

RESEARCH

Open Access



Beyond binary: a machine-learning classification of childhood COVID-19 vaccination intentions using behavioural data

Chiara Chiavenna^{1,2,3*}, Laura P. Leone^{1,4,5}, Paolo Pin^{3,6}, Maria Cucciniello^{1,2,7,8} and Alessia Melegaro^{1,2,3}

Abstract

Background Immunisation is one of the most cost-effective tools for preventing infectious diseases. Yet, vaccine hesitancy, defined as a delayed acceptance or refusal of vaccination despite availability, has grown in recent years, threatening global public health efforts. This study investigates how socio-demographic and behavioural factors relate to willingness to vaccinate children against COVID-19, moving beyond binary pro-/anti-vaccine classifications to explore a more nuanced spectrum of intentions.

Methods Using a large-scale survey conducted in summer 2021 among 5,552 adults (2,041 parents and 3,511 non-parents) in Italy and the UK, we applied supervised machine learning models (XGBoost, Random Forest, and Multinomial Logistic Regression) to identify population segments based on their willingness to vaccinate children against COVID-19. We emphasise the importance of intention-based segmentation by distinguishing between “unwilling”, “undecided”, and “willing” respondents, a classification that better reflects the continuum of vaccination intentions.

Results Our findings, based on SHAP value analysis, show that friends’ opinion, the age of the child, and trust in vaccines are the strongest predictors of parental stances, with friends’ opinion emerging as the top factor across all models for parents. Overall, behavioural indicators played a key role in distinguishing between willingness groups.

Conclusions By integrating survey data with interpretable machine learning, this study highlights the importance of behavioural profiling and data collection for tailoring public health messages and targeting interventions to the most responsive segments of the population. While our empirical analysis is situated in the context of childhood COVID-19 vaccination, the framework has broader relevance for understanding parental decision-making and designing communication strategies in future vaccination campaigns.

Keywords Machine learning, Vaccination intention, Childhood vaccination, Parental decision, COVID-19, Human behaviour, Homophily

*Correspondence:

Chiara Chiavenna
chiara.chiavenna@unibocconi.it

¹DONDENA Centre for Research on Social Dynamics and Public Policy, Bocconi University, Via Sarfatti 25, Milan 20136, Italy

²Department of Social and Political Science, Bocconi University, Via Sarfatti 25, Milan 20136, Italy

³BIDSA Bocconi Institute for Data Science and Analytics, Bocconi University, Via Sarfatti 25, Milan 20136, Italy

⁴Department of Sociology, University of Oxford, Park End Street 42-43, Oxford OX1 1JD, United Kingdom

⁵Nuffield College, University of Oxford, New Rd, Oxford OX1 1NF, United Kingdom

⁶Department of Political Economics and Statistics, University of Siena, Piazza S. Francesco, 7-8, Siena 53100, Italy

⁷SDA Bocconi School of Management, Via Sarfatti 10, Milan 20136, Italy

⁸CERGAS Centre for Research on Health and Social Care Management, Via Sarfatti 10, Milan 20136, Italy



© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Introduction

Vaccinating children against infectious diseases is critical not only for individual protection but also to safeguard population health. While paediatric vaccination rates are high in high-income countries, global coverage remains suboptimal, particularly in regions vulnerable to epidemics [1]. In 2023, coverage for the diphtheria-tetanus-pertussis (DTP3) vaccine plateaued at 84%, and over 22 million children missed their first measles dose, contributing to more than 100,000 measles deaths, mostly among children under five [2]. These shortfalls highlight ongoing challenges in ensuring widespread vaccine uptake. The COVID-19 pandemic further heightened the urgency of paediatric immunisation. While children often experienced milder symptoms, their potential to transmit the virus—particularly through asymptomatic spread in schools and households—made their vaccination a crucial component of epidemic control [3, 4]. Yet, uptake remained limited: in the EU/EEA, only 26.7% of individuals under 18 received a COVID-19 vaccine [5]. While a global meta-analysis estimated that 60% of parents reported an intention to vaccinate their children and 26% were undecided [6], willingness appeared lower in European studies, where indecision was often substantial. These patterns highlight the importance of addressing parental indecision regarding vaccination intentions through effective public health communication.

To better understand the determinants of parental decision-making around childhood vaccination, an integrated approach is needed. Prior studies have shown that no single factor explains reluctance to vaccinate children; rather, it emerges from the interplay between behavioural, social, and demographic variables [7]. The COVID-19 pandemic further underscored the need to include psychosocial and behavioural dimensions in public health data collection and policy-making [8]. In this context, machine learning (ML) has emerged as a promising tool to analyse complex, high-dimensional data and to classify individuals into meaningful behavioural categories [9, 10]. By learning typical features of each intention group (e.g. unwilling vs. willing to vaccinate), ML models can identify patterns and predictors from a broad range of variables, including beliefs, risk perception, and social context. However, many existing applications have focused on binary outcomes (vaccinated vs. not vaccinated or willing vs. unwilling to vaccinate), often overlooking more nuanced, intermediate positions [11, 12]. Yet, undecided individuals—the so-called “fence-sitters”—are often the most persuadable segment [13]: compared to individuals with entrenched opposition, they are more likely to respond to accurate information, hence representing a key target for public health interventions [13–15]. Identifying and characterizing this

group requires models that go beyond binary classifications and explicitly capture heterogeneity in vaccination intentions.

Building on this framework, our study leveraged a large-scale survey conducted in Italy and the UK during the summer of 2021 to examine adults’ willingness to vaccinate children against COVID-19. In earlier work, we documented sharp divides in parental intentions, between those in favour of and those against vaccination, but also identified a sizeable segment of respondents expressing uncertainty or ambivalence [16]. While several studies have explored the predictors of outright vaccine refusal, few have systematically profiled this intermediate group [17–21]. Our dataset includes rich information on behavioural tendencies, vaccine-specific beliefs, social norms, trust in institutions, risk perception, moral values, and detailed demographics, allowing for a multi-dimensional analysis of vaccine decision-making. Hence, we extend this line of research by applying multiple ML models—eXtreme Gradient Boosting (XGBoost), Random Forests (RF), and Multinomial Logistic Regression (MNL)—to classify individuals into three meaningful intention-based categories: unwilling, undecided, and willing. Our approach emphasises the importance of modelling heterogeneity, and of distinguishing parents, the actual decision-makers, from non-parents, individuals who may hold opinions but can only enact them through peer influence. In addition, we examine how well individuals’ vaccine intentions can be predicted using only demographic variables typically available in administrative datasets, compared to models that include behavioural indicators. While existing evidence points to the importance of factors such as beliefs, social influence, and risk perception in shaping health outcomes [8], further empirical assessment of their predictive value is essential to substantiate the case for systematically collecting behavioural data as an integral part of future pandemic preparedness strategies.

Methods

Data

Study design

We conducted an online cross-sectional survey in Italy and the United Kingdom during June and July 2021. Participants were recruited through the professional survey company Lucid, which maintains large national online panels of adults (aged 18 or older). Quota sampling was applied within these panels to match national distributions by age, gender, education level, COVID-19 vaccination status and geographic region. At the time, eligibility to obtain a COVID-19 vaccine was limited to adults (18+) in the UK and extended to 16–17-year-olds in Italy. Although the European Medicines Agency had authorized vaccination for 12–15-year-olds in late May

2021, rollout for this group had only just begun in Italy and did not start in the UK until September 2021. To capture a broader range of childhood vaccination intentions, we targeted representativeness by adults' COVID-19 vaccination status, under the a priori assumption that adults delaying or hesitating about their own vaccination might also be more reticent about vaccinating children. While interim recruitment slightly overrepresented vaccinated respondents, Supplementary Table S1 shows that by the end of data collection the distribution of vaccinated and unvaccinated respondents was close to population coverage in both countries. The sample included two groups: (1) parents of children under 18 years, and (2) individuals without children or with children aged 19 or older, referred to collectively as “non-parents” for simplicity. After row-wise elimination of respondents with missing values, the final sample comprised 5,552 unique respondents, 2,041 parents and 3,511 non-parents, across both countries.

Outcome variables

The outcome variable was respondents' *willingness to vaccinate children against COVID-19*, measured on a 0–100 slider scale (0 = definitely not, 100 = definitely). This approach has been shown to preserve the full information of participants' responses compared to coarsened ordinal scales and is well suited to capturing the continuum of vaccine hesitancy [22]. Moreover, visual analogue sliders are widely used in health and psychology research and are often perceived as intuitive and engaging [23, 24], mitigating the potential cognitive burden of longer response times [25]. Questions were tailored to specific child age groups (0–2, 3–5, 5–11, 12–16, and 17–18 years). The question wording differed slightly by respondent type. Parents were asked: “*Suppose you could obtain a MHRA-approved COVID-19 vaccine for your children of age X. How likely are you to get a COVID-19 vaccine for your children of age X?*”. Non-parents were asked: “*Suppose some friends of yours could obtain a MHRA-approved COVID-19 vaccine for their children of age X, and that they asked you for advice. How likely are you to recommend them to vaccinate their children?*”. Parents could contribute multiple observations if they had more than one child, whereas non-parents responded to the question for all predefined age groups. As part of the survey design, participants were randomly assigned to one of three short information treatments (a neutral control, or a message highlighting either herd immunity or infection risk) immediately before answering this outcome question. A detailed description of these treatments and their effects, found to be small, is provided in a prior experimental study [16].

Predictor variables

We collected a wide range of predictor variables, grouped into two main categories: socio-demographic and health-related characteristics on one hand, and behavioural measures on the other.

Socio-demographic and health-related characteristics

We recorded comprehensive socio-demographic information, including age, gender, region of residence, residential setting (urban, suburban, town, rural, or other), marital status, educational attainment, employment and parental status. We also captured cultural and political background by assessing religious affiliation, frequency of religious service attendance, interest in politics, and self-placement on a left-right ideological scale. Health-related information comprised self-reported COVID-19 infection history, influenza vaccination history over the previous five years, and COVID-19 vaccination status at the time of the survey. A checklist asked about chronic medical conditions, including hypertension, diabetes, obesity, immunosuppression, and others.

Behavioural measures

We included a range of behavioural variables capturing both *individual beliefs* and *external influences* relevant to vaccine decision-making, mostly collected through Likert-type agreement scales. **Belief-related measures** included perceptions of vaccine safety, effectiveness, and necessity, concerns about vaccine overload and side effects, and beliefs about the potential consequences of personal vaccination behaviour for others. Additional items measured trust in medical professionals and the quality of the respondent's relationship with their healthcare provider. Emotional dimensions were assessed through separate items on worry about becoming infected with COVID-19 and about spreading it to others. Beliefs about who should have decision-making authority over childhood vaccination (e.g. parents vs. public institutions) were measured through a single-choice item.

External influences were assessed through measures of perceived norms and trusted information sources. **Perceived norms** included friends' attitudes toward vaccination (0–100 scale) and the proportion of people in the community wearing masks. **Trust** in institutions was measured by asking how often the national government could be trusted to do what is right. Participants also identified up to three actors they would trust most for COVID-19 vaccine information (e.g. health professionals, local authorities, media figures), and ranked seven sources of information used to make vaccination decisions in order of importance.

Several of these behavioural and attitudinal measures were adapted from validated constructs in the vaccine

hesitancy literature, including the 5C psychological antecedents of vaccination [26], institutional trust [27], and health behaviour models such as the Health Belief Model [28]. Some belief-related items were framed in general vaccine terms (e.g. perceptions of safety or necessity), whereas other measures (e.g. worry about COVID-19 infection, trusted actors for COVID-19 vaccine information) directly referenced the pandemic context. Initial data cleaning was performed using Stata/SE version 18.0. Descriptive statistics of all survey items included in the analysis are provided in Supplementary Information section S2.

Modelling approach

Model choice and training procedure

We applied and compared three predictive modelling approaches: XGBoost, RF and MNL. XGBoost and RF are tree-based ensemble algorithms suited for capturing complex, non-linear patterns [29, 30], while MNL offers a simpler, interpretable alternative that assumes a linear relationship between the predictors and the log-odds of each class. Each model was trained separately for parents and non-parents, using the same outcome specification. To assess out-of-sample performance, we implemented a stratified group-wise train-test split (80% training, 20% testing), stratified by outcome category and grouped by individual. This ensured that all observations from the same respondent were allocated to a single subset and that class proportions were preserved in the train and test sets. The training set was used for hyperparameter tuning and model fitting, and the test set for final performance evaluation.

Hyperparameter tuning

Each tree-based model was tuned via 5-fold cross-validation on the training set, using grid search to identify the Hyperparameter settings that optimized the macro-F1 score. As with the train-test split, in cross-validation we used a grouped stratified split, where stratification was performed on the outcome variable and grouping on individual identifiers. This procedure helps prevent overfitting and ensures that performance estimates are not

dependent on a specific data split. All machine learning analyses were conducted in R version 4.3.3 using the caret package version 6.0.94, which offers flexible functions to perform training and evaluation for different ML methods [31]. Final hyperparameter configurations selected for XGBoost and RF are reported in Table 1. Full tuning details and performance metrics on the training set for XGBoost, RF, and MNL are available in Supplementary Information sections S4.1, S12.1 and S13.1, respectively.

Model evaluation

After hyperparameters tuning, each model was retrained on the full training set and evaluated on the held-out test set. We assessed performance using Accuracy, macro-averaged Precision, Recall, F1 score, Matthews correlation coefficient (MCC), and Area Under the Receiver Operating Characteristic Curve (AUC). These metrics were selected for their relevance to multi-class classification and class imbalance: while Accuracy reflects the overall proportion of correct predictions, macro-averaged metrics assign equal weight to each class. Specifically, MCC has been shown to outperform both Accuracy and F1 in multi-class scenarios [32]. AUC evaluates the model's ranking ability and general discriminative power [33]. All macro-averaged metrics were computed using a one-vs-rest strategy.

Model interpretation

To interpret model predictions and identify key behavioural drivers of childhood COVID-19 vaccination willingness, we used SHapley Additive exPlanations (SHAP), a unified framework for interpreting machine learning models [34]. In a multi-class problem, for an observation x , the SHAP value $\phi_{j,c}(x)$ for feature j and class c quantifies the average change in the predicted probability of class c when the actual value x_j is taken into account, compared with the expected prediction if j were unknown. The sign of $\phi_{j,c}(x)$ reflects directionality: a positive value indicates that the observed value of j increased the probability of assignment to class c , whereas a negative value indicates that it decreased that probability. We computed approximate SHAP values for all observations from the tuned XGBoost models for parents and non-parents using a Monte Carlo permutation sampling approach [35], a model-agnostic method suitable for both tree-based and regression-based models, implemented in the R package fastshap (version 0.1.1). Following standard practice, we displayed the relationship between observed feature values and their corresponding SHAP contributions using beeswarm and dependence plots. We also summarized variable importance for each class by the mean absolute SHAP value $MAS_{j,c} = \mathbb{E}_x[|\phi_{j,c}(x)|]$, which captures the average

Table 1 Tuned hyperparameters of XGboost and RF models

Panel A: Parents

XGBoost nrounds = 200, max_depth = 9, eta = 0.05, gamma = 0.1, colsample_bytree = 0.7, min_child_weight = 1, subsample = 0.8

RF mtry = 18, ntree = 250, maxnodes = 360, nodesize = 10

Panel B: Non-parents

XGBoost nrounds = 400, max_depth = 9, eta = 0.1, gamma = 0.05, colsample_bytree = 0.6, min_child_weight = 1, subsample = 0.8

RF mtry = 30, ntree = 250, maxnodes = 2460, nodesize = 10

magnitude of these probability shifts across respondents, with larger values indicating greater overall influence of feature j on model predictions.

Administrative vs. behavioural models

To evaluate the feasibility of classifying individuals using only variables typically available in administrative datasets, and to assess the added value of behavioural profiling for targeted vaccine outreach, we trained two additional sub-models within both the parent and non-parent groups. The first model included only administrative or health record-type variables (e.g. age, chronic conditions, vaccination history), while the second relied solely on behavioural variables (e.g. attitudes, trust, beliefs). Since the administrative-only model contained 25 predictors, we selected the top 25 behavioural predictors from an initial full-variable model (which included 33 predictors), ranking them by their mean absolute SHAP value across outcome classes to ensure a fair comparison. Both sub-models were tuned and evaluated using the same procedure as the main models. The training set performance and a complete list of predictors included in the two sub-models is provided in Supplementary Information section S5.

Sensitivity analyses

To assess the robustness of our results, we conducted several sensitivity analyses.

Class rebalancing In our main analysis no re-balancing of cases in the three outcome classes was performed (the distribution of responses across classes is reported in the Results). Imbalanced target distributions can lead machine learning algorithms to favour the majority class: in a three-class setting, classification decisions are benchmarked against a default cut-off of 0.33. We therefore repeated model tuning using two strategies: (1) random up-sampling (with replacement) of minority classes so that they matched the size of the majority class and (2) random down-sampling of the majority classes to match the size of the least frequent class (Supplementary Information section S6).

Alternative outcome cut-offs As the discretization of the continuous outcome involved somewhat arbitrary thresholds, we repeated the analyses using alternative cut-off values (35–65, 30–70, and 20–80) to assess the consistency of model performance. Results are presented in Supplementary Information section S7.

Assessment within untreated group Because information treatments were presented immediately before the outcome question, we included treatment assignment as a covariate in all models. As an additional robustness

check, we assessed the SHAP ranking on the subsample of untreated respondents, i.e. those randomly assigned to the control information message. Predictor rankings in this restricted sample were largely consistent with the main analysis (Supplementary Information section S8).

Alternative variable importance measures We additionally computed other variable importance metrics alongside the main SHAP values analysis. Specifically, we obtained the Gain measure (XGBoost's internal importance metric) through the caret package [29]. Gain represents how much each variable improves the model's predictions across all decision trees. In addition, we computed permutation importance of variables using the vip package in R (version 0.4.1) [36]. This method measures the decrease in predictive performance when the values of a variable are randomly shuffled, thereby breaking its association with the outcome. The six top-ranked variables for parents and non-parents based on these methods are reported in Supplementary Information section S9.

Missing data imputation We examined the impact of missing data by imputing incomplete predictors' values using the k-nearest neighbours (k-NN) algorithm with Gower's distance. We then assessed the model performance and the stability of SHAP predictor rankings in the imputed dataset (Supplementary Information section S10).

Variable encoding and classification setup

For analysis, the continuous 0–100 outcome variable reflecting intentions about childhood COVID-19 vaccination was discretized into three categories, in line with our goal of profiling the “fence-sitters” class: we labelled the groups as unwilling [0, 25), undecided [25, 75] and willing (75, 100]. These thresholds were chosen to distinguish between strong opposition, ambivalence, and strong willingness, while maintaining roughly balanced group sizes. We assessed the robustness of model performance to alternative cut-offs; details are provided in the Sensitivity analyses paragraph below. Ordinal predictors were treated as numeric; categorical predictors with more than two levels were dummy-coded via full one-hot encoding, and binary predictors were represented using a single dummy. Two indicators (vaccine confidence and trust in doctors) were computed as the first principal component from a principal component analysis (PCA) on six and two items, respectively. From 52 original variables we obtained 89 single predictors after encoding. More details on variables' encoding can be found in Supplementary Information section S3.

Table 2 Macro performance of XGBoost, RF, and MNL models

Model	Accuracy	Precision	Recall	F1	MCC	AUC
Panel A: Parents						
XGBoost	0.664	0.659	0.659	0.658	0.490	0.846
RF	0.655	0.651	0.646	0.646	0.475	0.842
MNL	0.599	0.587	0.594	0.588	0.389	0.795
Panel B: Non-parents						
XGBoost	0.580	0.565	0.568	0.566	0.357	0.764
RF	0.585	0.567	0.570	0.565	0.362	0.767
MNL	0.571	0.544	0.553	0.542	0.335	0.749

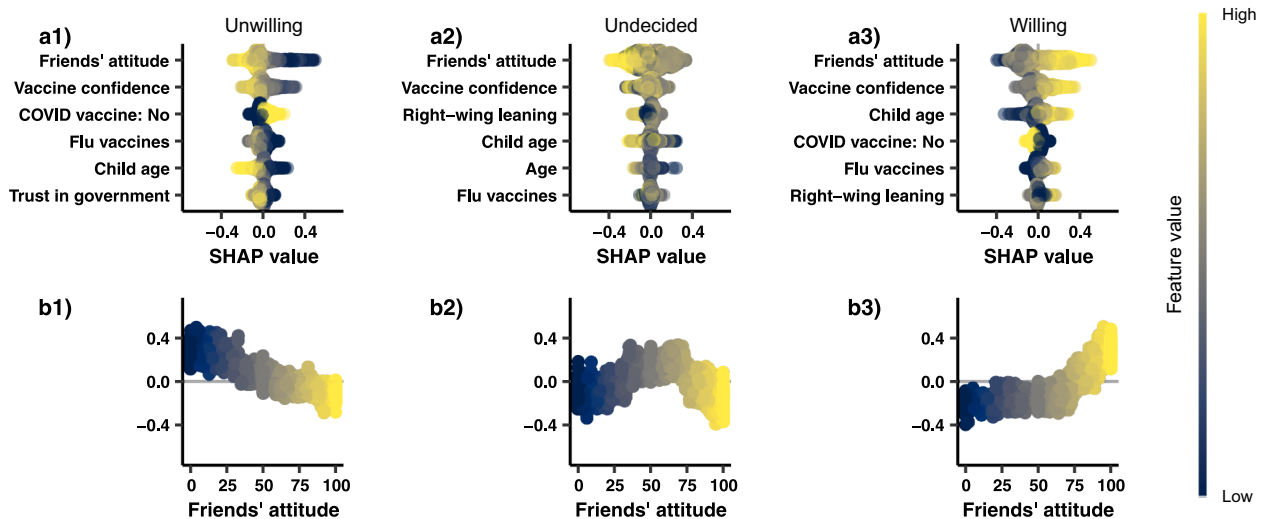


Fig. 1 Factors associated with intentions regarding childhood COVID-19 vaccination among parents. Top: Beeswarm plots of the six most important predictors for classifying unwilling (a1), undecided (a2) and willing (a3) according to SHAP values. Bottom: Dependence plots showing the relationship between friends' attitude toward vaccination and SHAP values for unwilling (b1), undecided (b2) and willing (b3)

Results

Sample characteristics and outcome distribution

We analysed 2,816 responses from the 2,041 parents and 17,902 from the 3,511 non-parents. The distribution across the unwilling, undecided, and willing categories was similar in the parent (29.4%, 31.9%, 38.8%) and non-parent (33.2%, 29.3%, 37.5%) samples. Sample characteristics by outcome group and parental status are provided in Supplementary Information section S2.

Model performance overview

Table 2 presents the performance of each classification algorithm (XGBoost, RF, and MNL) on the held-out test set for parents and non-parents. Each metric ranges from 0 to 1 (worst to best performance), except for the Matthews correlation coefficient (MCC), which ranges from -1 (total disagreement) to +1 (perfect prediction), with 0 indicating chance-level performance. In our results, the XGBoost model achieves AUC values above 0.84 in parents and 0.76 in non-parents, indicating strong discriminatory ability. Similarly, F1 scores and MCC values for XGBoost (F1 = 0.658/0.566; MCC = 0.490/0.357) demonstrate that our models achieve high predictive validity,

particularly when using XGBoost. By contrast, the MNL performs relatively worse, especially among parents (MCC = 0.389), underscoring the value of using non-linear tree-based methods for this classification task.

Key predictors and SHAP interpretation

We used SHAP values to examine predictor contributions across outcome classes, based on the best-performing XGBoost model. Supplementary Information section S11 reports the top predictors' mean absolute SHAP values. Figure 1, panels a1-a3, displays SHAP results for parents. Factors are ranked by mean absolute SHAP values, and the dots represent the distribution of individual effects. Friends' attitude toward vaccination was the most influential predictor across all classes, with average probability shifts of about 12.5 p.p. for the undecided, 16 p.p. for the willing, and 9 p.p. for the unwilling; these effects were 1.6–2.7 times larger than those of vaccine confidence, the second-ranked predictor. The direction of this effect varied (panels b1-b3): high values of friends' attitude (around 75–100), reflecting pro-vaccine social circles, increased the probability of willingness and reduced the probability of indecision or unwillingness. Intermediate

values (25–75) were linked to indecision, while low values (0–25) predicted unwillingness or ambivalence. General vaccine confidence showed a similar pattern: greater confidence increased the probability of willingness and decreased the likelihood of indecision or refusal. Child age also shaped parental intentions: older children were more likely to elicit willingness to vaccinate, while younger age groups were linked to indecision and, to a lesser extent, unwillingness. Its effect was smallest for the undecided (3 p.p.) and largest for the willing (5.4 p.p.). Parental uptake of seasonal flu vaccines was another strong behavioural marker: regular flu vaccination correlated with higher willingness to vaccinate children against COVID-19, whereas a negative flu vaccination history modestly increased indecision and strongly predicted unwillingness. Other predictors showed more class-specific effects. COVID vaccine refusal helped distinguish both the unwilling and the willing (a1, a3); low trust in government was more strongly associated with unwillingness (a1); and right-wing political orientation and respondent age had a complex, mixed effect, particularly in the undecided class (a2). The sixth-ranked predictors—trust in government (unwilling class), flu vaccination history (undecided), and right-wing political leaning (willing)—had mean SHAP values three to eight times smaller than that of friends' attitude, underscoring the latter's dominant role.

Among non-parents, SHAP values from the XGBoost model highlighted child age, friends' attitude, vaccine confidence and trust in doctors as the four most influential predictors across all outcome classes (Fig. 2, panels a1–a3). Child age was the strongest determinant, shifting

probabilities by about 14 p.p. for the unwilling, 9 p.p. for the undecided, and 18 p.p. for the willing—exceeding the effects of friends' attitude (6–9 p.p.). The direction of the top variables' effects varied by class. Vaccine confidence and trust in doctors similarly pushed predictions toward willingness and away from indecision or opposition. Child age had positive SHAP values across classes, with older age brackets increasing the probability of willingness. Right-wing political orientation, respondent's age, COVID-19 and flu vaccination showed class-specific contributions. Together, these patterns point to a constellation of cognitive, social, and experiential factors all shaping non-parents' vaccination intentions. Panels b1–b3 display SHAP values for the friends' attitude variable. Compared to parents (Fig. 1, panels b1–b3), this feature showed less discriminatory power, with a more uniform SHAP value distribution, though the same directional trend was evident.

Variable importance across models

Table 3 compares the top predictors of childhood COVID-19 vaccination intentions across models (XGBoost, RF, MNL), outcome classes (unwilling, undecided, willing), and respondent groups (parents and non-parents). Among parents, friends' attitude and vaccine confidence consistently ranked among the top three predictors across models and outcome classes. Models also showed strong agreement on other features such as COVID vaccine status, flu vaccination history, and child age, particularly to distinguish the unwilling and willing classes. For non-parents, the level of agreement was similarly high: child age, vaccine confidence, and friends'

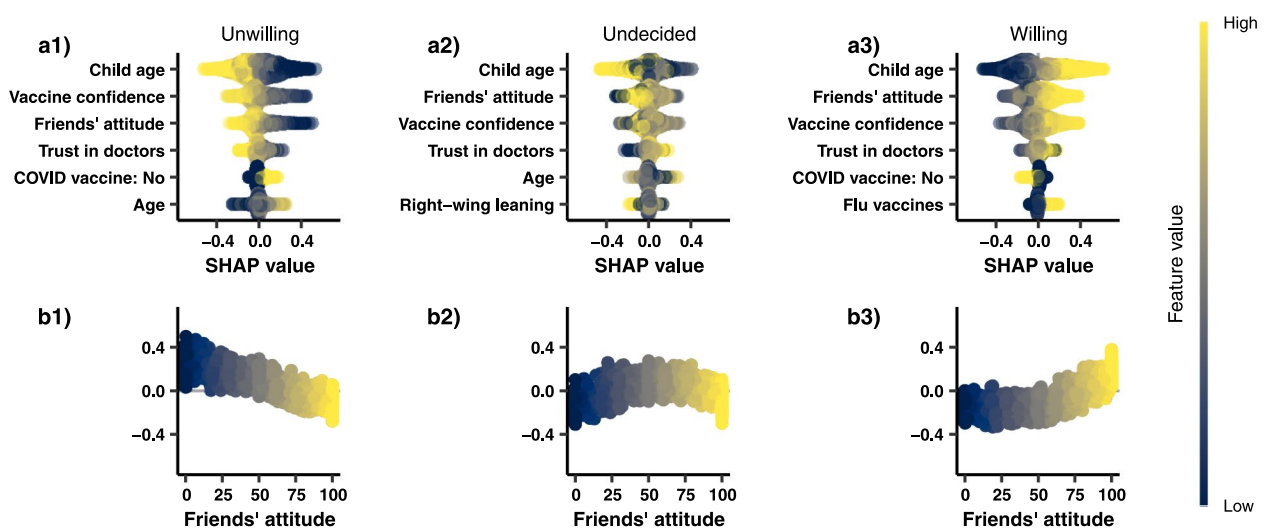


Fig. 2 Factors associated with intentions regarding childhood COVID-19 vaccination among non-parents. A random sample of one-third of the observations is plotted due to figure size constraints. The SHAP ranking of the top six predictors and the qualitative shape of the plots are consistent with those of the overall sample. Top: Beeswarm plots of the six most important predictors for classifying unwilling (a1), undecided (a2) and willing (a3) according to SHAP values. Bottom: Dependence plots showing the relationship between friends' attitude toward vaccination and SHAP values for unwilling (b1), undecided (b2) and willing (b3)

Table 3 Top-ranked predictors by model, outcome class, and respondent group

Group	Feature	Unwilling			Undecided			Willing		
		XGB	RF	MNL	XGB	RF	MNL	XGB	RF	MNL
Parents	Friends' attitude	1	1	1	1	1	1	1	1	1
	Vaccine confidence	2	2	3	2	2	3	2	2	3
	COVID vaccine: No	3	3	8				4	3	8
	Flu vaccines	4	4	9	6	4	9	5	4	9
	Child age	5	8	5	4	8	5	3	8	5
	Trust in government	6	5	14						
	Right-wing leaning				3	11	58	6	11	58
Non-parents	Age				5	7	17			
	Child age	1	1	1	1	1	1	1	1	1
	Vaccine confidence	2	2	3	3	2	3	3	2	3
	Friends' attitude	3	3	2	2	3	2	2	3	2
	Trust in doctors	4	5	13	4	5	13	4	5	13
	COVID vaccine: No	5	4	10				5	4	10
	Age	6	7	11	5	7	11			
Right-wing leaning				6	14	64				
	Flu vaccines						6	8	14	

Table 4 Class-specific performance scores of XGBoost models by respondent group

Class	Precision	Recall	F1	MCC	AUC
Panel A: Parents					
Willing	0.692	0.741	0.715	0.529	0.863
Undecided	0.553	0.503	0.527	0.324	0.764
Unwilling	0.731	0.735	0.733	0.618	0.911
Panel B: Non-parents					
Willing	0.659	0.711	0.684	0.484	0.826
Undecided	0.430	0.393	0.411	0.179	0.669
Unwilling	0.607	0.600	0.603	0.410	0.798

attitude appeared in the top three across models for all three outcome classes. This consistency reinforces the robustness and generalizability of our findings. Full SHAP rankings for all models are reported in Supplementary Information sections S4.2, S12.2 and S13.2.

Challenges in classifying the undecided

Performance by outcome class reveals substantial heterogeneity in classification accuracy (Table 4). In both parent and non-parent samples, the undecided group is consistently the most difficult to classify, with the lowest F1 and MCC scores (e.g. F1 = 0.527, MCC = 0.324 among parents; F1 = 0.411, MCC = 0.179 among non-parents). These results suggest greater ambiguity and overlap in

the predictor patterns associated with indecision. By contrast, unwilling individuals are classified with the highest accuracy, especially among parents (F1 = 0.733, MCC = 0.618). Among non-parents, performance remains highest for the willing (F1 = 0.684, MCC = 0.484) and unwilling (F1 = 0.603, MCC = 0.410) groups. This ranking is mirrored in the AUC values, which exceed 0.91 for unwilling, 0.82–0.86 for willing, and are lower for undecided (AUC = 0.67–0.76).

Added value of behavioural profiling

Table 5 compares the predictive performance of XGBoost models trained separately on administrative and behavioural variables. Among parents, the behavioural model outperformed the administrative model across all metrics, with improvements of approximately 7–11 percentage points in Accuracy, F1, and MCC (e.g. MCC: 0.424 vs. 0.291). This suggests that beliefs, attitudes, and trust-related variables provide substantial additional explanatory power when predicting parental willingness around childhood COVID-19 vaccination. Among non-parents, the performance gap was narrower: the behavioural model also outperformed the administrative model, though the improvements were more modest (e.g. Accuracy: 0.535 vs. 0.456; AUC: 0.707 vs. 0.643). These

Table 5 Macro performance of XGBoost models using administrative or behavioural predictors only

Model	Accuracy	Precision	Recall	F1	MCC	AUC
Panel A: Parents						
Administrative	0.532	0.521	0.532	0.523	0.291	0.732
Behavioural	0.620	0.618	0.613	0.614	0.424	0.809
Panel B: Non-parents						
Administrative	0.456	0.438	0.443	0.439	0.168	0.643
Behavioural	0.535	0.521	0.520	0.516	0.289	0.707

results indicate that while both types of information are useful to classify individuals, behavioural features may be more critical for parents, whereas administrative data alone captures a large share of predictive signal among non-parents.

Sensitivity analyses

We confirmed the robustness of our findings using multiple strategies. Model performance remained comparable across alternative data balancing approaches (upsampling and downsampling, with respect to no rebalancing in the main analysis; Supplementary Information section S6). When testing alternative cut-offs for discretizing the 0–100 outcome scale (Supplementary Information section S7), overall performance metrics (mean F1) were largely stable across thresholds. Nonetheless, a marked drop in the F1 score for the undecided class, particularly among non-parents, emerged when narrower thresholds were applied (e.g. 35–65, 30–70, 25–75 compared to 20–80), while gains for the willing and unwilling classes were marginal. The ranking of predictors based on SHAP values among the untreated respondents was largely consistent with the main results (Supplementary Information section S8). Likewise, alternative variable importance metrics (Gain and permutation importance) identified similar top determinants (Supplementary Information section S9). Finally, analyses replicated on the imputed dataset produced comparable predictive performance and stable SHAP-based variable rankings (Supplementary Information section S10).

Discussion

Public health responses to vaccine uptake often rely on binary distinctions between the “willing” and “hesitant” [11, 12]. However, recent guidance from the European Centre for Disease Prevention and Control (ECDC) recommends avoiding rigid labels such as “hesitant”, which risk oversimplifying a continuum of attitudes [37]. Our study contributes to this shift by treating vaccine willingness as a continuous construct and identifying individuals in ambivalent or transitional states through a three-category classification. Specifically, we pursued three objectives: (1) to identify individuals with ambivalent willingness regarding childhood COVID-19 vaccination; (2) to assess the added predictive value of behavioural variables over administrative and socio-demographic features; and (3) to uncover the most influential predictors of vaccine intentions. To our knowledge, this is one of the first large-scale studies to robustly address these questions using interpretable, cross-validated machine learning models in representative adult samples, including both parents and non-parents.

Key findings

Our models accurately distinguished individuals across the three outcome classes, with tree-based algorithms (XGBoost and RF) consistently outperforming traditional MNL. This supports the use of machine learning techniques in public health research to guide targeted communication policies. Classification performance was higher among parents, suggesting that intentions regarding COVID-19 childhood vaccination may be more clearly segmented in this group. While we optimized predictive performance through cross-validation, our primary goal was to identify which predictors mattered most, for which group, and in what direction. Our SHAP analyses identified peer attitude, especially friends’ views, as the most influential determinant of vaccine intentions, echoing recent work on social homophily and opinion clustering in vaccine behaviour [38–40]. Other key predictors included personal vaccination behaviour, vaccine confidence, trust in healthcare providers, political orientation, and child age. Convergence on the most important predictors across machine learning algorithms (especially XGBoost and RF, the best-performing ones) reinforced the robustness of our results. While peer attitudes were consistently important, they had greater discriminatory value among parents, possibly due to more salient social comparisons within parenting networks. Conversely, child age was the strongest predictor among non-parents, indicating that their support for vaccination varied depending on the age of the child presented. This pattern likely reflects broader societal beliefs or normative expectations about when childhood vaccination is perceived as most appropriate or necessary.

Despite strong overall model performance, classification was consistently less accurate for the undecided group than for the willing or unwilling. This finding aligns with prior segmentation research on vaccine hesitancy [41] and reinforces the notion that ambivalent intentions are less stable, more context-dependent, and not strongly anchored in identity or ideology. Such heterogeneity poses challenges for predictive modelling but also highlights the undecided as a potentially receptive target for tailored outreach. Nonetheless, SHAP analyses showed that the key predictors for this group mostly overlapped with those of both the willing and the unwilling: undecided respondents were more likely to be embedded in peer groups that were also undecided, suggesting the possibility to leverage social influence. Finally, our results underscored the value of behavioural profiling: variables such as attitudes, trust, and peer norms consistently outperformed administrative indicators in identifying unwilling and undecided individuals.

Policy and practical implications

While administrative data remain crucial for operational planning, they offer limited insight into the belief structures shaping vaccine decisions. Integrating behavioural indicators into population surveys, health registries, and real-time monitoring systems would substantially improve the ability to identify individuals who remain ambivalent or at risk of disengagement. Such integration could support the dynamic adaptation of communication campaigns to evolving patterns of trust, social norms, and vaccine sentiment, ultimately improving outreach effectiveness and minimizing the risk of stigmatization. This effort could be further enhanced by incorporating complementary behavioural data streams, such as mobility patterns or digital trace data, to increase the granularity, precision, and timeliness of public health strategies.

Our evidence also calls for a shift in how public health communication campaigns are designed. Rather than relying on uniform messaging, strategies should be adapted to behavioural profiles and local social contexts. For example, leveraging community influencers, such as friends, trusted local figures, or peer group leaders, could help shape vaccine narratives and normalize acceptance. Given the predictive importance of child age, outreach could also benefit from age-specific messaging tailored to perceived risks and benefits at different developmental stages.

Limitations and future research

This study has several limitations. First, it was conducted in two European countries shortly before COVID-19 vaccines became available for children. As such, the findings are valid in the context of COVID-19 childhood vaccination during the pandemic and may not generalize to other vaccines or future pandemics since vaccine intentions are highly context-dependent [42]. Indeed, the relatively low burden of severe COVID-19 among children and the intensity of societal debate may have created a decision-making environment distinct from that of routine immunisations. Moreover, while policy milestones and media representations may have triggered public debate at different times in Italy and the UK, our data do not allow us to identify the impact of these country-specific dynamics, and we therefore report pooled estimates. Nevertheless, the methodological framework developed here is adaptable to other contexts, and future research should examine whether similar patterns are observed across cultural settings, healthcare systems, and points in time. A second limitation stems from the cross-sectional design, which limits our ability to examine possible temporal dynamics in vaccine intentions. Vaccine sentiment is influenced by evolving public discourse, media narratives, and policy measures [43]; thus, longitudinal or repeated cross-sectional studies would enable

the identification of trajectories of intentions and the predictors of change. Future work could also investigate how shifting peer norms and communication strategies shape transitions along the willingness spectrum, including through experimental or quasi-experimental designs. Third, the study relied on self-reported survey data for all variables, including vaccine confidence, social network attitudes, and prior vaccination behaviour. These measures may be affected by recall bias, social desirability, or motivated reasoning [44], which could influence both model accuracy and the interpretability of SHAP attributions. Moreover, because the survey was conducted before vaccine rollout for children, we could not observe actual uptake, which is known to diverge from stated intentions [45, 46]. Future efforts should be directed toward linking survey responses with administrative vaccination records to provide objective behavioural benchmarks for attitudinal data. Fourth, although quota sampling was applied, participation was voluntary. As with most survey research, self-selection bias cannot be excluded, and certain relevant but hard-to-reach subgroups (e.g. individuals endorsing conspiracy beliefs about vaccination) may be under-represented. Addressing this limitation and succeeding in engaging less accessible populations remains a future challenge. Fifth, although tree-based models such as XGBoost demonstrated superior predictive accuracy, their complexity necessitates the use of post hoc interpretation tools. While SHAP values enhanced model interpretability by attributing feature contributions at the individual level, they remain conditional on the fitted model structure; as such, they should be interpreted as exploratory rather than explanatory. In addition, the undecided group, despite being a central focus of this study, remained particularly difficult to classify, and model performance was sensitive to how the continuous willingness scale was discretized. This suggests a need for cautious and transparent threshold selection in future applications. We speculate that the observed heterogeneity likely reflects the existence of meaningful subgroups (e.g. hesitant but persuadable vs. indifferent or disengaged). However, qualitative and mixed-methods research is needed to explore the cognitive and social processes underlying indecision. Finally, while our survey included measures on emotions, information sources, and perceived societal versus individual consequences, our analysis did not seek to disentangle reflective, “rational” decision-making from hesitancy rooted in emotional concerns or (mis)information. Reconstructing decision-making processes is inherently difficult with survey data, and exposure to misinformation cannot be directly inferred from self-reported sources. Future work could build on these measures to investigate how emotional, informational, and rational factors interact in shaping vaccine intentions,

and to identify what types of support parents may need to make informed decisions.

Abbreviations

DTP3	Three-dose diphtheria, tetanus and pertussis vaccine
ML	Machine learning
PCA	Principal component analysis
XGBoost/XGB	Extreme gradient boosting
RF	Random forest
MNL	Multinomial logistic regression
MCC	Matthews correlation coefficient
AUC	Area under the curve
SHAP	Shapley additive explanations
p.p.	Percentage points
k-NN	k-nearest neighbours
ECDC	European centre for disease prevention and control

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12963-025-00437-2>.

Supplementary file 1

Acknowledgements

We thank Tiziano Rotesi for his contribution to data collection; Lorenzo Lucchini and Duilio Balsamo for valuable discussions on methods; and Vittoria Offeddu for helpful feedback on the manuscript.

Author contributions

CC, LPL, AM, and MC contributed to the conceptualisation of the study. CC and LPL were responsible for methodology, formal analysis, data curation, validation, and software development. AM, MC, and PP contributed to investigation activities, including data collection. CC and LPL drafted the original manuscript. All authors (CC, LPL, AM, MC, and PP) participated in writing - review and editing. Supervision was provided by AM and CC. Project administration was carried out by CC. AM and PP were responsible for funding acquisition.

Funding

This work was supported by the PRIN project 2017ELHNNJ "The economics of vaccination", by the ERC Grant "IMMUNE" (ID: 101003183), by the PNRR Grant BEHAVE-MOD (ID: SA-02. P0001), and the Bocconi Covid Crisis Lab, supported by the "Fondazione Romeo ed Enrica Invernizzi". LPL additionally acknowledges the resources provided by Nuffield College, the Clarendon Fund, and the International Max Planck Research School for Population, Health and Data Science (IMPRS-PHDS).

Data availability

The dataset supporting the conclusions of this article will be available upon article acceptance in the Harvard Dataverse repository.

Declarations

Ethics approval and consent to participate

The surveys were approved by the Bocconi Research Ethics Committee. Informed consent was obtained from all individual participants included in the study.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Received: 21 July 2025 / Accepted: 30 November 2025

Published online: 24 December 2025

References

- World Health Organization.: Global childhood immunization levels stalled in 2023, leaving many without life-saving protection. Accessed: 2025-05-16. Available from: <https://www.who.int/news/item/15-07-2024-global-childhood-immunization-levels-stalled-in-2023-leaving-many-without-life-saving-protection>.
- World Health Organization.: Measles cases surge worldwide, infecting 10.3 million people in 2023. Accessed: 2025-05-16. Available from: <https://www.who.int/news/item/14-11-2024-measles-cases-surge-worldwide-infecting-10-3-million-people-in-2023>.
- Gimma A, Munday JD, Wong KL, Coletti P, van Zandvoort K, Prem K, et al. Changes in social contacts in England during the COVID-19 pandemic between March 2020 and March 2021 as measured by the CoMix survey: A repeated cross-sectional study. *PLoS Med.* 2022;19(3):e1003907.
- Trentini F, Manna A, Balbo N, Marziano V, Guzzetta G, O'Dell S, et al. Investigating the relationship between interventions, contact patterns, and SARS-CoV-2 transmissibility. *Epidemics.* 2022;40:100601.
- ECDC.: COVID-19 Vaccine Tracker. Accessed 2025-13-06. Available from: <https://vaccinetracker.ecdc.europa.eu/public/extensions/COVID-19/vaccine-tracker.html#uptake-tab>.
- Galanis P, Vraka I, Siskou O, Konstantakopoulou O, Katsiroumpa A, Kaitelidou D. Willingness, refusal and influential factors of parents to vaccinate their children against the COVID-19: a systematic review and meta-analysis. *Prev Med.* 2022;157:106994.
- Dubé E, Vivion M, MacDonald NE. Vaccine hesitancy, vaccine refusal and the anti-vaccine movement: influence, impact and implications. *Expert Rev Vaccines.* 2015;14(1):99–117.
- Bavel JJV, Baicker K, Boggio PS, Capraro V, Cichocka A, Cikara M, et al. Using social and behavioural science to support COVID-19 pandemic response. *Nat Hum Behav.* 2020;4(5):460–71.
- Hallsworth M. A manifesto for applying behavioural science. *Nat Hum Behav.* 2023;7(3):310–22.
- Plonsky O, Apel R, Ert E, Tennenholtz M, Bourgin D, Peterson JC, et al. Predicting human decisions with behavioral theories and machine learning. *arXiv preprint arXiv:1904.06866.* 2019.
- Carrieri V, Lagravinese R, Resce G. Predicting vaccine hesitancy from area-level indicators: a machine learning approach. *Health Econ.* 2021;30(12):3248–56.
- Cheong Q, Au-Yeung M, Quon S, Concepcion K, Kong JD. Predictive modeling of vaccination uptake in US counties: a machine learning-based approach. *J Med Internet Res.* 2021;23(11):e33231.
- Leask J. Target the fence-sitters. *Nature.* 2011;473(7348):443–5.
- Larson HJ. Negotiating vaccine acceptance in an era of reluctance. *Hum Vaccin Immunother.* 2013;9(8):1779–81.
- Betsch C, Korn L, Holtmann C. Don't try to convert the antivaccinators, instead target the fence-sitters. *Proc Natl Acad Sci U S A.* 2015;112(49):E6725–6.
- Chiavenna C, Leone LP, Melegaro A, Rotesi T, Bokemper SE, Paintlis EE, et al. Personal risk or societal benefit? Investigating adults' support for COVID-19 childhood vaccination. *Vaccine.* 2023;41(25):3683–7.
- Ward PR, Attwell K, Meyer SB, Rokkas P, Leask J. Understanding the perceived logic of care by vaccine-hesitant and vaccine-refusing parents: a qualitative study in Australia. *PLoS One.* 2017;12(10):e0185955.
- Rossen I, Hurlstone MJ, Dunlop PD, Lawrence C. Accepters, fence sitters, or rejecters: moral profiles of vaccination attitudes. *Soc Sci Med.* 2019;224:23–7.
- Enkel SL, Attwell K, Snelling TL, Christian HE. 'Hesitant compliers': qualitative analysis of concerned fully-vaccinating parents. *Vaccine.* 2018;36(44):6459–63.
- Rozbroj T, McCaffery K. The importance of addressing social inequalities and targeting the undecided to promote vaccination against COVID-19. *Lancet Reg Health Am.* 2021. <https://doi.org/10.1016/j.lanwpc.2021.100250>.
- Saldarriaga EM. Using machine learning to identify COVID-19 vaccine-hesitancy predictors in the USA. *BMJ Public Health.* 2023. <https://doi.org/10.1136/bmjph-2023-000456>.
- MacDonald NE, et al. Vaccine hesitancy: definition, scope and determinants. *Vaccine.* 2015;33(34):4161–4.
- DeCastellarnau A. A classification of response scale characteristics that affect data quality: a literature review. *Quality & quantity.* 2018;52(4):1523–59.
- Funke F, Reips UD. Why semantic differentials in web-based research should be made from visual analogue scales and not from 5-point scales. *Field Methods.* 2012;24(3):310–27.
- Krosnick JA. Questionnaire design. In: *The Palgrave handbook of survey research.* Springer; 2017:439–455.

26. Betsch C, Schmid P, Heinemeier D, Korn L, Holtmann C, Böhm R. Beyond confidence: development of a measure assessing the 5C psychological antecedents of vaccination. *PLoS One*. 2018;13(12):e0208601.
27. Larson HJ, Clarke RM, Jarrett C, Eckersberger E, Levine Z, Schulz WS, et al. Measuring trust in vaccination: a systematic review. *Hum Vaccin Immunother*. 2018;14(7):1599–609.
28. Rosenstock IM. The health belief model and preventive health behavior. *Health Educ Monogr*. 1974;2(4):354–86.
29. Chen T, Guestrin C. Xgboost: A scalable tree boosting system. In: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*; 2016:785–794.
30. Breiman L. *Random Forests Machine learning*. 2001;45:5–32.
31. Kuhn M. Building predictive models in R using the caret package. *J Stat Softw*. 2008;28:1–26.
32. Chicco D, Jurman G. The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*. 2020;21:1–13.
33. Hand DJ, Till RJ. A simple generalisation of the area under the ROC curve for multiple class classification problems. *Mach Learn*. 2001;45:171–86.
34. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. *Advances in neural information processing systems*. 2017;30.
35. Štrumbelj E, Kononenko I. Explaining prediction models and individual predictions with feature contributions. *Knowl Inf Syst*. 2014;41(3):647–65.
36. Greenwell BM, Boehmke BC. Variable importance plots—an introduction to the vip package. *R J*. 2020;12(1):343–66. <https://doi.org/10.32614/RJ-2020-013>.
37. ECDC.: Tools and methods for promoting vaccination acceptance and uptake: a social and behavioural science approach. Available from: <https://www.ecdc.europa.eu/sites/default/files/documents/vaccination-acceptance-uptake-tools-social-behavioural-science.pdf>.
38. Hâncean MG, Lerner J, Perc M, Molina JL, Geantă M. Assortative mixing of opinions about COVID-19 vaccination in personal networks. *Sci Rep*. 2024;14(1):3385.
39. Stefkovics Á, Albert F, Ligeti AS, Dávid B, Rudas S, Koltai J. Vaccination homophily in ego contact networks during the COVID-19 pandemic. *Sci Rep*. 2024;14(1):15515.
40. Thompson O, Cristea M, Tamariz M. COVID-19 vaccination attitudes and uptake: a sociocultural perspective focusing on parents and peers. *PLoS One*. 2024;19(7):e0300771.
41. Lincoln TM, Schlier B, Strakeljahn F, Gaudiano BA, So SH, Kingston J, et al. Taking a machine learning approach to optimize prediction of vaccine hesitancy in high income countries. *Sci Rep*. 2022;12(1):1–12.
42. Larson HJ, Jarrett C, Eckersberger E, Smith DM, Paterson P. Understanding vaccine hesitancy around vaccines and vaccination from a global perspective: a systematic review of published literature, 2007–2012. *Vaccine*. 2014;32(19):2150–9.
43. Jolley D, Douglas KM. The effects of anti-vaccine conspiracy theories on vaccination intentions. *PLoS One*. 2014;9(2):e89177.
44. Krumpal I. Determinants of social desirability bias in sensitive surveys: a literature review. *Quality & quantity*. 2013;47(4):2025–47.
45. Sheeran P, Webb TL. The intention-behavior gap. *Soc Personal Psychol Compass*. 2016;10(9):503–18.
46. Sheeran P. Intention-behavior relations: a conceptual and empirical review. *Eur Rev Soc Psychol*. 2002;12(1):1–36.

Publisher's note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.