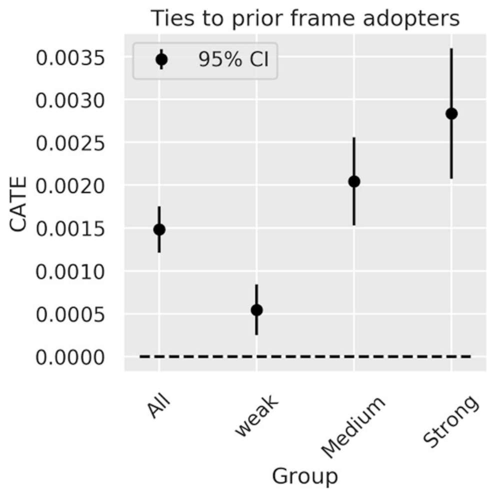

Fig. 6 Results from a randomized field experiment testing the effect of social proof on VPF adoption. Users in the treatment arm received the friend aggregation post as a means of social proof for VPFs. The three friends for this format were selected at random, enabling estimation of conditional average treatment effects conditioned on approximated tie strengths to the friends in the aggregation. The findings show an increasing trend in CATE correlated with increasing levels of tie strength.

diverse areas such as job searches (Rajkumar et al., 2022), scientific publications (Fronczak et al., 2022), novel information propagation (Bakshy et al., 2012), and many others. For vaccine acceptance, the results are mixed, with some studies showing that weak ties do matter (Moehring et al., 2021; Rabb et al., 2022), and others showing that they do not (Sinclair and Agerström, 2023). Our results support the latter conclusion, with very little effect found from the exposure to VPF adoption by weak ties. In addition, the value of influencers, which are generally weaker ties with high follower count, has not yet been established although many vaccine messaging campaigns utilize such celebrities (Ives, 2021; Lorenz, 2021). Our results show that these users are not an influential choice for providing VPF social proof (Fig. 4). One reason for these findings could be that for socially sensitive topics such as vaccines, the deeper affinity that a user has with their strong ties is a necessary precondition for being influenced to publicly disclose one's views. While weak ties and influencers may still hold value in providing novel social capital to influence downstream decision-making (Krämer et al., 2021), they by themselves do not seem to trigger the advertising of beliefs publicly for this socially sensitive issue.

Finally, our study is also novel in our ability to estimate the "dose effect" of social influence on a vaccine-related decision in our findings of complex diffusion dynamics. These dynamics have implications for public health messaging campaigns, motivating designs that plan for multiple exposures per user to saturate conversion rates. When the campaign aims to make in-roads with those lacking vaccine confidence, our results showing the heterogeneity that comes with prior beliefs suggest that far more exposures will be needed to reach comparable conversion rates. If the campaign is budget constrained and cannot reach such high levels of exposure, it may instead be a better use of resources to go after one of the other factors from the 4C model.

On the opposite side of positive behavior change (influencing frame adoption), it was possible that exposing one's vaccination beliefs via a VPF could also lead to unsolicited, negative reactions on Facebook. Despite vaccinations being a polarizing issue in the United States, we found no evidence of a backfire effect in which users exposed to their friend's adoption responded by limiting social ties (RQ2, Fig. 9). While there can still exist other forms of such an effect, the fact that we did not observe increases in aggregate unfriending, unfollowing, and blocking suggests that

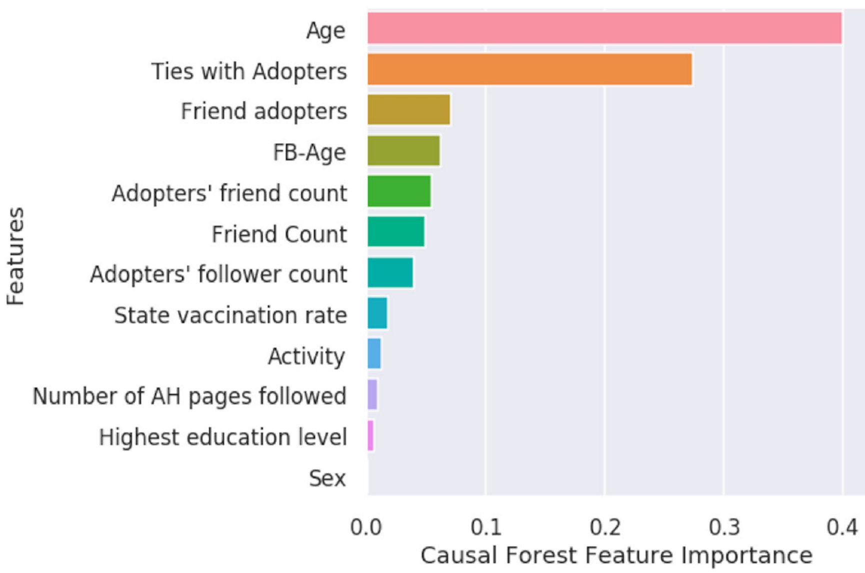

Fig. 7 Causal feature importance for HTEs. Each bar represents the importance of the associated feature in maximizing the heterogeneous treatment effect.

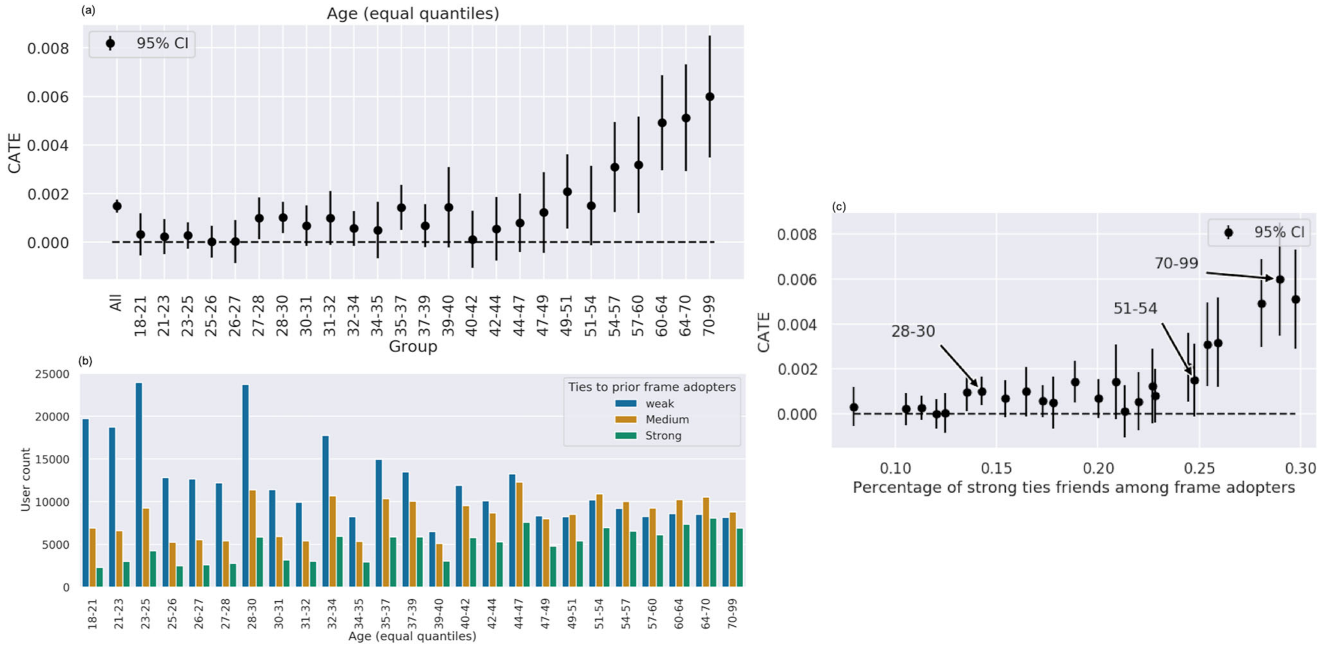

Fig. 8 The relationship between age and average treatment effect. Heterogeneous treatment effects grow with user age (**a**). This pattern is likely driven by strong ties, as older cohorts tend to have proportionally more stronger tie friends who have adopted (**b**). As the age cohort's strong tie friend proportion exceeds ~0.25, we see increasing CATE, significantly different from 0 (**c**; text annotation shows select cohort age ranges).

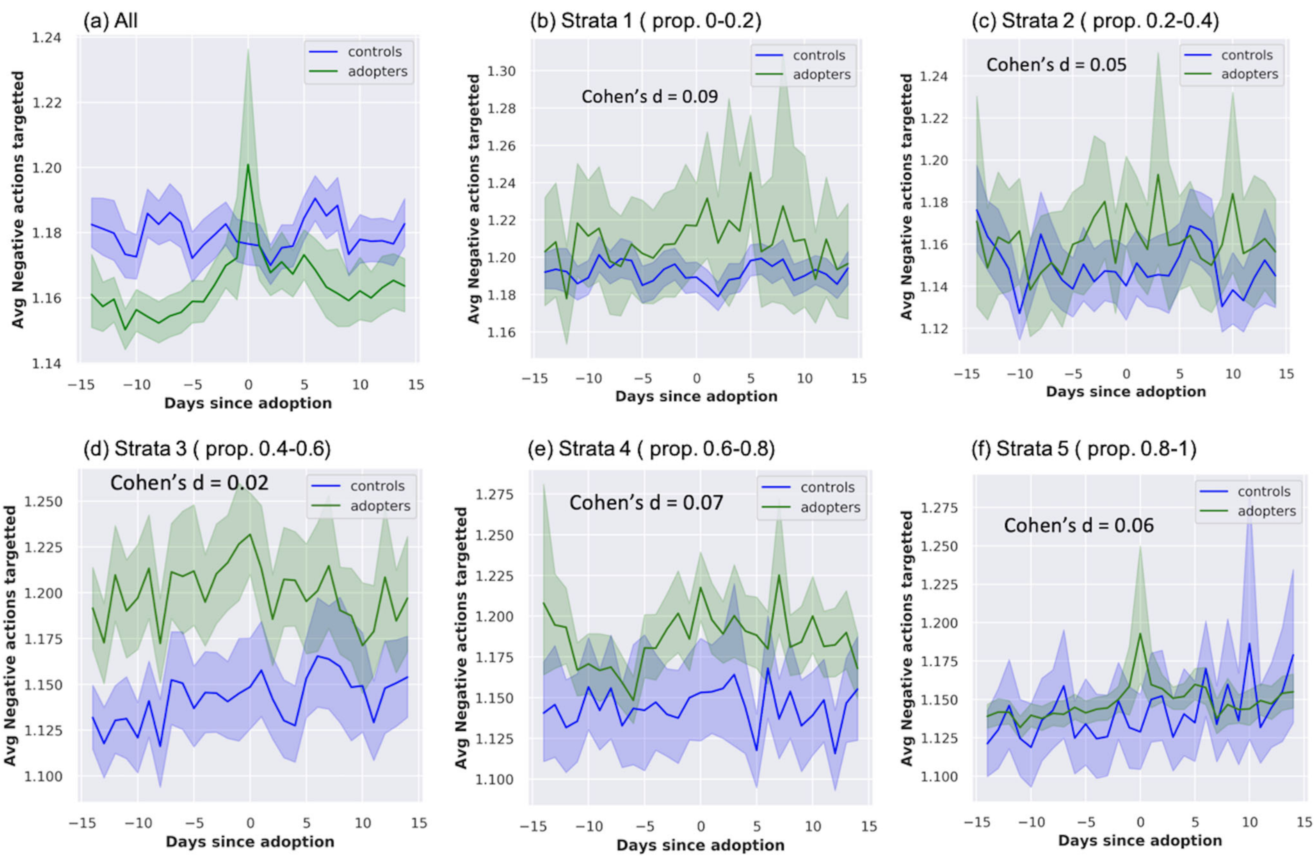

Fig. 9 Comparing backfire effects experienced by VPF adopters and controls in the two weeks preceding and succeeding VPF adoption. a Shows the average negative actions per day across all adopters and controls in the sample. **b**–**f** Show the average negative actions per day, across adopters and controls in each stratum matched based on the propensity score to receive treatment (VPF adoption).

campaigns in which identifiable social proof of vaccination is provided may not need to be overly concerned about large scale observable social contraction as an unintended downstream effect. Given the fact that VPF adoption awareness came largely from social promotions or friend posts on Facebook's News Feed, the absence of backfire effects cannot be attributed simply to the natural consequence of homophily or echo chambers as the Facebook friend graph has previously been shown to be cross-cutting with respect to user beliefs and interests (Goel et al., 2010; Lu and Lee, 2021).

This study has a number of limitations. First, we do not know to what extent these findings generalize to other health behaviors beyond VPF adoption. While it's unlikely that our findings idiosyncratically only apply to profile frames, and more likely that the learnings transfer to a variety of other health communications aimed at behavior change, it's unclear where the boundaries of generalization are. To define these, further work which varies the behavior, exposures, and tie strengths is needed.

Another limitation of our work is that we are not generally able to distinguish to what extent the mechanism of adoption is driven by homophily versus social influence (outside of having the field experiment which allows us to control for homophily via randomization in one promotional format). This differentiation has significant implications where, if homophily is the dominant driver, such campaigns are largely converting people who already hold open views towards vaccination and have come together to form ties on social media, but not necessarily making gains to influence those lacking confidence and in areas of the network which may be in most need for behavior change to drive public health objectives such as minimum thresholds for herd immunity.

The data available for this study also presented some additional limitations. We only looked at controlled exposures from a fixed set of promotional formats for which data was cleanly logged and available for analysis. Of course, users may see that a friend has adopted a VPF outside of these opportunities, such as when they are visiting a friend's profile page or via an organic friend post in their feed. This presents opportunities for "treatment" exposure that we were not able to detect, although we don't believe there is any systematic under-estimation that would skew our conclusions. We also did not control for all confounders in our observational data analysis, either because they were unknown or because we did not have proxies for all known ones, and therefore the results from the regression model cannot be interpreted causally. Building a complete dependency graph for variables (and their operational proxies) which may be influencing both exposure and adoption would allow us to more completely control for possible confounders, and bring the interpretation of the model coefficients closer to causality.

While the use of the field experiment did bridge this correlation/causality gap to some degree, we also note that the experiment design could have been improved to include multiple factors representing additional promotional formats and cohort properties, ties could have been chosen in the treatment arms more systematically to introduce controlled variation, and we could have selected additional endpoints to collect from users via pre/post treatment surveys to segment treatment effects by key pre-treatment variables and to estimate intent changes.

With respect to our backfire analysis, we looked for increases in specific events that limited direct connections on the Facebook social graph (unfriending, unfollowing, and blocking) upon adoption of an official VPF. We did not study other forms of negative social interactions, such as counter-speech in comments, negative reactions, adoption of anti-vaccination frames, or negative actions against non-adopters. Therefore, we cannot rule out these and other forms of backfire effects.

Finally, we note that VPF adoption is not the final endpoint of interest for public health purposes, and this study did not look at how increased adoption led to increases in intent or uptake.

Despite these opportunities for improving the study design in future work, our present results strengthen previous findings (based largely on small-scale surveys) that there is heightened value in positive vaccine messages containing social proof from close friends and family and that online delivery of such messages can help drive health-related behavior change at scale. We believe this result can help inform design choices made by policymakers and campaign designers to optimize public health communications. Overall, when there is the opportunity to deliver messages containing social support, and there is a choice in which ties to select, our results argue for including the social proof from the strongest ties possible to most effectively leverage the social influence causal channel (RQ1), and that providing this social proof does not result in social contraction as an unintended side effect (RQ2).

Data availability

This study was conducted using de-identified data logged by Facebook in the normal usage, testing, and launch of VPFs in accordance with Facebook's data use policy. The authors are Meta employees (Nadav Rakocz was a Meta research intern during this project). The data is not shared for public use. One exception to this is the vaccination rate data which was obtained from the CDC website.

Received: 2 May 2022; Accepted: 18 April 2023;
Published online: 07 July 2023

References

Agranov M, Elliott M, Ortoleva P (2021) The importance of social norms against strategic effects: the case of Covid-19 vaccine uptake. Econ Lett 206:109979. https://doi.org/10.1016/j.econlet.2021.109979

Aral Sinan, Walker Dylan (2014) Tie strength, embeddedness, and social influence: a large-scale networked experiment. Manag Sci 60:1352–1370. https://doi.org/10.1287/mnsc.2014.1936

Athey S, Tibshirani J, Wager S (2019) Generalized random forests. Ann Stat. https://doi.org/10.1214/18-aos1709

Bakshy E et al. (2012) The role of social networks in information diffusion. In: Proceedings of the 21st international conference on World Wide Web (WWW'12). Association for Computing Machinery, New York, NY, USA, pp. 519–528

Betsch C, Böhm R, Chapman GB (2015) Using behavioral insights to increase vaccination policy effectiveness. Policy Insights Behav Brain Sci 2(1):61–73. https://doi.org/10.1177/2372732215600716

Bond RM et al. (2012) A 61-million-person experiment in social influence and political mobilization. Nature 489(7415):295–298. https://doi.org/10.1038/nature11421

Brewer NT et al. (2017) Increasing vaccination: putting psychological science into action. Psychol Sci Public Interest 18(3):149–207. https://doi.org/10.1177/1529100618760521

Bruine de Bruin W et al. (2019) Reports of social circles' and own vaccination behavior: a national longitudinal survey. Health Psychol 38(11):975–983. https://doi.org/10.1037/hea0000771

Brunson EK (2013) The impact of social networks on parents' vaccination decisions. Pediatrics 131(5):e1397–e1404. https://doi.org/10.1542/peds.2012-2452

Christakis NA, Fowler JH (2013) Social contagion theory: examining dynamic social networks and human behavior. Stat Med 32(4):556–577. https://doi.org/10.1002/sim.5408

Facebook AI (no date) CS 4803/7643: deep learning guest lecture: embeddings and world2vec. https://www.cc.gatech.edu/classes/AY2020/cs7643_spring/slides/L13_Embedding_world2vec_final_version.pdf

Frey D (1986) Recent research on selective exposure to information. In: Berkowitz L (ed) Advances in experimental social psychology. Academic Press, pp. 41–80

Fronczak A, Mrowinski MJ, Fronczak P (2022) Scientific success from the perspective of the strength of weak ties. Sci Rep 12(1):5074. https://doi.org/10.1038/s41598-022-09118-8

Goel S, Mason W, Watts DJ (2010) Real and perceived attitude agreement in social networks. J Personal Soc Psychol 99(4):611–621. https://doi.org/10.1037/a0020697

Goldberg MH et al. (2020) Social norms motivate COVID-19 preventive behaviors. https://doi.org/10.31234/osf.io/9whp4

Granovetter MS (1973) The strength of weak ties. Am J Sociol 1360–1380. https://doi.org/10.1086/225469

Graupensperger S, Abdallah DA, Lee CM (2021) Social norms and vaccine uptake: College students'COVID vaccination intentions, attitudes, and estimated peer norms and comparisons with influenza vaccine. Vaccine 39(15):2060–2067. https://doi.org/10.1016/j.vaccine.2021.03.018

Hernán and Robins (no date) Selection bias. Causal inference: what if. Chapman and Hall, Boca Raton, FL

Ives M (2021) Celebrities are endorsing Covid vaccines. Does it help? N Y Times https://www.nytimes.com/2021/05/01/health/vaccinated-celebrities.html. Accessed 13 Feb 2023

Konstantinou P et al. (2021) Transmission of vaccination attitudes and uptake based on social contagion theory: a scoping review. Vaccines 9(6). https://doi.org/10.3390/vaccines9060607

Krämer NC, Sauer V, Ellison N (2021) The strength of weak ties revisited: further evidence of the role of strong ties in the provision of online social support. Soc Media+Soc 7(2):20563051211024958. https://doi.org/10.1177/20563051211024958

Lau BHP et al. (2022) Understanding the societal factors of vaccine acceptance and hesitancy: evidence from Hong Kong. Public Health 207:39–45. https://doi.org/10.1016/j.puhe.2022.03.013

Lazarus JV et al. (2021) A global survey of potential acceptance of a COVID-19 vaccine. Nat Med 27(2):225–228. https://doi.org/10.1038/s41591-020-1124-9

Loomba S et al. (2021) Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. Nat Hum Behav 5(3):337–348. https://doi.org/10.1038/s41562-021-01056-1

Lorenz T (2021) To fight vaccine lies, authorities recruit an "Influencer Army". N Y Times https://www.nytimes.com/2021/08/01/technology/vaccine-lies-influencer-army.html. Accessed 28 Sep 2021

Lu Y, Lee JK (2021) Determinants of cross-cutting discussion on Facebook: political interest, news consumption, and strong-tie heterogeneity. New Media Soc 23(1):175–192. https://doi.org/10.1177/1461444819899879

Meta (2020) Keeping people safe and informed about the Coronavirus. Meta. https://about.fb.com/news/2020/12/coronavirus/. Accessed 21 Jan 2022

Meta (2021) Encourage your friends to get a COVID-19 vaccine. Meta. https://about.fb.com/news/2021/04/encourage-your-friends-to-get-a-covid-19-vaccine/. Accessed 21 Jan 2022

Moehring A et al. (2021) Surfacing norms to increase vaccine acceptance. SSRN Electron J. https://doi.org/10.2139/ssrn.3782082

Oz M (2018) Discussing controversial issues on social media: examining the role of affordances, fear of isolation and de-individuation. Thesis. Available at: https://doi.org/10.26153/tsw/1274

Rabb N et al. (2022) The influence of social norms varies with "others" groups: evidence from COVID-19 vaccination intentions. Proc Natl Acad Sci USA 119(29):e2118770119. https://doi.org/10.1073/pnas.2118770119

Rajkumar K et al. (2022) A causal test of the strength of weak ties. Science 377(6612):1304–1310. https://doi.org/10.1126/science.abl4476

Schmidt AL et al. (2018) Polarization of the vaccination debate on Facebook. Vaccine 36(25):3606–3612. https://doi.org/10.1016/j.vaccine.2018.05.040

Seabold S, Perktold J (2010) Statsmodels: econometric and statistical modeling with Python. In S. van der Walt and J. Millman (eds) Proceedings of the 9th Python in science conference

Sinclair S, Agerström J (2023) Do social norms influence young people's willingness to take the COVID-19 vaccine? Health Commun 38(1):152–159. https://doi.org/10.1080/10410236.2021.1937832

Solís Arce JS et al. (2021) COVID-19 vaccine acceptance and hesitancy in low- and middle-income countries. Nat Med 27(8):1385–1394. https://doi.org/10.1038/s41591-021-01454-y

State B, Adamic L (2015) The diffusion of support in an online social movement: evidence from the adoption of equal-sign profile pictures. In: Proceedings of the 18th ACM conference on computer supported cooperative work & social computing (CSCW '15). Association for Computing Machinery, New York, NY, USA, pp. 1741–1750

Sun T, Taylor SJ (2020) Displaying things in common to encourage friendship formation: a large randomized field experiment. Quant Mark Econ 18(3):237–271. https://doi.org/10.1007/s11129-020-09224-9

Yom-Tov E, Fernandez-Luque L (2014) Information is in the eye of the beholder: seeking information on the MMR vaccine through an Internet search engine. In: AMIA… Annual symposium proceedings/AMIA symposium. AMIA Symposium. pp. 1238–1247

Author contributions

NR and AB led the study design, data analysis, and manuscript preparation. SE contributed to study design, data analysis, and manuscript preparation. IN and UW contributed to the study design. All authors read, commented, and approved the submitted manuscript.

Competing interests

This study was conducted by researchers at Meta. Otherwise, the authors declare no competing interests.

Ethical approval

This article does not contain any studies with human participants performed by any of the authors.

Informed consent

This article does not contain any studies with human participants performed by any of the authors.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1057/s41599-023-01692-0.

Correspondence and requests for materials should be addressed to Nadav Rakocz or Amit Bahl.