



ELSEVIER

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Health Economics

journal homepage: www.elsevier.com/locate/jhealeco

Persistence and heterogeneity of the effects of educating mothers to improve child immunisation uptake: Experimental evidence from Uttar Pradesh in India

Stephen O'Neill^{a,*}, Richard Grieve^a, Kultar Singh^b, Varun Dutt^c, Timothy Powell-Jackson^d

^a Department of Health Services Research and Policy, London School of Hygiene & Tropical Medicine, London, United Kingdom

^b Sambodhi Research and Communications, Noida, Uttar Pradesh, India

^c ConveGenius Insights Pvt. Ltd, Hyderabad, India

^d Department of Global Health and Development, London School of Hygiene & Tropical Medicine, London, United Kingdom

ARTICLE INFO

JEL classification:

I1
I12

Keywords:

Vaccination
Randomised controlled trial
Heterogeneity machine learning
Causal forest

ABSTRACT

Childhood vaccinations are among the most cost-effective health interventions. Yet, in India, where immunisation services are widely available free of charge, a substantial proportion of children remain unvaccinated. We revisit households 30 months after a randomised experiment of a health information intervention designed to educate mothers on the benefits of child vaccination in Uttar Pradesh, India. We find that the large short-term effects on the uptake of diphtheria–pertussis–tetanus and measles vaccination were sustained at 30 months, suggesting the intervention did not simply bring forward vaccinations. We apply causal forests and find that the intervention increased vaccination uptake, but that there was substantial variation in the magnitude of the estimated effects. We conclude that characterising those who benefited most and conversely those who benefited least provides policy-makers with insights on how the intervention worked, and how the targeting of households could be improved.

1. Introduction

Enormous progress has been made in reducing child mortality and disability over the last two decades in low- and middle-income countries (LMIC), and childhood vaccinations have played an important part in this success story (Bhutta et al., 2013). They represent one of the most cost-effective health technologies, in that they prevent mortality and disability at relatively low cost (Bärnighausen et al., 2014). Yet, despite the well-documented evidence and consistent investment in national immunisation programmes, the WHO estimates that globally 25 million infants were not fully vaccinated in 2021 (World Health Organisation, 2022). More than 60 % of these children live in 10 countries: Angola, Brazil, the Democratic Republic of the Congo, Ethiopia, India, Indonesia, Myanmar, Nigeria, Pakistan and the Philippines. Not since 2009 has the number of children who are unvaccinated been so high (World Health Organisation, 2022). Understanding how to increase the uptake of vaccines is especially pressing following COVID-19, which not only interrupted routine vaccination services, but also highlighted the need to better understand the determinants of vaccine uptake.

* Corresponding author at: Department of Health Services Research and Policy, London School of Hygiene and Tropical Medicine, 15-17 Tavistock Place, WC1H 9SH London, United Kingdom.

E-mail address: Stephen.ONeill@lshtm.ac.uk (S. O'Neill).

<https://doi.org/10.1016/j.jhealeco.2024.102899>

Received 12 June 2023; Received in revised form 13 May 2024; Accepted 17 May 2024

Available online 18 May 2024

0167-6296/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Indeed, strategies that are successful in improving vaccine coverage for one disease may prove effective in raising acceptance of vaccines to protect against future outbreaks of COVID-19 or other emerging diseases.

The setting for our study is Uttar Pradesh, one of the most populous and poorest states in India, and in which 70 % of children aged 12 to 23 months are fully vaccinated against common childhood disease (International Institute for Population Sciences (IIPS) and ICF, 2021). While this represents a marked improvement over the past five years, it is clear that the widespread availability of free immunisation services in public facilities has been insufficient to guarantee high coverage and the benefits of herd immunity in the population. Consistent with this picture is a growing body of evidence that suggests demand-side factors, including poor parental knowledge, distrust, time costs and procrastination, are important barriers to vaccination uptake (Larson et al., 2014; Mills et al., 2005). This, in turn, has prompted investigation of light-touch behavioural interventions that target parents with health information messages, cash or in-kind incentives.

Various systematic reviews and meta-analyses (Shea et al. 2009; Johri et al., 2015; Oyo-Ita et al., 2016), as well as more recent studies (Banerjee et al., 2021; Gibson et al., 2017), have evaluated strategies for increasing coverage of childhood vaccinations. The most recent review included 14 studies evaluating a range of interventions such as health education, monetary incentives, home visits and supportive supervision (Oyo-Ita et al., 2016). Some of these interventions were found to be effective in the short term – health education interventions, for example, improved immunisation coverage by 68 % – although the quality of evidence varied. Despite this body of literature, the evidence is limited on two important questions.

First, are the effects of demand-side interventions sustained over time? It may be the case that the initial effects of an intervention are attenuated over time, if the intervention merely brings forward vaccinations that would have happened anyway. Under such a scenario, any health benefits of vaccination would be temporary, and so if estimates of the relative cost-effectiveness of the intervention are based on immediate impacts of uptake, then this would overstate the cost-effectiveness of the intervention. An additional perspective on the question of sustainability concerns the persistence of behaviour change in response to temporary or one-off interventions (Celhay et al., 2019). Behaviour change of this nature can be thought of as habit formation, which has obvious relevance for smoking cessation and exercise interventions (Charness and Gneezy, 2009; Volpp et al., 2009). In the context of childhood immunisation, an intervention that led to a sustained change in parental behaviour could deliver benefits – in terms of vaccine uptake – to children who were not born at the time of intervention. Such evidence from follow-up of additional children would serve to increase the cost-effectiveness of an intervention.

Second, who benefits from the intervention? Evidence on heterogeneous treatment effects has a number of uses. It can add to existing knowledge and a priori reasoning to help inform policymakers as to who should be targeted by the intervention to maximise uptake, and in the context of immunisation, reach the herd immunity threshold. It can be informative as to which groups the effects of the intervention are most persistent for. It can provide insights into potential inequities, for example by indicating how widely the benefits of an intervention are felt. It can potentially shed light on the mechanisms through which the intervention worked and may be informative as to which groups the intervention has or does not have persistent effects for. Finally, it can offer policymakers insights on how to adapt the intervention or what other forms of intervention may be needed in tandem with demand-side strategies to increase effectiveness in certain subgroups. For example, if the intervention is shown not to work for remote households far from public health facilities, additional strategies such as community outreach or immunisation camps may be needed. While distinct, the two questions are of keen interest to policymakers seeking to implement interventions at scale. Interventions whose effects are short-lived or undermined by the passage of time are of limited value to public health officials. Evidence on what drives variation in intervention effects can provide valuable information on how a public health programme should be designed and delivered outside the confines of a research project.

This paper addresses these questions in the context of a brief health education intervention that was undertaken in rural Uttar Pradesh, India from 2015 to 2016 (Powell-Jackson et al., 2018). The setting has general appeal; it is relevant to low-income countries in which immunisation services are free at the point of use, yet immunisation uptake is low. The intervention provided the mothers of unvaccinated or incompletely vaccinated children aged 0 to 36 months with health information on the benefits of vaccination through home visits. It was implemented as an individually randomised controlled trial and outcomes were measured seven months after the information was given. The intervention led to a large immediate increase in vaccination uptake. Coverage of three doses of diphtheria–pertussis–tetanus vaccine (DPT3) was 28 % in the control group and 43 % in the intervention group (risk difference of 15 percentage points, $p < 0.001$), and coverage of measles vaccination was 42 % in the control group and 64 % in the intervention group (risk difference of 22 percentage points, $p < 0.001$). The cost per disability-adjusted life year averted of providing information was US \$186, implying that the intervention was highly cost-effective.

We use new data and analytical methods to assess the sustainability and the heterogeneity of effects of the educational intervention. Fieldworkers returned to the study participants approximately 30 months after the intervention, and measured outcomes amongst the original sample of children (hereon referred to as the index children) and younger siblings that were not yet born at the time of intervention. Levels of attrition over the follow-up period were low: 93 % of the 722 study participants who were randomised completed follow-up at 30 months, and for this subsample baseline characteristics remained well balanced between the randomised groups.

To study heterogeneity, we estimate individual treatment effects using causal forests (Wager and Athey, 2018), an ensemble machine learning approach that is becoming increasingly popular. The Causal Forest approach is a non-parametric method that builds on causal trees (Athey and Imbens, 2016), which recursively splits individuals into groups with a rule tailored towards the estimation of heterogeneous treatment effects. The method allows for high dimensional interactions between covariates while avoiding overfitting by repeatedly estimating causal trees from random subsets of the data, using the remainder of the data to predict effects, and then averaging the predictions to obtain an overall predicted outcome for each individual under each treatment state (Wager and

Athey, 2018). The difference between the two predictions is then the individual level-effect estimate. By considering variation in these estimated individual-level effects with respect to covariates, we can characterise the groups that benefit most/least from the intervention. Alternatively, the individual-level effect estimates for pre-specified subgroups can be aggregated to obtain subgroup effect estimates. Subgroup analyses have traditionally been seen as controversial. For good reason, there is much scepticism of *ex post* analysis of subgroups because of the risk that results are selectively reported owing to their statistical significance. Studies typically have substantially less statistical power to estimate subgroup effects than overall effects, and must fully acknowledge any lack of precision in reporting and interpreting subgroup-level effects (Burke et al., 2015). To guard against data mining, there are established norms around trial registration and pre-specification of subgroup analyses. But relying solely on information about those subgroups that are prespecified risks discarding potentially valuable information, which may be useful to help target interventions, but also to inform priorities for future research. Moreover, there may be genuine uncertainty regarding what factors may influence effectiveness. When interest is around hypothesis generating, existing approaches may be excessively conservative, and fail to raise new hypotheses. Machine learning allows the researcher to stay neutral as to the source of heterogeneity (i.e. the effect modifiers) and discover patterns in the data by searching over high-dimensional functions of covariates. Such a machine learning method can complement the approach of *a priori* specification of a limited set of outcome models or subgroups of interest, and the extent to which particular subgroups have been specified *a priori*, and the use of a theory or intuition to inform the likely direction of effects for particular subgroups is important in interpreting the strength of recommendations for policy-making and further research.¹ Effects can still be aggregated for a limited number of pre-specified subgroups to test pre-specified hypotheses (e.g. by mothers perception of vaccine efficacy), and also to build on 'theory' or 'intuition' for further subgroups that were not pre-specified (e.g. age and vaccination history), in a way that is useful for helping target interventions and future research priorities. The use of 'honest' estimation, where the same data are never used for estimation and sample splitting helps to protect against false discovery and yields confidence intervals with correct coverage (Athey and Imbens, 2016). Nonetheless, studies should be careful in the strength of the policy recommendations and further research recommendations that are made from the results of subgroup analyses according to whether the subgroups are pre-specified, and the extent to which they are predicted by theory or prior reasoning.

We report three key findings. First, the large initial effect of the intervention on vaccination outcomes was maintained at 30 months follow-up. The magnitudes of the effect in absolute terms were similar to those previously reported at seven months follow-up, suggesting that the intervention did not simply bring forward vaccinations that would have happened anyway. In this sense, the effects of the brief intervention were sustained. We are unable to make firm conclusions as to the persistence of parental behaviour change because the confidence intervals on the effects of the intervention on vaccination uptake of younger siblings are wide. The effect on uptake of DPT3, for example, was 9 percentage points amongst younger siblings (compared with 32.5 % in the control group) but confidence intervals ranged from -2.2 to 20.3 percentage points.

Second, estimates of individual level treatment effects show that the majority of participants benefited from the intervention. For DPT3 vaccination at 30 months follow-up, individual treatment effects ranged from 1.4 to 28.6 percentage points, with a statistically significant effect observed for 56.4 % of the participants. For measles vaccination at 30 months follow-up, individual treatment effects ranged from 16.3 to 44.5 percentage points, with a statistically significant effect observed for all the participants, after excluding those children that had already received the measles vaccine at baseline. These findings therefore suggest that the intervention did not cause harm by reducing the chances of children being vaccinated, as would be expected with an intervention of this nature.

Third, we examined whether the heterogeneity was associated with baseline characteristics. When looking at DPT3 uptake, individuals that benefitted most (top 25 % of effects) from the intervention tended to be older, had received previous vaccinations in the schedule (e.g. the first and second doses of DPT), were more likely to be located closer to a public rather than private health facility, and had mothers that demonstrated less knowledge about the causes, symptoms and prevention methods for tetanus compared to those that benefitted least (bottom 25 % of effects). Age and the receipt of the first and second doses of DPT were also strongly associated with larger effects on measles vaccine uptake, while mothers' knowledge regarding tetanus or proximity to a public health facility did not explain variation in treatment effects.

In a broad sense, our study contributes to the literature on demand-side interventions for immunisation uptake in low- and middle-income countries, providing novel insights on the persistence and heterogeneity of effects (Banerjee et al., 2020, 2021; Gibson et al., 2017; Johri et al., 2015). There is a small literature on whether temporary incentives can lead to healthy habit formation, such as smoking cessation and exercise (Charness and Gneezy, 2009; Volpp et al., 2009), and a rich body of theoretical work on how to maintain behaviour change (Kwasnicka et al., 2016). There is, however, a need for more evidence on whether one-off health education interventions can lead to sustained changes in behaviour in the uptake of health care technologies. We find little evidence that the intervention caused harm by reducing the chances of children being vaccinated in contrast to another study in India, where a combination of small incentives, reminders and persuasion was found to reduce immunisation rates in some villages, possibly because the interventions crowded out existing intrinsic motivation of parents to vaccinate their children (Chernozhukov et al., 2018).

The paper is structured as follows: Section 2 provides background information on the study setting, and the information intervention. Section 3 discusses the original experimental design and the data. Section 4 describes the econometric methods used. Section 5 presents the results and Section 6 discusses the findings with respect to the limitations of the study and the broader literature.

¹ Machine learning does not, of course, provide the means to make a causal interpretation of the heterogeneity results since covariates could be proxying for other characteristics.

2. Background

2.1. Context

The study took place in the following six districts of the state of Uttar Pradesh: Kannauj, Kanpur Nagar, Kanpur Dehat, Auraiya, Etawah, and Fatehpur. Uttar Pradesh has more than 230 million people with GDP per capita and levels of literacy typical of a low- or lower-middle income country. At the time of the study, 67 % of children aged 12–23 months had received the DPT3 vaccine, 71 % had received the measles vaccine, and 51 % were fully vaccinated. There was considerable dropout between the first and third doses of DPT vaccine. The study districts had a population of 13.7 million people.² DPT3 coverage amongst children 12–23 months ranged from 58 % to 78 %, and full vaccination coverage ranged from 34 % to 62 % in the study districts (International Institute for Population Sciences (IIPS) and ICF, 2017). The reasons for the insufficient coverage of childhood vaccinations are multi-faceted. However, there is a general consensus that it is not because of a lack of availability in the supply of vaccines. The national programme provides childhood vaccinations at no cost to the parents and there is established infrastructure and health personnel of different cadres to deliver vaccines to rural areas. The delivery system is aided by an extensive network of accredited social health activists in the community who are expected to keep an up-to-date list of households eligible for immunisation, and encourage parents to get their children vaccinated.

2.2. Brief health education intervention

The original study tested a health education intervention, designed and implemented in partnership with Sambodhi Research and Communications, a research organisation in Uttar Pradesh. The intervention focused on tetanus, a serious disease and one of the leading causes of death amongst newborns in India. There is a highly effective vaccine against tetanus given as a combined shot. The Indian Academy of Paediatrics recommends that three doses of DPT should be given, at 6 weeks, 10 weeks, and 14 weeks, with a minimum age of 6 weeks. If any of these doses are missed, there is a catch-up range of up to seven years of age (Kasi et al., 2021).

Field staff provided mothers with information on childhood tetanus and the benefits of the tetanus vaccine through door-to-door visits. Information was conveyed using a structured script, alongside visual aids. The script described the causes and symptoms of tetanus, possible health consequences, the individual health benefit of the combination DPT, and the wider community benefits associated with herd immunity. We tested two versions of the script that varied how the information was framed. The first script framed information on tetanus vaccination as gains, emphasising that the child would be less likely to get tetanus and more likely to be healthy once vaccinated. The second script framed the information as a loss, highlighting that an unvaccinated child would be more likely to get tetanus and suffer the health consequences of the disease. There was a question and answer session and a Hindi leaflet was left with the mother at the end. (See Appendix A for further details and the leaflets). The intervention was brief; it took about 10 min to deliver to each household (Powell-Jackson et al., 2018).

Mothers were eligible for inclusion in the study if their child was alive, was aged 0–36 months, had not received three doses of DPT vaccine, and if the mother intended to remain in the study area for at least six months. Eligibility was determined using two sources of data: 1) a household survey conducted in the same villages prior to the original study; and 2) a list of mothers who had recently given birth provided by accredited social health activists working in each of the study villages. The intervention was delivered to eligible households in 180 clusters (villages) in the six study districts in September 2015. The baseline survey was conducted at the same time as the door-to-door home visit, prior to treatment group assignment and the provision of the information.

3. Experimental design and data

3.1. Experimental design

The trial randomised 722 mothers of children aged 0 to 36 months, in a ratio of 1:1:1 to one of three study arms: mothers in the first treatment group received information framed as a gain, mothers in the second treatment group received information framed as a loss, and the third arm acted as a control group, with no information given to the mother. The comparison groups were well balanced at baseline (Powell-Jackson et al. 2018). In this paper, we combine the first two arms into a single 'treatment' group since this was pre-specified as the primary analysis in the study protocol and the original evaluation did not find evidence that the framing mattered. The interventions were delivered during the first home visit, once the mother had given their consent to participate and had been interviewed for the baseline survey.

3.2. Data

Three rounds of data collection were undertaken, at baseline (September 2015) and in two follow-up surveys, seven months (April 2016) and 30 months (March 2018) after the information intervention. In the first follow-up, the study team completed interviews with 706 mothers whose child was still alive, resulting in a loss to follow-up of 16 mothers and a rate of attrition of 2.2 %. In the second follow-up, the study team completed interviews with 674 mothers, resulting in an additional loss to follow-up of 32 mothers, and an

² The sampling procedures to select the districts and study clusters are described in detail elsewhere (Tougher et al., 2018).

overall attrition rate of 6.6 % against the original sample. The study team also obtained data on outcomes for 299 younger siblings of the index child. It should be noted that the original study did not seek to have adequate power to assess effects for siblings. In this paper, we draw primarily on the new data from the second follow-up survey. Each round of data collection captured the child's immunisation status, and the mother's knowledge of the causes of, symptoms of, and prevention methods against tetanus.

We focus on two vaccination outcomes. The first outcome is the proportion of children who had received three doses of DPT vaccine. This was the pre-specified primary outcome of the original trial. The second is the proportion of children who had received the measles vaccine. The previously published results of the trial reported that the intervention had a large positive effect on this outcome, even though the information intervention itself was focused entirely on tetanus (Powell-Jackson et al., 2018). Measles is also the last vaccine in the standard immunisation schedule, and hence a good guide as to whether the child has received the recommended vaccines. We measured the vaccination status of the child following international best practice of relying on the vaccination card or, where unavailable, self-reported information from the mother (International Institute for Population Sciences (IIPS) and ICF., 2021).³ It is important to note that, of the 706 children included in the first follow-up survey, 116 had received the measles vaccine at baseline. The effect of the intervention on measles vaccine uptake amongst these children must be zero, since they were already vaccinated. In other words, the relationship is mechanical. There is therefore nothing to be learnt about the targeting of the intervention for this subgroup of children. Hence, we exclude children that had received the measles vaccine at baseline in the analysis of heterogeneous treatment effects for measles vaccine uptake. We do, however, carry out a validation/falsification test where we include these individuals in the estimation and assess whether we correctly estimate that the intervention had no effect on this group.

The original study protocol mentioned the possibility of conducting subgroup analyses with respect to five variables. Three of the five variables (sex of child, household wealth, education of mother) were not measured in the baseline survey to reduce the time burden on the participants and hence cannot be considered here. Two of the five variables are considered below the mother's perception of efficacy, and the closest health facility). The intervention sought to increase demand for vaccination by addressing misperceptions of the efficacy of the tetanus vaccine such that mothers with perceptions of efficacy (as measured by an index of 0 to 10 based on responses to interactive games on hypothetical questions relating to different immunisation coverage scenarios) lower than the true efficacy may be more responsive to the intervention. The type of health facility (government or private) that was closest to them was considered since this may reflect access to the publicly-provided immunisation services.

We also used all available information from the baseline household survey to consider a number of further subgroups that a priori reasoning suggests may modify the relative effectiveness of the intervention (See Table 1). Children further along in their vaccination schedule – as measured by prior vaccinations and proxied by age – may be more affected by the intervention since they are closer to being ready for DPT3 and measles vaccination. The intervention also sought to increase mothers' knowledge of the causes, symptoms, and prevention methods of tetanus such that those with poor knowledge at baseline may be more affected by the intervention.

4. Econometric methods

4.1. Estimands of interest

Defining $Y_i(1)$ and $Y_i(0)$ as individual i 's potential outcomes with and without treatment respectively (Rubin, 1974), the individual's treatment effect can be defined as $\tau_i = Y_i(1) - Y_i(0)$. However, since it is not possible to observe both potential outcomes simultaneously (Holland, 1986), this effect is generally unidentifiable. In a trial, we rely on randomisation to allow us to identify the ATE:

$$ATE = E(\tau_i) = E(Y_i(1) - Y_i(0)) \quad (1)$$

Here, we are interested in conditional effects, that is the contrast between the two treatment arms, conditional on observed baseline covariates. Formally, the estimand of interest is the conditional average treatment effect (CATE):

$$\tau(x) = E(Y_i(1) - Y_i(0)|X = x) \quad (2)$$

where X can consist of a combination of the observed covariates (i.e. X can be a vector). By considering an individual's covariates we can use $\tau(x)$ to estimate an *individualised* treatment effect, which can be aggregated to estimate the CATE for subgroups of interest. The sample average of these individualised effects can be taken over the full sample to obtain an estimate of the average treatment effect (ATE). While in principle $\tau(x)$ can be estimated using standard regression approaches, it is challenging to correctly specify the relationship between confounders and outcomes, and the form of effect modification. We therefore adopt a causal machine learning method, Causal Forest, to estimate these individualised effects.

4.2. Estimation of persistence of effects

We examine two aspects of persistence. First, to assess whether the effects of the intervention on the index children are sustained over time, we estimate the ATE at 30 months follow-up and compare it to the estimated effect at seven months follow-up. We present estimates from an OLS regression without and then with adjustment for baseline covariates. If the initial effects of the intervention are

³ 43% of mothers had a vaccination card. There was no difference between treatment and control.

Table 1

Descriptive statistics for baseline covariates across the comparator groups for those included in the seven and 30 months follow-up.

	At seven months follow-up			At 30 months follow-up		
	All Mean (SD)	Treated Mean (SD)	Control Mean (SD)	All Mean (SD)	Treated Mean (SD)	Control Mean (SD)
Age of index child at baseline (months)	10.34 (7.69)	10.38 (7.82)	10.26 (7.44)	10.24 (7.69)	10.32 (7.85)	10.08 (7.37)
Perception of tetanus vaccination efficacy (index)	7.26 (1.98)	7.20 (2.02)	7.39 (1.90)	7.27 (2.00)	7.20 (2.03)	7.40 (1.92)
	%	%	%	%	%	%
Received DPT 1st dose	65.4	63.3	69.8	65.6	63.8	69.2
Received DPT 2nd dose	41.8	39.3	46.8	41.7	39.3	46.4
Received BCG vaccine	84.3	82.4	88.1	84.1	82.2	87.9
Received Measles vaccine	16.4	16.1	17.0	16.5	15.8	17.9
Mother knows a cause of tetanus	43.1	41.2	46.8	43.0	41.3	46.4
Mother knows a symptom of tetanus	8.4	8.3	8.5	8.3	8.0	8.9
Mother knows a prevention method of tetanus	40.4	39.7	41.7	40.7	40.2	41.5
Closest health facility: Government	86.8	86.8	86.0	86.3	86.8	85.3
Closest health facility: Private/Other	13.2	13.2	14.0	13.7	13.2	14.7
Number of observations	706	471	235	674	450	224

Notes: Table reports the mean and standard deviations of baseline covariates for the full sample, and for the treated and control groups, at seven months and 30 months follow-up. The measles subsample excludes individuals that had received the measles vaccine at baseline.

attenuated over time, this may indicate that the intervention merely brings forward vaccinations that would have happened anyway. To test this, we pool the data from the two follow-up surveys, restricting to those present at both timepoints and estimate the difference in effects between follow-up periods ($\tau_{30 \text{ months}} - \tau_{\text{seven months}}$), and test the null hypothesis that this difference is non-negative (i.e. the effect has not attenuated) using a one-tailed paired *t*-test.

Second, to assess whether the intervention led to a sustained change in parental behaviour, we take advantage of additional data collected in the 30 month follow-up survey on the vaccination status of the index child's sibling that was born after the intervention ($N = 293$). We assess whether sibling's vaccination status differed according to whether or not the mother was assigned to the intervention ($N_{\text{treated}} = 176$; $N_{\text{control}} = 117$). We hypothesise that if the intervention had persistent effects on parents' behaviour, the intervention would positively influence vaccine uptake for siblings. We note that the siblings themselves were not randomised to the intervention groups, raising the possibility of imbalance in baseline characteristics between the comparison groups of siblings. However, it seems unlikely, that the information intervention, or vaccination of the index child, could influence whether and when a subsequent child was born.

4.3. Heterogeneous treatment effects

4.3.1. Causal forest method

To estimate individual-level effects of the intervention, we apply the Causal Forest algorithm. Causal Forests are an ensemble of non-parametric causal trees (Athey and Imbens, 2016), which are built recursively by splitting observations into groups based on whether a particular covariate exceeds a threshold. Each time a split is made, the covariate and threshold used are chosen to maximise the variance of the estimated treatment effect, $\hat{\tau}(x_i)$ for the sample used to define the split. Thus, sample splits are formed so that the estimated treatment effect is as homogenous as possible within a leaf (created by splitting at the thresholds), and as different as possible between leaves. Under an unconfoundedness assumption, the mean of the observed outcomes for the individuals in the treated and control groups within a leaf represent the estimates of the mean potential outcomes for that leaf defined by the covariates that determined that split ($X = x$) allowing the individualised effect $\tau(x)$ within the leaf L to be calculated as:

$$\hat{\tau}(x) = \left(\frac{1}{|\{i : D_i = 1, X_i \in L\}|} \sum_{\{i: D_i=1, X_i \in L\}} Y_i \right) - \left(\frac{1}{|\{i : D_i = 0, X_i \in L\}|} \sum_{\{i: D_i=0, X_i \in L\}} Y_i \right) \quad (3)$$

Thus, the estimated effect for the subgroup is the difference in average outcomes for treated versus control units within the leaf of the tree, L , in which the unit lies.

Causal trees, like decision trees, are prone to overfitting. This can be mitigated by using a Causal Forest, defined as an ensemble of B causal trees, analogous to the use of random forests to mitigate overfitting by decision trees. This implies averaging predictions $\hat{\tau}_b(x)$ over a large number of different possible covariate splits to estimate a CATE for each individual in the sample (Wager and Athey, 2018), reducing variance and smoothing sharp decision boundaries (Bühlmann, 2002; Wager and Athey, 2018):

$$\hat{\tau}(x) = \frac{\sum_{b=1}^B \hat{\tau}_b(x)}{B} \quad (4)$$

If the same data are used to both decide splits and to estimate effects, inference will be biased because the splits are specifically chosen to give more different effects across groups. Valid asymptotic confidence intervals for the true underlying treatment effect (Wager and Athey, 2018) are thus obtained using a sample splitting or 'honest' estimation approach (Athey and Wager, 2019), where an observation is never used to both determine splits and estimate effects at the same time. This approach yields valid confidence

intervals with coverage rates that do not deteriorate as the data generating process becomes more complex, or more covariates are added to the forests.

We can also view Causal Forests as a locally weighted estimator that uses forest based weights – i.e. gives more weight to observations that are similar to the unit of interest when estimating effects (Athey and Wager, 2019), where units are deemed to be similar if they tend to lie in the same leaf of the trees in the Causal Forest. More specifically, we fit two separate regression forests to estimate response functions for propensity for treatment $[\hat{e}(x)]$ and outcome $[\hat{m}(x)]$. As the intervention is randomised, the $\hat{e}(x)$ is less important here than in observational studies that rely on ‘selection on observables’. Next we make ‘out-of-bag’ predictions (i.e. predictions based on trees that did not include the i^{th} observation) using these treatment propensity and outcome forests. The CATE is then estimated using:

$$\hat{\tau}(x) = \frac{\sum_{i=1}^n \alpha_i(x) (Y_i - \hat{m}^{(-i)}(X_i)) (W_i - \hat{e}^{(-i)}(X_i))}{\sum_{i=1}^n \alpha_i(x) (W_i - \hat{e}^{(-i)}(X_i))^2} \quad (5)$$

where $\alpha_i(x)$ is the learned adaptive weight for individual i capturing how often individual i falls into the same leaf as x (Athey et al., 2019). By considering each individual’s covariates x , we can estimate an effect for each individual using this approach.

We can then aggregate the individual-level effects to obtain subgroup effects, with a variant of doubly robust estimators already implemented in the generalized random forest R package *grf* (Tibshirani et al., 2020). Here we use augmented inverse propensity weighting (AIPW) (Athey and Wager, 2019; Robins et al., 1994) to account for imbalance in covariates that were not used to split on when forming a particular tree, providing efficient estimates. AIPW can lead to unstable estimates where propensity scores are close to 0 or 1, however this is not the case here, with propensity scores tending to be close to the rate of assignment to treatment (P(Treated) = 2/3). Tuning parameters (e.g., minimum node size for individual trees) are chosen by cross-validation.⁴ Following Athey & Wager (2019) and Basu et al. (2018), we estimate a second Causal Forest, using only the variables that saw a reasonable number of splits in the original Causal Forest, to improve precision. That is, those variables that had low variable importance scores in the original Causal Forest are excluded,⁵ allowing the forest to make more splits on the most important features in low-signal situations (Athey and Wager, 2019). This is likely to be important here given the relatively small sample sizes.

4.3.2. Implementation of estimation approach

We estimate CATEs for each individual based on their covariate values, using the Causal Forest method for each outcome according to the following steps:

1. We estimate regression forests to predict the outcome and treatment with 50,000 trees.
2. These predictions are used to form debiasing weights which are used in an initial Causal Forest consisting of 50,000 causal trees, estimated using the *causal_forest* function in version 2.1.0 of the *grf* package for R (Tibshirani et al., 2020), after tuning all hyperparameters.
3. We retain those variables whose importance in determining splits within the initial Causal Forests’ trees was above 20 % of the mean importance (as recommended by the package authors) and re-estimate the forest as described in step 2 to obtain the final Causal Forest.
4. We use this Causal Forest to estimate the CATE for each individual, along with their standard errors as described above.
5. We aggregate estimates for groups of individuals using AIPW.

4.3.3. Falsification test

For those children who had received the measles vaccine prior to study entry, uptake of this vaccine could not have been influenced by the intervention. This provides a falsification test for the machine learning estimation approach since detecting a substantive effect on measles vaccine uptake for this group would indicate confounding due to model misspecification. By contrast, if the estimates for the model were close to zero, this would provide support for the estimation approach.

4.4. Exploring heterogeneity of treatment effects

To understand the heterogeneity of treatment effects, we group individuals by quartile of their estimated treatment effects, and then explore the characteristics of those that benefit most (with largest 25 % of individual CATEs) and least (with smallest 25 % of individual CATEs), first for DPT3 and then separately for measles. We also report conditional average treatment effects for all the subgroups, together with appropriate measures of precision, albeit caution should be exercised given the reduced statistical power associated with smaller subgroups.

4.5. Sensitivity of estimates to sparsity

Given concerns regarding the small sample sizes, we re-estimate effects using Shrinkage Bayesian Causal Forest (Caron et al. 2022)

⁴ Results are robust to tuning various subsets of the hyperparameters.

⁵ Variables whose importance scores were below 0.2 times the mean variable importance score were excluded.

as a sensitivity analysis. This approach uses a Dirichlet prior over the splitting probabilities, in addition to the priors used in the Bayesian Additive Regression Trees for the original Bayesian Causal Forest (Hahn et al. 2020) that induce sparsity in the estimation of prognostic and moderating effects. The Shrinkage Bayesian Causal Forest was implemented using the SparseBCF package in R using default settings and 5000 Markov Chain Monte Carlo iterations (with 10,000 burn-in iterations).

5. Results

Table 1 shows baseline balance between the comparison groups of study participants followed-up at months seven and 30 month, respectively. The baseline characteristics are similar between treatment arms at both time points. In particular, differences between treatment groups remain small at 30 months when there was more, albeit still limited, loss to follow-up.⁶ In the follow-up survey at 30 months, the study team also collected data on sibling children born after the intervention. As Table A1 shows, amongst this subsample of study households, balance between the treatment groups was reasonably good, consistent with the notion that the intervention is unlikely to have influenced subsequent fertility decisions. On average siblings were approximately 26 months younger than the index child. We additionally note that the characteristics of the subsample with a younger sibling differ somewhat from the full sample, reflecting the fact that the decision to have another child is likely influenced by household factors.

5.1. Persistence of effects

Table 2 presents the effects of the intervention for our outcomes of interest for the index child by follow-up period. We report both unadjusted and adjusted estimates. We also report Causal Forest and Shrinkage Bayesian Causal Forest estimates of the ATE, for the purposes of comparison and to benchmark the individual treatment effects presented later. For measles vaccine uptake, we report effect estimates from two samples: the full sample and the subsample of children who were not already vaccinated at baseline. The unadjusted results are that the intervention increased uptake of DPT3 vaccination by 14.6 (95 % CI: 7.3, 21.9) percentage points at seven months follow-up and 15.2 (95 % CI: 7.4, 23.0) percentage points at 30 months follow-up, with little evidence of a reduction in the effect on DPT3 vaccination over time ($p = 0.567$). Although there was an increase in coverage of DPT3 over time in the control group, the results suggest that this increase did not reflect the control group “catching-up”, as a similar increase over time occurred in the intervention group. The corresponding results for measles were of an increase in vaccination uptake of 22.0 (95 % CI: 14.3, 29.7) percentage points at seven months follow-up and 22.2 (95 % CI: 14.6, 29.8) percentage points at 30 months follow-up. There was again little evidence of a reduction in the effect by follow-up period ($p = 0.703$).⁷

Results were qualitatively similar for both versions of the intervention, albeit effects were somewhat larger when the information was negatively framed (Supplement Table A4). Effects tended to be larger among those for whom vaccination status was self-reported, albeit this group differed in a number of ways (e.g. age of child, vaccination history) from those for whom a vaccination card was available making differences in effects challenging to interpret.

We formally assessed the potential impact of attrition using Lee bounds (Lee, 2009) implemented using the user-written command *leebounds* for the Stata software package as described in Tauchmann (2014). The group (treated/control) that suffers less from attrition is trimmed from above or below to create similar attrition in both groups and then the differential between the groups' outcomes yields the lower and upper bound. We then estimate from the data which treatment group is subject to the higher probability of selection. The estimated treatment-effect bounds at 30 months are narrow for both DPT3 (15.0 (95 % CI: 7.0 to 22.0) to 15.6 (95 % CI: 7.3 to 23.8) percentage points), and for measles (21.9 (95 % CI: 13.9 to 29.8) to 22.5 (95 % CI: 14.6 to 30.3) percentage points).

For both vaccinations, the results were similar after adjustment for any residual baseline differences between groups for both the estimates from the OLS regression and the two Causal Forest approaches (Table 2). When we focus on the subsample of children who had not received the measles vaccine at baseline, we see the estimates of effect are larger (27.5 percentage points at seven months follow-up and 28.2 percentage points at 30 months).⁸ Such estimates give a sense of the effect had the intervention targeted children who were not already vaccinated against measles.

Table 3 presents the results for the sibling sample for which outcomes were measured at 30 months follow-up. There was some evidence that the intervention had an effect on uptake of either the DPT3 or measles vaccine, although the estimated effects while clinically meaningful did not meet criteria of statistical significance at 5 % or 10 % levels. Due to the small sample size, confidence intervals are wide such that we cannot rule out large effects that would be regarded as meaningful from a public health perspective. Note that to detect a 9 percentage point difference between the groups would have required a sample size of 942, based on the observed incidence in the control group of 32.48 % and in the treated group of 41.48 %, and the observed ratio of units in the treated and control arms in the sibling sample ($= 176 / 117$) (Rosner, 2011, page 381).

⁶ While the differences between groups are modest, it is helpful to note that in the Causal Forest estimation, units are reweighted by their propensity for treatment, to improve balance within each estimated ‘leaf’ of the trees in the forest.

⁷ While the evidence suggest an absence of ‘catching-up,’ it is important to note that we cannot conclusively rule this out due to our sample size. As shown in the note accompanying Table 2, catching up by 4.9 (DPT3) and 6 (measles) percentage points would be consistent with the data.

⁸ Table A2 shows that children in both treated and control arms that had been vaccinated for measles before the baseline (compared to those not vaccinated for measles) tended to be older, were much more likely to have received other vaccinations (BCG and 1st & 2nd doses of DPT) and have mothers that were slightly more likely to know causes of tetanus but less likely to know symptoms or prevention methods of tetanus. Some difference in terms of their closest health facility are also observed.

Table 2

Absolute difference in probability of DPT3 and measles vaccination uptake with versus without the intervention at seven and 30 months follow-up.

	DPT3		Measles (full sample)		Measles (subsample)	
	At seven months follow-up	At 30 months follow-up	At seven months follow-up	At 30 months follow-up	At seven months follow-up	At 30 months follow-up
	ATE (95 % CI)	ATE (95 % CI)	ATE (95 % CI)	ATE (95 % CI)	ATE (95 % CI)	ATE (95 % CI)
Unadjusted difference	0.146 (0.073, 0.219)	0.152 (0.074, 0.230)	0.22 (0.143, 0.297)	0.222 (0.146, 0.298)	0.275 (0.193, 0.356)	0.282 (0.198, 0.367)
Adjusted difference	0.168 (0.099, 0.236)	0.174 (0.101, 0.247)	0.247 (0.182, 0.313)	0.244 (0.179, 0.309)	0.295 (0.219, 0.371)	0.292 (0.217, 0.367)
Causal Forest	0.145 (0.090, 0.200)	0.152 (0.091, 0.213)	0.233 (0.178, 0.287)	0.234 (0.180, 0.288)	0.298 (0.225, 0.351)	0.287 (0.225, 0.348)
Shrinkage Bayesian Causal Forest	0.138 (0.061, 0.213)	0.123 (0.029, 0.208)	0.235 (0.0164, 0.306)	0.224 (0.158, 0.287)	0.276 (0.191, 0.360)	0.277 (0.200, 0.352)
Mean of control group	0.281	0.353	0.417	0.545	0.308	0.451
Number of observations	706	674	706	674	590 ⁱ	563 [†]

Notes: Table reports the average treatment effect of the intervention at seven months and 30 months follow-up. Confidence intervals in parentheses. Unadjusted differences were estimated using OLS. Adjusted differences were estimated using OLS and include the same baseline covariates as used in the Causal Forest approach. Causal Forest estimates were obtained by aggregating individual level CATE estimates using AIPW as described in the text. The measles subsample excludes individuals that had received the measles vaccine at baseline.

For the adjusted regression analysis, we estimate differences in the ATEs at 30 months versus at seven months using pooled regression with interaction terms for each period with the treatment indicator, after restricting the sample to those present in both waves. For DPT3 the difference in ATEs was 0.002 (−0.049, 0.052), for measles in the full sample it was 0.006 (−0.049, 0.061) and after excluding those that had already received the measles vaccine the difference was 0.004 (−0.060, 0.069). We could not reject the null that the ATE at 30 months was at least as large as at 7 months (i.e. there was no catching up) with p-values of 0.523, 0.580 and 0.555 respectively. The corresponding p-values for the unadjusted analysis were 0.567, 0.703, and 0.672. Assessing paired differences is more challenging in the context of (Bayesian) Causal Forests where analyses are conducted separately by follow-up period making it difficult to account for correlation within individuals.

Table 3

Absolute difference in probability of vaccination uptake with versus without the intervention amongst younger siblings at 30 months follow-up.

	Younger sibling sample (30 months follow-up)	
	DPT3 ATE (95 % CI)	Measles ATE (95 % CI)
Unadjusted difference	0.090 (−0.022, 0.203)	0.064 (−0.054, 0.182)
Adjusted difference	0.081 (−0.030, 0.192)	0.034 (−0.074, 0.143)
Causal Forest	0.060 (−0.034, 0.154)	0.021 (−0.058, 0.100)
Shrinkage Bayesian Causal Forest	0.032 (−0.041, 0.145)	0.016 (−0.046, 0.097)
Mean for control group	0.325	0.470
N	293	293

Notes: Table reports the average treatment effect of the intervention on vaccination uptake amongst siblings at 30 months follow-up. Confidence intervals in parentheses. Unadjusted differences were estimated using OLS. Adjusted differences were estimated using OLS and include baseline covariates. Causal Forest estimates were obtained by aggregating individual level CATE estimates using AIPW as described in the text. The measles subsample excludes individuals that had received the measles vaccine at baseline.

5.2. Heterogeneity of intervention effects

Figs. 1 and 2 show the individual level effect estimates for DPT3 and measles vaccination, respectively, with individuals ranked in order of estimated treatment effects, with effects reported as absolute difference in the probability of vaccination uptake under treatment versus under control.⁹ The measles vaccination sample excludes those vaccinated prior to study entry. The confidence intervals suggest that the intervention had a statistically significant effect on DPT3 vaccine uptake for 56 % of children, and on measles vaccine uptake for 100 % of children. None of the point estimates are below zero, suggesting that the intervention did not reduce the probability of being vaccinated for any of the individuals sampled.

In Table 4, we split the sample into quartiles according to the estimated CATE for each outcome. We report the CATE for the least affected and the most affected subgroups. The magnitude for the estimated mean differences in the CATEs for the most versus least affected quartiles are large, 15.7 percentage points (DPT3), and 19.9 percentage points (measles).

⁹ Figs. A1 and A2 display the corresponding figures at seven months. For DPT3 uptake, there is a large group of approximately 200 children for whom the effects tend to be smaller and imprecisely estimates at seven months, although this is less evident at 30 months. Inspection of the data reveals that all of these children were aged between 0 and 6 months at the baseline.




Fig. 1. Individual level CATE estimates based on Causal Forests for DPT3 vaccination at 30 months follow-up ranked by magnitude of CATE. Note: Individuals ranked from minimum to maximum HTE; Black line = point estimate, Grey line = 95 % Confidence interval, Blue line = zero effect

To characterise this heterogeneity, [Table 5](#) reports summary statistics for the characteristics of individuals in the groups that are estimated to benefit least (quartile 1) and most (quartile 4) from the intervention. We first examine the DPT3 vaccination results. Children in the quartile who benefited most from the intervention were older and more likely to have received other vaccinations, including the first and second doses of DPT, the BCG vaccine, and the measles vaccine, compared to those who benefited least. Those affected most had mothers with lower levels of baseline knowledge about the causes, symptoms, and methods of prevention of tetanus, and their closest health facility was more likely to be a government primary care facility. There was little evidence that the mother's perception of the effectiveness of the tetanus vaccination influences the effectiveness of the intervention. Turning to the measles vaccination results, the estimates of heterogeneity in the interventions effectiveness was associated with fewer baseline characteristics. Again, children who benefited most were older and more likely to have received the first and second doses of DPT. For the other baseline variables considered, differences in the means of the most and least affected groups were moderate or small in magnitude, with high levels of uncertainty, and the differences in means for these characteristics between the 'most' and 'least' affected groups were not statistically significant. The estimated subgroup effects, generated by aggregating the individual effect estimates for each of the subgroups using AIPW, reveals a similar pattern of results to those reported in [Table 5](#), i.e. the magnitude of effect for age group,

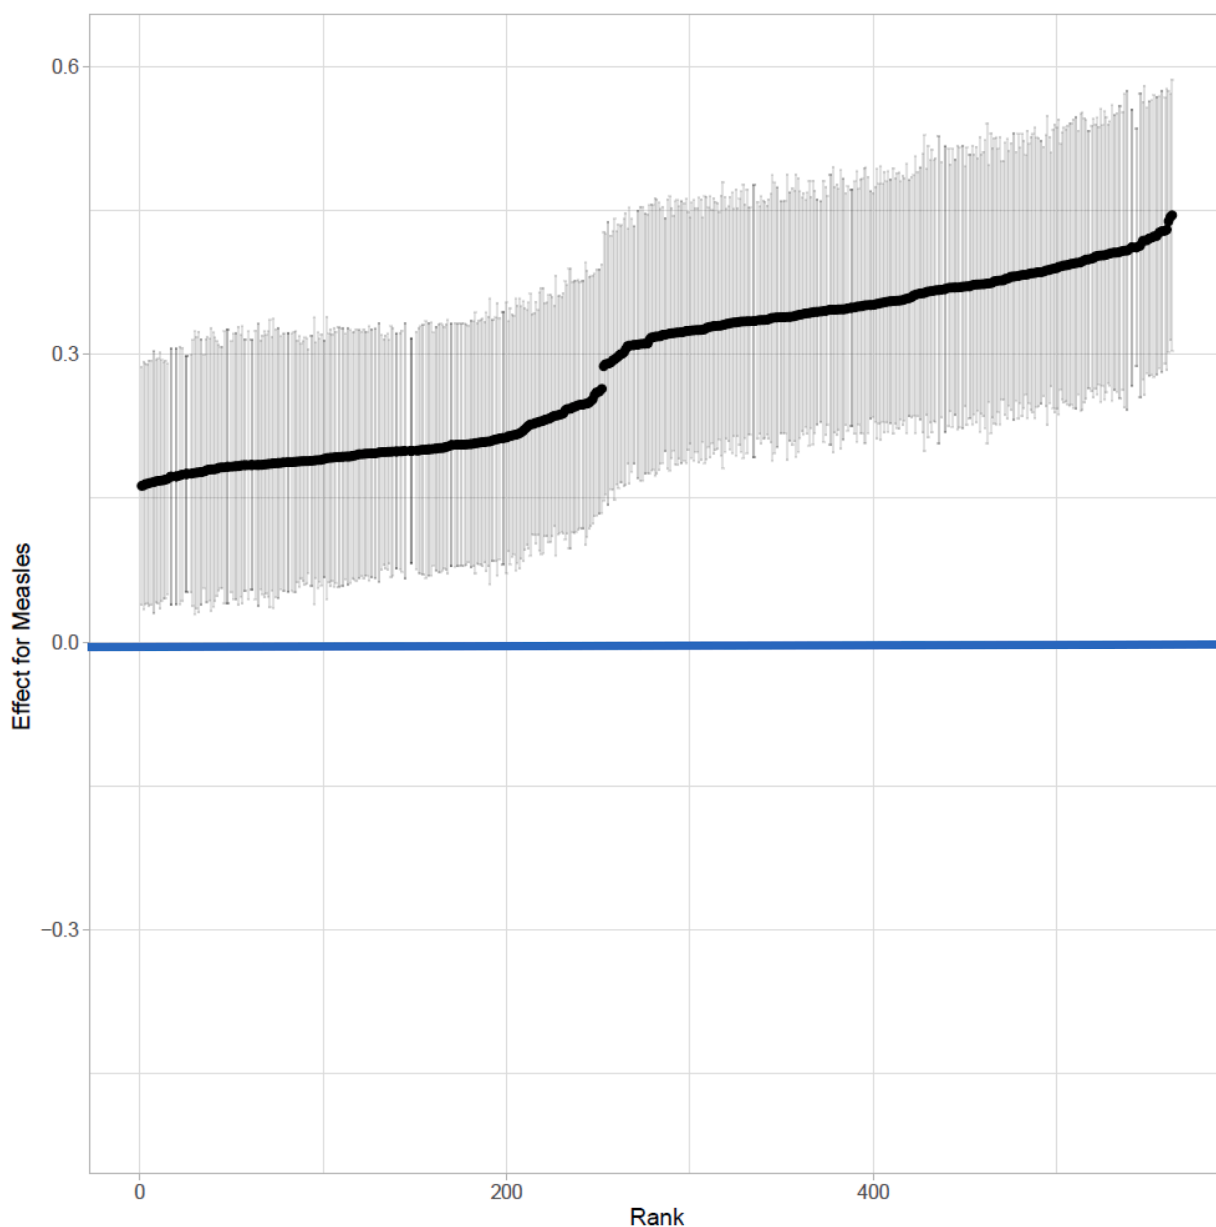


Fig. 2. Individual level CATE estimates based on Causal Forests for measles vaccination at 30 months follow-up ranked by magnitude of CATE. Note: Individuals ranked from minimum to maximum HTE; Black line = point estimate, Grey line = 95 % Confidence interval, Blue line = zero effect.

which was a pre-specified variable were fairly large and precisely estimated, whereas for other baseline variables the differential effects of the intervention versus comparator were of small or moderate magnitude and were imprecisely estimated (see Fig. A4 and Fig. A5). Table A3 reports results that are similar based on the Shrinkage Bayesian Causal Forest estimates.

5.3. Falsification test

For the falsification test, we report the individual treatment effects of the intervention on the uptake of measles vaccination for the subgroup of individuals that had already received the measles vaccine prior to study entry. Fig. A3 shows that none of these individual level effect estimates significantly differ from zero, and the overall ATEs for this subgroup is of small magnitude, -5.0 percentage points (95 % CI: $-12.2, 2.3$; $N = 111$). The analysis thus passes our falsification test, increasing the plausibility of estimates obtained using the Causal Forest approach.

Table 4

Absolute difference in probability of vaccination uptake with versus without the intervention for the least and most affected groups of participants at 30 months follow-up.

	DPT3 vaccination 25 % Least affected CATE (95 % CI) [p- value]	25 % Most affected CATE (95 % CI) [p- value]	Difference (95 % CI) [p-value]	Measles vaccination 25 % Least affected CATE (95 % CI) [p- value]	25 % Most affected CATE (95 % CI) [p- value]	Difference (95 % CI) [p-value]
Causal Forest	0.084 (0.035, 0.134) [0.001]	0.242 (0.195, 0.288) [<0.001]	0.157 (0.089, 0.225) [<0.001]	0.211 (0.173, 0.249) [<0.001]	0.410 (0.363, 0.458) [<0.001]	0.199 (0.138, 0.260) [<0.001]
Shrinkage Bayesian Causal Forest	0.055 (0.052, 0.058) [0.001]	0.174 (0.171, 0.177) [<0.001]	0.119 (0.115, 0.123) [<0.001]	0.172 (0.169, 0.175) [<0.001]	0.357 (0.354, 0.360) [<0.001]	0.185 (0.181, 0.190) [<0.001]
Control mean	0.655	0.170		0.820 [†]	0.179	
N observations	169	176		141	140	

[†] While it is counterintuitive that sum of the control mean and the effect for those least affected exceeds 1, we attribute this to our use of AIPW.

Notes: Table reports the average treatment effect of the intervention at 30 months follow-up for the 25 % of individuals least and most affected by the intervention (based on estimated individual level CATEs). Confidence intervals in parentheses. The measles subsample excludes individuals that had received the measles vaccine at baseline.

Table 5

Baseline characteristics of those quartiles of children for whom the intervention had least versus most effect on the uptake of DPT3 and measles vaccination at 30 months based on Causal Forest estimates after Augmented Inverse Probability Weighting.

	DPT3 vaccination			Measles vaccination		
	25 % Least affected CATE (95 % CI)	25 % Most Affected CATE (95 % CI)	Difference (95 % CI) [p-value]	25 % Least affected CATE (95 % CI)	25 % Most Affected CATE (95 % CI)	Difference (95 % CI) [p-value]
Age of child (months)	1.935 (1.053, 2.817)	12.070 (11.186, 12.954)	10.135 (8.886, 11.384) [<0.001]	1.894 (1.192, 2.596)	14.760 (14.056, 15.464)	12.866 (11.872, 13.860) [<0.001]
Received DPT 1st dose	0.462 (0.395, 0.529)	0.899 (0.832, 0.966)	0.437 (0.343, 0.531) [<0.001]	0.539 (0.463, 0.615)	0.871 (0.795, 0.947)	0.332 (0.224, 0.440) [<0.001]
Received DPT 2nd dose	0.160 (0.091, 0.229)	0.673 (0.604, 0.742)	0.513 (0.416, 0.610) [<0.001]	0.284 (0.211, 0.357)	0.636 (0.563, 0.709)	0.352 (0.249, 0.455) [<0.001]
Received BCG vaccine	0.823 (0.768, 0.878)	0.923 (0.868, 0.978)	0.100 (0.022, 0.178) [0.012]	0.865 (0.802, 0.928)	0.886 (0.823, 0.949)	0.021 (-0.068, 0.110) [0.643]
Received measles vaccine	0.041 (-0.014, 0.096)	0.238 (0.183, 0.293)	0.197 (0.119, 0.275) [<0.001]	- -	- -	- -
Perception of tetanus vaccination efficacy	6.923 (6.623, 7.223)	7.185 (6.885, 7.485)	0.262 (-0.162, 0.686) [0.226]	7.638 (7.313, 7.963)	7.521 (7.196, 7.846)	-0.117 (-0.577, 0.343) [0.618]
Mother knows a cause of tetanus	0.521 (0.448, 0.594)	0.208 (0.135, 0.281)	-0.313 (-0.416, -0.210) [<0.001]	0.418 (0.336, 0.500)	0.350 (0.268, 0.432)	-0.068 (-0.184, 0.048) [0.252]
Mother knows a symptom of tetanus	0.101 (0.060, 0.142)	0.012 (-0.029, 0.053)	-0.089 (-0.147, -0.031) [<0.002]	0.071 (0.024, 0.118)	0.064 (0.017, 0.111)	-0.007 (-0.074, 0.060) [0.837]
Mother knows a prevention method of tetanus	0.574 (0.501, 0.647)	0.226 (0.153, 0.299)	-0.348 (-0.451, -0.245) [<0.001]	0.390 (0.310, 0.470)	0.336 (0.254, 0.418)	-0.054 (-0.169, 0.061) [0.358]
Closest facility is government PHC facility	0.840 (0.789, 0.891)	0.964 (0.913, 1.015)	0.124 (0.052, 0.196) [<0.001]	0.851 (0.794, 0.908)	0.879 (0.822, 0.936)	0.028 (-0.052, 0.108) [0.495]

Notes: Table reports the mean and standard deviations of baseline covariates for the 25 % of individuals least and most affected by the intervention (based on estimated individual level CATEs obtained using Causal Forest with Augmented Inverse Probability Weighting. See [Table A3](#) for corresponding results using Shrinkage Bayesian Causal Forests). Confidence intervals in parentheses. The measles subsample excludes individuals that had received the measles vaccine at baseline. Differences between the groups were assessed using independent samples *t*-tests.

6. Discussion

Immunisation is one of the most effective health technologies to prevent mortality and disability, yet millions of children are not fully vaccinated despite the fact that immunisation services are available for free in many countries. In this study, we asked new questions of a previously published randomised controlled trial, specifically we asked whether the effects of a brief information intervention, implemented through door-to-door visits across 180 villages in Uttar Pradesh, India, persisted, and whether there was heterogeneity according to characteristics of the child, the mother's perceptions, and the health care context. Our findings, based on new data collected 30 months after the intervention was delivered, complement those reported in a previous analysis of earlier trial data (Powell-Jackson et al., 2018). In a broad sense, our study contributes to the literature on demand-side interventions for immunisation uptake in low- and middle-income countries, providing novel insights on the persistence and heterogeneity of effects (Banerjee et al., 2020, 2021; Gibson et al., 2017; Johri et al., 2015).

Our findings show that the intervention was highly effective in raising uptake of both DPT3 and measles immunisation over a sustained period. There was no evidence that the early effects of the intervention were short-lived, in the sense that it brought forward vaccinations that would have happened anyway. The findings indicate that those who did respond to the intervention, did so promptly, given that the difference in the point estimates of the treatment effects in the seven to 30 month window were negligible. The magnitude of the difference in outcomes between the intervention and control strategy did not meaningfully narrow over time, which also implies that there was probably no effective immunisation catch-up programme in place. Our 30 month follow-up period was sufficiently long to be confident that the control group is not likely to ever "catch up" with the intervention group such that the long-lasting health benefits of vaccination should be realised. Unlike interventions targeting reversible behaviours such as smoking cessation, (lack of) persistence in effects on vaccine uptake is likely to be driven by changes in the control group's behaviour, making it challenging to identify testable mechanisms underlying persistence. To examine the question of whether the intervention had a sustained impact on parental behaviour, we used data collected on the vaccination outcomes of siblings who were born after the time of the intervention. There was no strong evidence that the intervention may have had an effect on DPT3 uptake for siblings, with the sample size too small to make firm conclusions. There is a small literature on whether temporary incentives can lead to healthy habit formation, such as smoking cessation and exercise (Charness and Gneezy, 2009; Volpp et al., 2009), and a rich body of theoretical work on how to maintain behaviour change (Kwasnicka et al., 2016). There is, however, a need for more evidence on whether one-off health education intervention can lead to sustained changes in behaviour in the uptake of health care technologies.

Our exploration of heterogeneity, drawing on modern machine learning methods, found that there was substantial variation in the effect of the intervention. For those who benefited most, the intervention effects were large. For those who benefited least, the point estimates were modest and, importantly, there was no evidence the intervention caused harm by reducing the chances of children being vaccinated. While this finding is to be expected, given the light touch nature of the intervention, it stands in contrast to another study in India, where a combination of small incentives, reminders and persuasion was found to reduce immunisation rates in some villages, possibly because the interventions crowded out existing intrinsic motivation of parents to vaccinate their children (Chernozhuikov et al., 2018). For DPT3, we found that a range of characteristics were associated with treatment effect heterogeneity. For measles vaccine uptake, fewer characteristics were associated with variation in the treatment effect.

For both vaccinations a general finding was that for older children, and those who had received two DPT doses, which were both specified a priori and supported by prior reasoning, the intervention led to relatively large increases in uptake. The prior reasoning was that for those children who, at baseline, were further along in their immunisation schedule (proxied by age) and hence closest to crossing the threshold to achieving DPT3 status, 'a soft nudge' from the intervention was sufficient to encourage those with enough previous doses to complete the DPT schedule. By contrast, this nudge from the intervention was insufficient to encourage those with less vaccine history, including younger children, to reach DPT3 status. Furthermore, the parents of children with fewer previous immunisations likely had less trust in the health system and were therefore less receptive to the intervention.

These results pertaining to knowledge are largely consistent with the assumed mechanism for the intervention. The individuals with worse understanding of tetanus prior to the intervention had the greatest capacity to acquire new knowledge about the benefits of vaccination. However, it does not appear that the effect of the intervention was mediated by changing perceptions of vaccine effectiveness, since those most affected by the intervention did not have more inaccurate perceptions at baseline. If anything those most affected by the intervention had relatively high perceptions of vaccine effectiveness at baseline. The finding that the most affected individuals were more likely to have a government primary care facility as their nearest type of facility points towards the importance of context. It suggests that stimulating demand was more effective for those with better access to the publicly-provided immunisation services. This mirrors other studies that have shown demand-side interventions are more effective when the supply-side is in place (Powell-Jackson et al., 2015). For measles immunisation uptake, those most affected were older and more likely to have had previous doses of DPT. These results are consistent with recommended age for measles immunisation of nine months or later.

The study had a number of limitations. First, the analysis of heterogeneity would have benefited from a richer set of baseline covariates such as income, wealth and education levels. Data on various dimensions of socioeconomic status would have helped provide evidence of the impact of the intervention on inequalities in immunisation uptake. Whether gender is associated with the effectiveness of the intervention is an especially important question in this particular setting, in light of the gender inequalities and strong gender discrimination in North India (Dhar et al., 2022). Second, as the intervention was randomised at the individual-level there is a risk of contamination in the effect of the intervention, in particular as women in the intervention group may have relayed information to counterparts in the control group. We do not have direct evidence to counter this concern, but the information was delivered in private and the fact we see large effects on knowledge indicates that contamination was not so severe as to balance out knowledge between treatment and control (Powell-Jackson et al., 2018). We note that the (large) estimates of effect on immunisation

uptake would be biased downwards in the presence of contamination. Third, while the study used a principled approach to explore heterogeneity that combines pre-specification and prior reasoning of potentially important subgroups, with machine learning approaches that avoids using the same subsample to define subgroups as to estimate effects, general concerns about the need for careful interpretation of subgroup results must be acknowledged (Brankovic et al., 2019; Brookes et al., 2001; Kent et al., 2010; Wang et al., 2007; Wang et al., 2021). It must be recognised that in most RCTs, the sample sizes are such that subgroup-level estimates will be imprecise, and that while those subgroup results that are 'post hoc' and not aligned with prior reasoning, may generate useful hypotheses for further research, they are unlikely to lead to direct policy recommendations. While findings here such as those regarding differences in effect by age and vaccination history are in line with prior reasoning, the analysis is underpowered to provide definitive policy recommendations. Value of information analysis can be informative as to whether further research is beneficial to inform policy in such contexts. Fourth, as previously mentioned, the number of siblings interviewed was small, which meant the analysis of the sibling sample was underpowered.

Machine learning methods are attractive given their ability to model data flexibly while reducing the risk of overfitting. Nonetheless they have some limitations. They may be viewed as more complex and less transparent than more familiar parametric methods. They can be sensitive to the choice of hyperparameters (such as the number of trees in causal forest) and inference can be more challenging. This is true also when comparing effects over time. While it is straightforward to compare estimates within the regression framework using paired *t*-tests to recognise we have repeated measures for the same individual, it is not straightforward to account for correlation between waves when comparing the Causal Forest estimates. With the increasing adoption of these approaches, methodological advances and the availability of guidance for their use (Padula et al., 2022) machine learning methods are likely to be recognised as a useful complement to existing approaches.

While here the number of variables is fairly modest, causal forest allows us to account for possible high dimensional interactions between the covariates while avoiding overfitting. For instance even with a modest number of covariates it is challenging to correctly specify the non-linear relationship between child age, vaccination history and mothers knowledge/beliefs regarding vaccination, and vaccination uptake or the effects of the intervention. Causal Forest allows us to account for such non-linearity while avoiding overfitting. However, we view machine learning methods as a complement to, rather than substitute for, existing parametric approaches - to the extent that effects are similar across methods this offers some reassurance regarding model misspecification.

From a policy perspective, there are three key messages. First, it will be reassuring for policymakers that the intervention increased immunisation uptake for the majority of households, there was no evidence that the intervention discouraged any women from having their children immunised, and the effects were maintained over time. There was considerable heterogeneity in the responses to the intervention and policymakers will need to consider carefully what this means for targeting and for the deployment of other complementary interventions. Recent evidence from a study in India suggests that various interventions implemented in combination are more effective in raising immunisation rates than when implemented in isolation (Banerjee et al., 2021). Second, the results suggest that the intervention could be better targeted at children in a certain age range. Not only was age, a pre-specified subgroup, found to modify the relative effectiveness of the intervention, it can easily be observed, and it is feasible to target according to age. One strategy may be to narrow the focus on families with children in the range of 6 to 18 months and carry out repeated visits to villages as new children enter this age cohort. But whether this proves more cost-effective is uncertain because the cost per household visit would likely increase. Third, if awareness and knowledge of vaccination amongst mothers are a key binding constraint, as the results in this study suggest, there may be alternative strategies that can deliver the information with the same fidelity, but at a lower cost than our intervention. Such strategies could involve greater use of community health workers who are already present in villages. It may be tempting to think that mobile phone technology could offer an alternative delivery platform. Indeed, improving the digital health infrastructure is a strategic priority of the Indian government. To-date, mobile phones have largely been used to deliver targeted text message reminders, for which the evidence on effectiveness is mixed (Banerjee et al., 2021; Mekonnen et al., 2019). With increasing ownership of smart phones and access to social media platforms, there is scope for richer (trustworthy) content to be delivered that may be more effective.

CRediT authorship contribution statement

Stephen O'Neill: Writing – review & editing, Writing – original draft, Software, Methodology, Formal analysis, Conceptualization. **Richard Grieve:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Kultar Singh:** Writing – review & editing, Conceptualization. **Varun Dutt:** Writing – review & editing, Conceptualization. **Timothy Powell-Jackson:** Writing – review & editing, Writing – original draft, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare no other conflicts of interest.

Acknowledgements

We would like to thank participants in the organized session on Exploring Heterogeneity of Treatment Effects at the International Health Economics Association 2021 for helpful comments and suggestions.

Appendix A

Further details on intervention

The intervention was brief and took about 10 min to deliver. Mothers were provided with information on the benefits of the tetanus vaccine by field staff through door-to-door visits. Field staff were trained to follow a standardised script to ensure fidelity in terms of the information given to mothers. The script described the causes and symptoms of tetanus, possible health consequences, the individual health benefit of the combination DPT, and the wider community benefits associated with herd immunity. Visual aids were shown to ensure the information was accessible to illiterate women, and a Hindi leaflet containing the information was left with the mother. There was a short question and answer session to ensure mothers had understood the information. These activities were brief, taking about 10 min to deliver.

We tested two versions of the script that varied how the information was framed. The first script framed information on tetanus vaccination as gains, emphasising that the child would be less likely to get tetanus and more likely to be healthy once vaccinated. The second script framed the information as a loss, highlighting that an unvaccinated child would be more likely to get tetanus and suffer the health consequences of the disease. The framing of information was a key aspect of the original study. However, because there were no significant or substantive differences in the overall effect with respect to framing, we pool the two information intervention groups in this paper.

In each round of data collection, we captured the immunisation status of the child and the mother's knowledge of the causes of, symptoms of, and prevention methods against tetanus. We followed standard methods to assess immunisation status [DHS and NRHS], using the vaccination card as the primary source of information and, if not available, self-reports from the mother. The interview also included 'games' with chickpeas designed to elicit women's perceptions of the efficacy of tetanus and measles vaccination, alongside verification questions to gauge understanding of these games. We used tablets and computer-assisted personal interviewing to collect the data, and field staff were blinded to randomised assignment in the two follow-up surveys.

Table A1
Descriptive statistics on baseline covariates for younger sibling sample (30 months).

	Full sample (30 months)		Younger sibling sample (30 months)	
	All Mean (Std. Dev.)	All Mean (Std. Dev.)	Treated Mean (Std. Dev.)	Control Mean (Std. Dev.)
Age of index child at 2nd follow up (months)	41.24 (7.69)	41.18 (7.61)	41.62 (7.94)	40.51 (7.08)
Age of sibling child (months)	N/A	15.27 (9.78)	15.86 (10.05)	14.38 (9.32)
Perception of tetanus vaccination efficacy (index)	7.27 (2.00)	7.40 (1.82)	7.41 (1.80)	7.38 (1.85)
	%	%	%	%
Mother knows a cause of tetanus	43.0	41.6	42.6	40.2
Mother knows a symptom of tetanus	8.3	7.2	7.4	6.8
Mother knows a prevention method of tetanus	40.7	39.6	40.9	37.6
Closest health facility: Government	86.3	82.5	83.5	81.2
Closest health facility: Private/Other	13.7	17.5	16.5	18.8
Number of observations	674	293	176	117

Notes: Table reports the mean and standard deviations of baseline covariates for the full sample, and for younger sibling subsample, and by treated and control groups, at 30 months follow-up.

Table A2
Comparison of characteristics of individuals present at seven months that were and were not vaccinated for measles at baseline.

	Individuals present at seven months (N = 706)			
	Not already vaccinated for measles		Already vaccinated for measles (excluded from analysis)	
	Treated Mean (Std. Dev.)	Control Mean (Std. Dev.)	Treated Mean (Std. Dev.)	Control Mean (Std. Dev.)
Age of index child at baseline (months)	9.4 (7.3)	9.7 (7.3)	15.6 (8.5)	12.8 (7.5)
Perception of tetanus vaccination efficacy (index)	7.2 (2.0)	7.4 (1.9)	7.4 (2.0)	7.2 (2.1)
	%	%	%	%
Received DPT 1st dose	57.0	65.1	96.1	92.5
Received DPT 2nd dose	31.9	39.0	77.6	85.0
Received BCG vaccine	79.7	85.6	96.1	100
Mother knows a cause of tetanus	40.8	46.2	43.3	50.0
Mother knows a symptom of tetanus	9.4	9.2	2.6	5.0

(continued on next page)

Table A2 (continued)

	Individuals present at seven months (N = 706)			
	Not already vaccinated for measles		Already vaccinated for measles (excluded from analysis)	
	Treated Mean (Std. Dev.)	Control Mean (Std. Dev.)	Treated Mean (Std. Dev.)	Control Mean (Std. Dev.)
Mother knows a prevention method of tetanus	42.0	45.1	27.6	25.0
Closest health facility: Government	85.8	86.6	92.1	82.5
Closest health facility: Private/Other	14.2	13.4	7.9	17.5
Number of observations	395	195	76	40

Notes: Table reports the mean and standard deviations of baseline covariates for individuals there were and were not already vaccinated for measles at baseline by treated and control group at seven months follow-up.

Table A3

Baseline characteristics of those quartiles of children for whom the intervention had least versus most effect on the uptake of DPT3 and measles vaccination at 30 months based on Shrinkage Bayesian Causal Forest estimates.

	DPT3 vaccination			Measles vaccination		
	25 % Least affected CATE (95 % CI)	25 % Most Affected CATE (95 % CI)	Difference (95 % CI) [p-value]	25 % Least affected CATE (95 % CI)	25 % Most Affected CATE (95 % CI)	Difference (95 % CI) [p-value]
Age of child (months)	4.982 (4.042, 5.922)	14.315 (13.373, 15.258)	9.333 (8.002, 10.664) [<0.001]	1.525 (0.857, 2.193)	17.971 (17.301, 18.642)	16.447 (15.500, 17.393) [<0.001]
Received DPT 1st dose	0.337 (0.281, 0.394)	0.946 (0.89, 1.003)	0.609 (0.529, 0.689) [<0.001]	0.376 (0.294, 0.458)	0.557 (0.475, 0.639)	0.181 (0.066, 0.297) [0.002]
Received DPT 2nd dose	0.142 (0.082, 0.202)	0.744 (0.684, 0.804)	0.602 (0.517, 0.687) [<0.001]	0.156 (0.087, 0.225)	0.300 (0.231, 0.369)	0.144 (0.047, 0.241) [0.004]
Received BCG vaccine	0.710 (0.656, 0.764)	0.946 (0.892, 1.001)	0.236 (0.159, 0.313) [<0.001]	0.830 (0.76, 0.899)	0.700 (0.63, 0.77)	-0.130 (-0.229, -0.031) [0.010]
Received measles vaccine	0.018 (-0.035, 0.07)	0.339 (0.286, 0.392)	0.322 (0.247, 0.396) [<0.001]	-	-	-
Perception of tetanus vaccination efficacy	0.740 (0.679, 0.8)	0.143 (0.082, 0.203)	-0.597 (-0.682, -0.511) [<0.001]	0.496 (0.416, 0.577)	0.329 (0.248, 0.41)	-0.168 (-0.282, -0.054) [0.004]
Mother knows a cause of tetanus	0.178 (0.136, 0.219)	0.006 (-0.036, 0.048)	-0.172 (-0.231, -0.112) [<0.001]	0.085 (0.039, 0.132)	0.086 (0.039, 0.132)	0.001 (-0.065, 0.066) [0.986]
Mother knows a symptom of tetanus	0.692 (0.632, 0.752)	0.113 (0.053, 0.173)	-0.579 (-0.664, -0.494) [<0.001]	0.433 (0.35, 0.515)	0.500 (0.417, 0.583)	0.067 (-0.050, 0.185) [0.259]
Mother knows a prevention method of tetanus	7.041 (6.739, 7.343)	7.488 (7.185, 7.791)	0.447 (0.019, 0.874) [<0.001]	7.241 (6.919, 7.563)	7.550 (7.227, 7.873)	0.309 (-0.147, 0.765) [0.183]
Closest facility is government PHC facility	0.722 (0.671, 0.773)	0.976 (0.925, 1.027)	0.254 (0.182, 0.326) [<0.001]	0.915 (0.856, 0.974)	0.771 (0.712, 0.831)	-0.143 (-0.227, -0.059) [0.001]


Notes: Table reports the mean and standard deviations of baseline covariates for the 25 % of individuals least and most affected by the intervention (based on estimated individual level CATEs). Confidence intervals in parentheses. The measles subsample excludes individuals that had received the measles vaccine at baseline. Differences between the groups were assessed using independent samples *t*-tests.

Table A4

Absolute difference in probability of vaccination uptake with versus without the intervention amongst younger siblings at 30 months follow-up by positive versus negative framing of information.


DPT3 vaccination	Positively framed Effect (95 % CI)	Negatively framed Effect (95 % CI)
Unadjusted	13.4 % (4.2 %, 22.5 %)	16.9 % (7.9 %, 25.9 %)
Adjusted	16.8 % (8.2 %, 25.3 %)	19.1 % (10.8 %, 27.4 %)
Measles vaccination		
Unadjusted	19.8 % (11.1 %, 28.6 %)	24.4 % (16.0 %, 32.8 %)
Adjusted	22.8 % (15.2 %, 30.3 %)	26.7 % (19.5 %, 33.9 %)

Notes: Table reports the average treatment effect of the intervention on vaccination uptake amongst siblings at 30 months follow-up by framing of information. Confidence intervals in parentheses. Unadjusted differences were estimated using OLS. Adjusted differences were estimated using OLS and include baseline covariates. The measles subsample excludes individuals that had received the measles vaccine at baseline.




Note: Individuals ranked from minimum to maximum HTE; Black line = point estimate, Grey line = 95% Confidence interval, Blue line = zero effect

Fig. A1. Individual level CATE estimates for DPT3 at seven months, ordered by magnitude.




Note: Individuals ranked from minimum to maximum HTE; Black line = point estimate, Grey line = 95% Confidence interval, Blue line = zero effect

Fig. A2. Individual level CATE estimates for Measles at seven months after excluding measles vaccinated at baseline, ordered by magnitude.



Note: Individuals ranked from minimum to maximum HTE; Black line = point estimate, Grey line = 95% Confidence interval, Blue line = zero effect

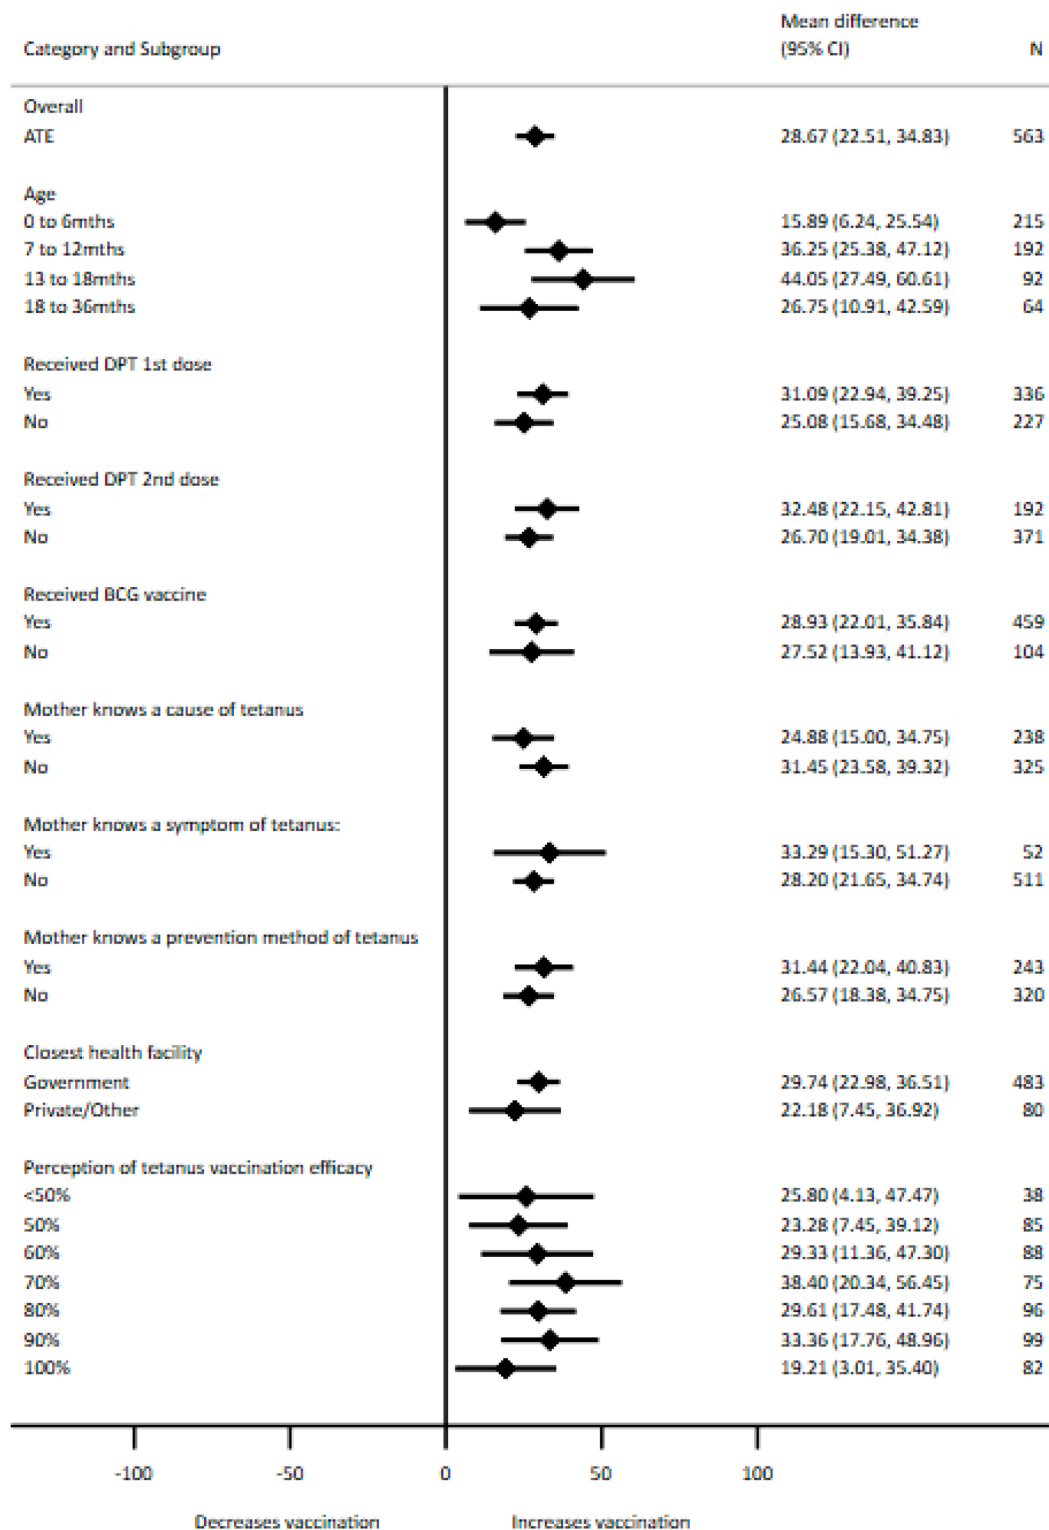
Fig. A3. Falsification test: Individual level CATE estimates for Measles at 30 months for those vaccinated for measles at baseline.



Notes: Individual level CATEs estimated using Causal Forests were aggregated for subgroups of interest using AIPW.

(caption on next page)

Fig. A4. Forest plot of CATE effects by subgroup: DPT3 (30 months).



Notes: Individual level CATEs estimated using Causal Forests were aggregated for subgroups of interest using AIPW.

Fig. A5. Forest plot of CATE effects by subgroup: Measles (30 months) (after excluding those vaccinated for measles at baseline).

Script

The mothers were addressed using the following script.

“Hello, my name is _____, I'd like to talk to you for a few minutes about an important health issue that you may have heard about. Please listen carefully to the information I am about to give you.” Then the information below was said to mothers depending on whether information was framed as a loss or as a gain.

Information leaflet 1: Information framed as a loss.

What is tetanus?

Tetanus is a serious disease that easily spreads from person to person. People of all ages can get tetanus but newborn babies are particularly vulnerable to tetanus – it is one of the leading causes of death among newborns in India causing almost 32,000 newborn deaths per year.

How does a person get tetanus?

Tetanus is caused by bacteria that live in many different substances including soil, house dust, and human and animal waste such as manure. The bacteria usually enter the body through a wound in the skin and signs of tetanus usually start showing between 4 and 21 days after infection.

What are the symptoms of tetanus?

The first sign of tetanus is usually stiffness in the jaw muscles, which makes it difficult to open the mouth. This sometimes spreads from the jaw to the neck causing difficulties in breathing or swallowing.

Other symptoms of tetanus include high fever, sweating, high blood pressure, and rapid heartbeat. The highly infectious bacteria is spread through contaminated wounds. Tetanus in newborns, which is mostly fatal, is particularly common in rural areas where deliveries take place at home often without safe birthing practices.

What are the consequences of tetanus disease?

Deaths from tetanus are usually caused by serious complications associated with the disease. The most serious complications include lasting brain damage in infants, broken bones, seizures, suffocation leading to cardiac arrest, severe respiratory infections such as pneumonia, or death.

Preventing tetanus with DPT vaccine

Did you know a vaccine that prevents your child from getting tetanus is widely available?

The vaccine is called “DPT” and is given by health workers as an injection over three doses when the child is typically in the first few months of life. It is very safe and is provided by the government free of cost.

Consequences of not of getting vaccinated

A child that has not been vaccinated against tetanus it is much more likely to get the disease and much more likely to die from the disease. In fact, because your child has not been vaccinated, he/she is much more likely to get tetanus and much more likely to die from the disease. Many children in India die from tetanus every year because they or their mothers do not receive the DPT/TT vaccine.

The consequences of not vaccinating your child go beyond poor health. Because your child is not completely vaccinated, he/she is at increased risk for more frequent and severe infections which can cause delays in physical and mental development. This has serious implications on your child's productivity later on in life, including fewer years of education and lower earnings as an adult. Not vaccinating your child may prevent them from having a healthy and productive life.

Source of vaccination

To get your child vaccinated free of cost you should immediately contact the ASHA, Anganwadi worker or ANM of your village. Free vaccinations for newborns are available in your village on Village Health and Nutrition Day (VHND) which is held every month. You can also visit your nearest sub health centre (SHC), primary health care centre (PHC) or community health centre (CHC) to get your child vaccinated.

Please contact your village Health workers i.e. ASHA, Anganwadi worker or ANM for further information.

Information leaflet 2: Information framed as a gain.

What is tetanus?

Tetanus is a serious disease that easily spreads from person to person. People of all ages can get tetanus but newborn babies are particularly vulnerable to tetanus – it is one of the leading causes of death among newborns in India causing almost 32,000 newborn deaths per year.

How does a person get tetanus?

Tetanus is caused by bacteria that live in many different substances including soil, house dust, and human and animal waste such as manure. The bacteria usually enter the body through a wound in the skin and signs of tetanus usually start showing between 4 and 21 days after infection.

What are the symptoms of tetanus?

The first sign of tetanus is usually stiffness in the jaw muscles, which makes it difficult to open the mouth. This sometimes spreads from the jaw to the neck causing difficulties in breathing or swallowing.

Other symptoms of tetanus include high fever, sweating, high blood pressure, and rapid heartbeat. The highly infectious bacteria is spread through contaminated wounds. Tetanus in newborns, which is mostly fatal, is particularly common in rural areas where deliveries take place at home often without safe birthing practices.

Preventing tetanus with DPT vaccine

Did you know a vaccine that prevents your child from getting tetanus is widely available?

The vaccine is called "DPT" and is given by health workers as an injection over three doses when the child is typically in the first few months of life. It is very safe and is provided by the government free of cost.

Since your child has not yet received the full three doses of DPT, talk to your ASHA, Anganwadi worker, or ANM to get information about how your child can receive catch-up vaccines to protect him/her from tetanus.

Advantages of getting vaccinated

Once a child has been vaccinated against tetanus it is very unlikely he or she will ever get the disease and it is very unlikely he or she will die from the disease. In fact, if you vaccinate your child, they are much less likely to get tetanus and much less likely to die from the disease. Let me explain. If every child in India received the tetanus vaccine, many more children would survive and lead a healthy life.

The benefits of vaccinating your child go beyond preventing disease. Children who are completely vaccinated have fewer childhood infections and are therefore able to achieve good physical and mental health. Additionally, children who are vaccinated remain in school longer which also translates to better earnings as an adult. Vaccinating your child is one step towards enabling your child to have a healthy and productive life.

Source of vaccination

You can get your child vaccinated free of cost on the Village Health and Nutrition Day (VHND) which is held every month at your nearest AWC.

You can also visit your nearest Sub Health Centre (SHC), Primary Health Centre (PHC) or Community Health Centre (CHC) to get your child vaccinated. To get your child vaccinated free of cost or to get further information about vaccines you should immediately contact the ASHA, Anganwadi worker or ANM of your village.

Appendix B

References

- Athey, S., Imbens, G., 2016. Recursive partitioning for heterogeneous causal effects. *Proc. Natl. Acad. Sci.* 113 (27), 7353–7360. <https://doi.org/10.1073/pnas.1510489113>.
- Athey, S., Tibshirani, J., Wager, S., 2019. Generalized random forests. *Ann. Stat.* (2), 47. <https://doi.org/10.1214/18-AOS1709>.
- Athey, S., Wager, S., 2019. Estimating treatment effects with causal forests: an application. *Obs. Stud.* 5 (2), 37–51. <https://doi.org/10.1353/obs.2019.0001>.
- Banerjee, A., Chandrasekhar, A., Dalpath, S., Duflo, E., Floretta, J., Jackson, M., Kannan, H., Loza, F., Sankar, A., Schrimpf, A., & Shrestha, M. (2021). *Selecting the most effective nudge: evidence from a large-scale experiment on immunization*. <https://doi.org/10.3386/w28726>.
- Banerjee, A., Chandrasekhar, A., Duflo, E., Dalpath, S., Floretta, J., Jackson, M., Kannan, H., Schrimpf, A., Shrestha, M., 2020. Evaluating the Impact of Interventions to Improve Full Immunisation Rates in Haryana, India. <https://doi.org/10.23846/TW10IE126>.
- Bärnighausen, T., Bloom, D.E., Cafiero-Fonseca, E.T., O'Brien, J.C., 2014. Valuing vaccination. *Proc. Natl. Acad. Sci.* 111 (34), 12313–12319. <https://doi.org/10.1073/pnas.1400475111>.
- Basu, S., Kumbier, K., Brown, J.B., Yu, B., 2018. Iterative random forests to discover predictive and stable high-order interactions. *Proc. Natl. Acad. Sci.* 115 (8), 1943–1948. <https://doi.org/10.1073/pnas.1711236115>.
- Bhutta, Z.A., Das, J.K., Walker, N., Rizvi, A., Campbell, H., Rudan, I., Black, R.E., 2013. Interventions to address deaths from childhood pneumonia and diarrhoea equitably: what works and at what cost? *Lancet* 381 (9875), 1417–1429. [https://doi.org/10.1016/S0140-6736\(13\)60648-0](https://doi.org/10.1016/S0140-6736(13)60648-0).
- Brankovic, M., Kardys, I., Steyerberg, E.W., Lemeshow, S., Markovic, M., Rizopoulos, D., Boersma, E., 2019. Understanding of interaction (subgroup) analysis in clinical trials. In: *European Journal of Clinical Investigation*, 49. Blackwell Publishing Ltd. <https://doi.org/10.1111/eci.13145>.

- Brookes, S.T., Whitley, E., Peters, T.J., Mulheran, P.A., Egger, M., Smith, G.D., 2001. Subgroup analyses in randomised controlled trials: quantifying the risks of false-positives and false-negatives methodology. In: HTA Health Technology Assessment NHS R&D HTA Programme Health Technology Assessment, 5. www.hta.ac.uk/htaacd.htm.
- Bühlmann, P., 2002. Analyzing bagging. *Ann. Stat.* 30 (4), 927–961.
- Burke, J.F., Sussman, J.B., Kent, D.M., Hayward, R.A., 2015. Three simple rules to ensure reasonably credible subgroup analyses. *BMJ* 351. <https://doi.org/10.1136/bmj.h5651>.
- Celhay, P.A., Gertler, P.J., Giovagnoli, P., Vermeersch, C., 2019. Long-run effects of temporary incentives on medical care productivity. *Am. Econ. J.: Appl. Econ.* 11 (3), 92–127. <https://doi.org/10.1257/app.20170128>.
- Charness, G., Gneezy, U., 2009. Incentives to exercise. *Econometrica* 77 (3), 909–931. <https://EconPapers.repec.org/RePEc:ecm:emetrp:v:77:y:2009:i:3:p:909-931>.
- Chernozhukov, V., Demirer, M., Duflo, E., & Fernández-Val, I. (2018). *Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India*. <https://doi.org/10.3386/w24678>.
- Dhar, D., Jain, T., Jayachandran, S., 2022. Reshaping adolescents' gender attitudes: evidence from a school-based experiment in India. *Am. Econ. Rev.* 112 (3), 899–927. <https://doi.org/10.1257/aer.20201112>.
- Gibson, D.G., Ochieng, B., Kagucia, E.W., Were, J., Hayford, K., Moulton, L.H., Levine, O.S., Odhiambo, F., O'Brien, K.L., Feikin, D.R., 2017. Mobile phone-delivered reminders and incentives to improve childhood immunisation coverage and timeliness in Kenya (M-SIMU): a cluster randomised controlled trial. *Lancet Glob. Health* 5 (4), e428–e438. [https://doi.org/10.1016/S2214-109X\(17\)30072-4](https://doi.org/10.1016/S2214-109X(17)30072-4).
- Holland, P.W., 1986. Statistics and causal inference. *J. Am. Stat. Assoc.* 81 (396), 945. <https://doi.org/10.2307/2289064>.
- International Institute for Population Sciences (IIPS) and ICF, 2017. *National Family Health Survey (NFHS-4)*, pp. 2015–2016.
- International Institute for Population Sciences (IIPS) and ICF, 2021. *National Family Health Survey (NFHS-5)*, India, pp. 2019–2021.
- Johri, M., Pérez, M.C., Arsenault, C., Sharma, J.K., Pai, N.P., Pahwa, S., Sylvestre, M.-P., 2015. Strategies to increase the demand for childhood vaccination in low- and middle-income countries: a systematic review and meta-analysis. *Bull. World Health Organ.* 93 (5), 339–346C. <https://doi.org/10.2471/BLT.14.146951>.
- Kasi, S.G., Shivananda, S., Marathe, S., Chatterjee, K., Agarwala, S., Dhir, S.K., Verma, S., Shah, A.K., Srirampur, S., Kalyani, S., Pemde, H.K., Balasubramanian, S., Parekh, B.J., Basavaraja, G.V., Gupta, P., 2021. Indian Academy of Pediatrics (IAP) Advisory Committee on Vaccines and Immunization Practices (ACVIP): recommended immunization schedule (2020–21) and update on immunization for children aged 0 through 18 years. *Indian Pediatr.* 58 (1), 44–53. <https://doi.org/10.1007/s13312-021-2096-7>.
- Kent, D.M., Rothwell, P.M., Ioannidis, J.P.A., Altman, D.G., Hayward, R.A., 2010. Assessing and reporting heterogeneity in treatment effects in clinical trials: a proposal. *Trials* 11 (1), 85. <https://doi.org/10.1186/1745-6215-11-85>.
- Kwasnicka, D., Dombrowski, S.U., White, M., Sniehotta, F., 2016. Theoretical explanations for maintenance of behaviour change: a systematic review of behaviour theories. *Health Psychol. Rev.* 10 (3), 277–296. <https://doi.org/10.1080/17437199.2016.1151372>.
- Larson, H.J., Jarrett, C., Eckersberger, E., Smith, D.M.D., Paterson, P., 2014. Understanding vaccine hesitancy around vaccines and vaccination from a global perspective: a systematic review of published literature, 2007–2012. *Vaccine* 32 (19), 2150–2159. <https://doi.org/10.1016/j.vaccine.2014.01.081>.
- Mekonnen, Z.A., Gelaye, K.A., Were, M.C., Gashu, K.D., Tilahun, B.C., 2019. Effect of mobile text message reminders on routine childhood vaccination: a systematic review and meta-analysis. *Syst. Rev.* 8 (1), 154. <https://doi.org/10.1186/s13643-019-1054-0>.
- Mills, E., Jadad, A.R., Ross, C., Wilson, K., 2005. Systematic review of qualitative studies exploring parental beliefs and attitudes toward childhood vaccination identifies common barriers to vaccination. *J. Clin. Epidemiol.* 58 (11), 1081–1088. <https://doi.org/10.1016/j.jclinepi.2005.09.002>.
- Oyo-Ita, A., Wiysonge, C.S., Oringanje, C., Nwachukwu, C.E., Oduwole, O., Meremikwu, M.M., 2016. Interventions for improving coverage of childhood immunisation in low- and middle-income countries. *Cochrane Database Syst. Rev.* (7), 2016. <https://doi.org/10.1002/14651858.CD008145.pub3>.
- Padula, W.V., Kreif, N., Vannest, D.J., Adamson, B., Rueda, J.-D., Felizzi, F., Jonsson, P., IJerman, M.J., Butte, A., Crown, W., 2022. Machine learning methods in health economics and outcomes research—the PALISADE checklist: a good practices report of an ISPOR Task force. *Value Health* 25 (7), 1063–1080. <https://doi.org/10.1016/j.jval.2022.03.022>.
- Powell-Jackson, T., Fabbri, C., Dutt, V., Tougher, S., Singh, K., 2018. Effect and cost-effectiveness of educating mothers about childhood DPT vaccination on immunisation uptake, knowledge, and perceptions in Uttar Pradesh, India: a randomised controlled trial. *PLoS Med.* 15 (3), e1002519. <https://doi.org/10.1371/journal.pmed.1002519>.
- Powell-Jackson, T., Mazumdar, S., Mills, A., 2015. Financial incentives in health: new evidence from India's Janani Suraksha Yojana. *J. Health Econ.* 43, 154–169. <https://doi.org/10.1016/j.jhealeco.2015.07.001>.
- Robins, J.M., Rotnitzky, A., Zhao, L.P., 1994. Estimation of regression coefficients when some regressors are not always observed. *J. Am. Stat. Assoc.* 89 (427), 846. <https://doi.org/10.2307/2290910>.
- Rosner, Bernard., 2011. *Fundamentals of Biostatistics, 7th Edition*. Brooks/Cole.
- Rubin, D.B., 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* 66 (5), 688–701. <https://doi.org/10.1037/h0037350>.
- Tauchmann, H., 2014. Lee (2009) treatment-effect bounds for nonrandom sample selection. *Stata J.: Promot. Commun. Stat.* 14 (4), 884–894. <https://doi.org/10.1177/1536867X1401400411>.
- Tibshirani, J., Athey, S., Friedberg, R., Hadad, V., Hirshberg, D., Miner, L., Sverdrup, E., Wager, S., Wright, M., 2020. Package grf: generalized random forests. Version., Available at the Comprehensive R Archive Network.
- Tougher, S., Dutt, V., Pereira, S., Haldar, K., Shukla, V., Singh, K., Kumar, P., Goodman, C., Powell-Jackson, T., 2018. Effect of a multifaceted social franchising model on quality and coverage of maternal, newborn, and reproductive health-care services in Uttar Pradesh, India: a quasi-experimental study. *Lancet Glob. Health* 6 (2), e211–e221. [https://doi.org/10.1016/S2214-109X\(17\)30454-0](https://doi.org/10.1016/S2214-109X(17)30454-0).
- Volpp, K.G., Troxel, A.B., Pauly, M.V., Glick, H.A., Puig, A., Asch, D.A., Galvin, R., Zhu, J., Wan, F., DeGuzman, J., Corbett, E., Weiner, J., Audrain-McGovern, J., 2009. A randomized, controlled trial of financial incentives for smoking cessation. *N. Engl. J. Med.* 360 (7), 699–709. <https://doi.org/10.1056/NEJMsa0806819>.
- Wager, S., Athey, S., 2018. Estimation and inference of heterogeneous treatment effects using random forests. *J. Am. Stat. Assoc.* 113 (523), 1228–1242. <https://doi.org/10.1080/01621459.2017.1319839>.
- Wang, R., Lagakos, S.W., Ware, J.H., Hunter, D.J., Drazen, J.M., 2007. Statistics in medicine — reporting of subgroup analyses in clinical trials. *N. Engl. J. Med.* 357 (21), 2189–2194. <https://doi.org/10.1056/NEJMsr077003>.
- Wang, X., Piantadosi, S., Le-Rademacher, J., Mandrekar, S.J., 2021. Statistical considerations for subgroup analyses. *J. Thorac. Oncol.* 16 (3), 375–380. <https://doi.org/10.1016/j.jtho.2020.12.008>.
- World Health Organisation, 2022. Immunization Coverage. July 14. <https://www.who.int/news-room/fact-sheets/detail/immunization-coverage>.