

SELECTING THE MOST EFFECTIVE NUDGE: EVIDENCE FROM A LARGE-SCALE EXPERIMENT ON IMMUNIZATION

ABHIJIT BANERJEE^{†,*;§}, ARUN G. CHANDRASEKHAR^{‡,*;§}, SURESH DALPATH[#],
ESTHER DUFLO^{†,*;§}, JOHN FLORETTA^{*}, MATTHEW O. JACKSON^{‡,*}, HARINI KANNAN^{*},
FRANCINE LOZA^{*}, ANIRUDH SANKAR[‡], ANNA SCHRIMPF^{*}, AND MAHESHWOR SHRESTHA[§]

ABSTRACT. Policymakers often choose a policy bundle that is a combination of different interventions in different dosages. We develop a new technique—*treatment variant aggregation* (TVA)—to select a policy from a large factorial design. TVA pools together policy variants that are not meaningfully different and prunes those deemed ineffective. This allows us to restrict attention to aggregated policy variants, consistently estimate their effects on the outcome, and estimate the best policy effect adjusting for the winner’s curse. We apply TVA to a large randomized controlled trial that tests interventions to stimulate demand for immunization in Haryana, India. The policies under consideration include reminders, incentives, and local ambassadors for community mobilization. Cross-randomizing these interventions, with different dosages or types of each intervention, yields 75 combinations. The policy with the largest impact (which combines incentives, ambassadors who are information hubs, and reminders) increases the number of immunizations by 44% relative to the status quo. The most cost-effective policy (information hubs, ambassadors, and SMS reminders but no incentives) increases the number of immunizations per dollar by 9.1% relative to status quo.

Date: June 18, 2024.

We are particularly grateful to the the Haryana Department of Health and Family Welfare for taking the lead on this intervention and allowing the evaluation to take place. Among many other, we acknowledge the tireless support of Rajeev Arora, Amneet P. Kumar, Sonia Trikha, V.K. Bansal, Sube Singh, and Harish Bisht. We are also grateful to Isaiah Andrews, Adel Javanmard, Adam McCloskey, Tyler McCormick, and Karl Rohe for helpful discussions. We thank Emily Breza, Denis Chetvetrikov, Paul Goldsmith-Pinkham, Ben Golub, Nargiz Kalantarova, Shane Lubold, Douglas Miller, Francesca Molinari, Suresh Naidu, Eric Verhoogen, and participants at various seminars for suggestions. Financial support from USAID DIV, 3iE, J-PAL GPI, and MIT-Givewell, and NSF grant SES-2018554 is gratefully acknowledged. Chandrasekhar is grateful to the Alfred P. Sloan foundation for support. We thank Rachna Nag Chowdhuri, Diksha Radhakrishnan, Mohar Dey, Maaïke Bijker, Niki Shrestha, Tanmayta Bansal, Sanjana Malhotra, Aicha Ben Dhia, Vidhi Sharma, Mridul Joshi, Rajdev Brar, Paul-Armand Veillon, and Chitra Balasubramanian, as well as Shreya Chaturvedi, Vasu Chaudhary, Shobitha Cherian, Louis-Mael Jean, Anoop Singh Rawat, Elsa Trezeguet, Simon Yaspo and Meghna Yadav for excellent research assistance.

[†]Department of Economics, MIT.

[‡]Department of Economics, Stanford University.

[#]Public Health Planning, Policy, and M&E, Health Department, Govt. of Haryana.

^{*}J-PAL.

[§]NBER.

^{*}Santa Fe Institute.

1. INTRODUCTION

In many settings, policymakers have to select the best policy among potential bundles that combine several interventions, each with different possible dosages or varieties. Similarly in medicine, a particular treatment regimen may combine several drugs in different potential dosages. For example, the management of HIV-AIDS was revolutionized in the mid 1990 by the combination of two or three drugs in subtle dosages, the famous “AIDS Cocktail.”

In this paper, we consider the problem of the state government in Haryana, India, who was looking to choose a new bundle of interventions to increase children’s immunization coverage. Immunization is recognized as one of the most effective and cost-effective ways to prevent illness, disability, and death. Yet nearly 20 million children under one do not receive critical immunizations each year (UNICEF and WHO, 2019). In 2016, 7 million of these children were in India, where only 60% of children received the basic set of vaccination within one year of life. Though resources directed towards immunization have increased steadily, there is mounting evidence that insufficient parental demand has contributed to stagnating immunization rates (WHO, 2019). Based on the existing research, the options considered were small incentives to parents, social network mobilization, and SMS reminders. The government also needed to determine the level and slope of incentives, the set of people to mobilize, and the intensity of the SMS reminder campaigns.

The ideal strategy, if time and implementation capacities were not constraints, would be to experiment iteratively in the context until the best bundle is found. There is a growing literature on how to conduct and analyze adaptive trials (Hadad et al., 2021; Kasy and Sautmann, 2021; Zhan et al., 2021). However, it is often not possible to conduct such sequential trials: the window for experimentation may be short before a policy must be chosen, or a decision on a treatment regimen or vaccine must be made quickly because of an urgent health threat. In such cases, the only option may be to conduct a large scale experiment that simultaneously tests many different policy bundles. This was certainly true in the case of HIV-AIDs, where there were huge pressures to rapidly identify and approve a treatment. In Haryana, it was possible to conduct a single experiment with multiple treatments, in over 900 villages.

However, with several alternative interventions and multiple possible dosages, there is an enormous number of potential combinations, each of which is a unique policy bundle. Altogether, in Haryana, there were 75 possible bundles of interventions with different dosages. There is no clear guidance in the literature on how to design and analyze such trials when the number of potential options is large.

We fill this gap by developing a methodology for *treatment variant aggregation* (TVA): a principled algorithm that pools together policy variants that have similar impact and prunes ineffective policy options. This reduces the dimensionality of the problem, and enables

identification of the best overall combination and consistent estimation of its impact. This methodology allows us to solve two problems that arise when evaluating large numbers of candidate policy bundles (henceforth we use “policy” as a short-hand for “policy bundle”).

First, the researcher must decide how many and which potential policies to include in an experimental design, and how to analyze the trial. One approach, recommended in McKenzie (2019), is to only include a limited number of bundles. For example, the ACTG-320 trials compared the three drug cocktail (protease inhibitor plus two nucleosides), in one specific dosage, to the two nucleosides in the same dosage (Hammer et al., 1997). However, the optimality of such an approach presumes an “oracle property:” that the researcher or policymaker *already knows* which policies are worth comparing. We consider situations in which any of many policies could turn out to be optimal, and the researcher would like to choose among all unique policy bundles in a fully-saturated factorial design. This reduces power since each policy may only be observed on a small sample, therefore, to increase precision, researchers often attempt to pool policy bundles ex-post based on observed outcomes. Without specific structure to the problem, however, this can be misleading in finite samples, especially when interaction effects are small, so that a test against zero has limited power, but are not quite zero (this is the “local-to-zero” problem, see Muralidharan et al. (2019)). Thus, we need to find minimal and realistic assumptions on the inferential environment that enable a principled, data driven approach to reducing the dimensionality of the problem.

The second problem is that the impact of a policy that is estimated to be the best out of a set of K unique policies can be over-estimated when K is large, due to the “winner’s curse” (Andrews, Kitagawa, and McCloskey, 2021). Some policy k^* could have the highest estimated impact partially due to a high true effect, but could also partially be due to randomness. Conditional on being the best in the data, some of the estimated impact is likely due to randomness. As a result, the expected impact of policy k^* is overestimated and the statistician must adjust for it.

Our main methodological contribution takes place before any estimation: it is a method for mapping a large number of unique policies into a small number of bundles in a data driven way. We develop a tool (a specific Hasse diagram) to represent a complicated factorial design while incorporating what the researcher knows about the structure of the problem, and enable treatment variant aggregation. We argue that this satisfies conditions required to apply appropriate existing regularization estimators (the Puffer LASSO) to collapse dosages into a smaller number of intervention bundles. This innovation thus allows us to combine existing methods to develop estimators that are consistent and asymptotically normally distributed for a large factorial design. In the case of the interventions to maximize the

number of measles immunizations in Haryana, this step reduces it from 75 candidate policies to 4.

The statistical setting that we analyze is as follows. There are M possible intervention arms, R possible “dosages” per arm (including zero) and therefore $K = R^M$ possible policies. The policymaker is uncertain about which policies are effective. However, it can be that in some circumstances an *incremental* dosage change on some arm does not have a meaningful effect on the outcome, for some other combination of other arms. For example, if there are three interventions to increase immunization demand (x, y, z) , and two adjacent dosages z and z' for the third intervention (e.g., SMS are sent to 33% of the parents, or to 66% of the parents), then it is possible that the bundles (x, y, z) and (x, y, z') are equally effective for some particular choices of x and y . We can thus pool those two policies, (x, y, z) and (x, y, z') , and treat them as one for all practical purposes. The policymaker conducts *treatment variant aggregation* (TVA). This *pools* together policy variants that are not meaningfully different (e.g., (x, y, z) and (x, y, z') are pooled together as above) and *prunes* all the combinations that are ineffective (those that pool with the null policy). TVA allows us to restrict attention to aggregated policy variants and only those that matter, which can improve estimation. We discuss how we can use TVA to consistently estimate policy effects and estimate the best policy effect adjusting for the winner’s curse. We proceed in several steps.

The first step is to represent the fully-saturated factorial regression of the outcome on unique policies in terms of another, equivalent specification that tracks the effects of incremental dosages. TVA utilizes a Hasse diagram lattice of policy variants to deduce how zeros in the marginal effects determines pruning and pooling of variants.

To fix ideas, consider a simple example with two arms ($M = 2$) and two non-zero dosages for each ($R = 3$), yielding $K = 9$ unique policies. So each arm can either be used or not, and used in either a low or high dose. Let us represent these by $[T_1, T_2] = [0, 0], [0, 1], [0, 2], [1, 0], \text{etc.}$, where the entries are the corresponding treatment levels with 0 being not used, and 1 being low and 2 high dosage. A standard regression would just have a dummy variable for each particular policy combination $[T_1, T_2]$, and then a corresponding coefficient $\beta_{[T_1, T_2]}$. An alternative representation breaks these into marginal effects:

$$\begin{aligned}
 (1.1) \quad y = & \alpha_{[0,0]} + \alpha_{[1,0]} \cdot 1\{T_1 \geq 1\} \cdot 1\{T_2 = 0\} + \alpha_{[0,1]} \cdot 1\{T_1 = 0\} \cdot 1\{T_2 \geq 1\} \\
 & + \alpha_{[2,0]} \cdot 1\{T_1 = 2\} \cdot 1\{T_2 = 0\} + \alpha_{[0,2]} \cdot 1\{T_1 = 0\} \cdot 1\{T_2 = 2\} \\
 & + \alpha_{[1,1]} \cdot 1\{T_1 \geq 1\} \cdot 1\{T_2 \geq 1\} + \alpha_{[2,1]} \cdot 1\{T_1 = 2\} \cdot 1\{T_2 \geq 1\} \\
 & + \alpha_{[1,2]} \cdot 1\{T_1 \geq 1\} \cdot 1\{T_2 = 2\} + \alpha_{[2,2]} \cdot 1\{T_1 = 2\} \cdot 1\{T_2 = 2\} + \epsilon.
 \end{aligned}$$

In this specification, the $\alpha_{[r_1, r_2]}$ are all marginal effects, and hence, inspecting the vector α and checking which $\alpha_{[r_1, r_2]} = 0$ tells us which adjacent policies can be pooled together, and which ones can be pruned (pooled with the null policy; for instance, if $\alpha_{[1,0]} = 0$). In a general factorial design of K unique policies, we have regressors of the form $1\{T_1 \geq r_1, T_2 \geq$

$r_2, \dots, T_M \geq r_M\}$ for treatment arm intensities T_m and thresholds r_m for arm m , with $K - 1$ regressors plus an intercept. At every stage we ask whether an incremental increase in dosage for a given arm of some policy causes a marginal change. That is, we check for zero effects: $\alpha_{[r_1, \dots, r_M]} = 0$ for some or multiple $[r_1, \dots, r_M]$.

This approach makes use of the researcher’s a priori knowledge of which policies can be pooled: these are policies that are dosage variations of the same treatment *profile*, or underlying policy type. Therefore, it places discipline on the problem. It ensures we are not mis-naming pooled choices by pooling non-comparable policy bundles, which is the issue implicitly raised in Muralidharan et al. (2019).¹ We assume that when there are non-zero marginal effects, those effect sizes are large enough—assuming away the local-to-zero range—so that we may actually discover and make inferences about the best policy combinations. Our approach works when these assumptions to allow regularization are palatable. When the assumptions, or reasonable relaxations, cannot be justified, sequential testing, rather than simultaneous testing with post-selection processing, is unavoidable.²

Our goal is to identify the support (the set of non-zero coefficients) of the regression equation (1.1). Under our maintained assumptions, a natural way to do this is to use LASSO. This requires an extra step, however, since the regressors in equation (1.1) are typically strongly correlated. For instance, $1\{T_1 = 2\} \cdot 1\{T_2 = 0\}$ implies $1\{T_1 \geq 1\} \cdot 1\{T_2 = 0\}$. In fact, the marginal effects specification may fail the necessary and sufficient condition for LASSO support consistency, of “irrepresentability,” which requires that regressors are not too correlated (Zhao and Yu, 2006). Thus, the second step is to apply the *Puffer transformation* to the variables to which LASSO is being applied (Rohe, 2014; Jia and Rohe, 2015). This decorrelates the design matrix that comes from (1.1). We show that the the specific structure of the RCT makes it particularly suitable for this technique.

Once LASSO has been applied on the Puffer-transformed variables to consistently estimate the marginal effects support, the third step is to reconstruct a set of unique policies taking into account the pooling and pruning implied by the LASSO results.

The fourth step is to estimate OLS on the new set of unique policies, post-selection. Using an argument adapted from Javanmard and Montanari (2013), we show that this estimator is consistent and asymptotically normally distributed.³

¹Specifically, Muralidharan et al. (2019) take issue with “short models” such that, for example, what is claimed as the effect of $(x, 0, 0)$ actually also includes some of the effect of $(x, y, 0)$. In this sense the treatment is “mis-named”. In TVA, the policy $(x, y, 0)$ is considered to be a categorically different treatment type from $(x, 0, 0)$ for $x, y > 0$. More generally, the pooled policy names always unambiguously indicate which unique policy combinations are pooled together.

²In practice, we show in Online Appendix E.3 through simulations that we may relax the local-to-zero assumption in several directions and still retain strong performance for this final objective.

³The convergence in distribution is not uniform (in the parameter space) Leeb and Pötscher (2005). Nevertheless, asymptotic normality holds pointwise (in the parameter space)—essentially, in our setting, the

This vector is of independent interest: these are the estimates and confidence interval of all the relevant policies. There is an optional fifth step, in case the policy maker is interested in identifying and getting estimate for a single “best” policy. This step is to re-estimate the effect of the best pooled and pruned policy, adjusting for the winner’s curse (Andrews et al., 2021). There are three advantages of conducting this adjustment post-TVA rather than on the full factorial design. First, when there are fewer potential alternatives to the best policy k^* , the odds of picking the best by chance are lower. In fact, in many cases, the winner’s curse adjustment may not be necessary, when there are sufficiently few policy bundles that survive at the TVA step. Second, with fewer alternatives, it is less likely that the second-best alternative has an effect that is similar to the k^* effect, which will reduce the shrinkage penalty. Third, there is the benefit of coherence: if two policies can be pooled, there is no point in applying a strong shrinkage penalty because of a competition between them.

We apply this method to the large-scale immunization experiment that we conducted in Haryana, India, from December 2016 to November 2017 in collaboration with the government of Haryana, which was interested in selecting the best policy for full scale adoption in the state . To stimulate demande for immunization, a large literature has found the effectiveness of “nudges,” including conventional ones such as small cash or in-kind incentives,⁴ SMS reminders,⁵ as well as more novel interventions such as symbolic social rewards⁶ or using influential individuals in a social network as “ambassadors.”⁷ We cross-randomized three arms with different nudges that had shown some promise in earlier work: (1) monetary incentives, (2) SMS reminders, and (3) seeding ambassadors. Incentives came in two types (linear and convex) with two dosages each (low and high in terms of value). SMS reminders had two dosages. Either 33% or 66% of caregivers received SMS reminders (and voicemails) about the next scheduled vaccination. Ambassadors where either randomly selected or selected through a nomination process. The nomination process was done in three ways, one of which came in two dosages (Information Hub). All together, we have 75 unique policies and 915 villages were at risk for all three treatments taken together.

Applying TVA, we find that when the outcome is the number of measles shot administered, four policies survive as candidate policies, including two with coefficient that are significantly different from zero, both of which involve the combination of ambassadors nominated by the social network, SMS, and incentives. The best policy is to use information hubs and either

non-uniformity does not have much bite since incorrect selection of the high-effect policies happens with probability tending rapidly to zero.

⁴See Banerjee et al. (2010); Bassani et al. (2013); Wakadha et al. (2013); Johri et al. (2015); Oyo-Ita et al. (2016); Gibson et al. (2017).

⁵See Wakadha et al. (2013); Domek et al. (2016); Uddin et al. (2016); Regan et al. (2017).

⁶See Karing (2018).

⁷See Alatas et al. (2019); Banerjee et al. (2019).

low or high SMS coverage, in combination with convex incentives that can be either low or high. This increases the number of immunizations by 44% ($p < 0.05$), after accounting for the winner’s curse. Choosing the cheapest among these suggests that the policymaker should chose low convex incentives, send SMS to 33% of caregivers, and identify information hubs to relay the message. To maximize the number of immunization per dollar spent, the best policy is using information hubs along with SMS reminders at 33% or more of caregivers covered. It increases the number of immunizations per dollar by 9.1% ($p < 0.05$) compared to the status quo with no additional intervention. It is the only policy that strictly increases the number of immunizations per dollar spent, and hence in this case the winner’s curse adjustment makes essentially no difference (and this step could be omitted).

The results highlight the importance of complementarities that may get lost had a factorial design not been used. Information hubs magnify the effect of other interventions and spark diffusion: neither incentives nor reminders are selected on their own, but are selected when combined with information hubs. Similarly, information hubs are not selected on their own, but are selected when combined with the conventional strategies. This suggests that in cases where there are no strong reasons to rule out interactions a priori, it is important to accommodate them in the design and the statistical analysis.

2. TREATMENT VARIANT AGGREGATION

2.1. Overview and Setup. We have a randomized controlled trial of M arms and R ordered dosages ($\{\text{none, intensity } 1, \dots, \text{intensity } R - 1\}$). This yields $K := R^M$ unique treatment combinations or *unique policies*. Let $T_{ik} \in \{0, 1\}$ be a dummy variable indicating that unit i is assigned to unique policy k . Unique policies are described as *variants* of each other when they differ only in the (non-zero) dosages of the treatments applied. This implies that two policies differing only in whether some arm is active or inactive (dosage is zero) are not considered variants, as formalized below in Section 2.1.1.

Assuming the same number of dosages per arm is for notational ease and without substantive loss of generality. In practice the number of dosages per arm can vary.

The *unique policy regression* is given by

$$(2.1) \quad y_i = T\beta^0 + \epsilon_i.$$

The support of this regression is given by the set of unique policies that have non-zero effect relative to control,

$$S_\beta := \{k \in [K] : |\beta_k^0| \neq 0\}.$$

Some of the variants have equivalent effects and ought to be considered as one policy. Some arms may be altogether ineffective and ought to be pruned (i.e., pooled with control).

We construct a method of *treatment variant aggregation* (TVA) in order to *pool and prune variants* systematically.

2.1.1. *Treatment Profiles and Policy Variants.* A fundamental concept is a *treatment profile*. This clarifies which unique policies are *variants* of each other and could potentially be pooled with one another (without being pooled with the control).

The treatment profile $P(k)$ of a unique policy k designates which of the M arms are active (having positive dosages), without regards to how high the dosage is. Two unique policies k, k' are variants of each other if and only if $P(k) = P(k')$, i.e., exactly the same arms are active for both policies. Thus, K unique policies are categorized into 2^M treatment profiles.

EXAMPLE 1. Consider observation i that has an assigned policy $k = (\text{No Ambassador, } 33\% \text{ SMS, low-value flat incentives})$ and observation j that has an assigned assigned policy $k' = (\text{No Ambassador, } 66\% \text{ SMS, low-value flat incentives})$. Though k and k' are distinct treatment combinations, they share the same treatment profile— $P(k) = P(k')$ —of $(\text{No Ambassador, Some SMS, Some incentives})$. Therefore k and k' are variants. They would not be variants if instead $k' = (\text{No Ambassador, } 66\% \text{ SMS, No incentives})$.

2.1.2. *Treatment Variant Aggregation: Pooling and Pruning.* Increasing the dosage in a treatment arm may be inconsequential after a point, and more generally policy variants may have the same impact. Here, we consider a re-specification of (2.1) that explicitly tracks the marginal effect of increasing dosages by grouping together policy variants that have the same effect on the outcome.⁸ When these marginal effects are zero, this means that a set of variants are to be either pooled or pruned (pooled with control).

Let \mathcal{P} denote the set of all partitions of the K policies. Elements of \mathcal{P} comprise every conceivable pooling of the K policies, with generic partition denoted Π . Whether two given policies k and k' are pooled corresponds to whether they are members of the same element of the partition, $\pi \in \Pi$.

Out of the full universe of all conceivable poolings only some make sense, and we refer to those as the *admissible* poolings. Informally, admissibility says that only policies affected by the same set of nonzero marginal effects *may* be pooled. The admissible pools are then a strict subset $\mathcal{P}_{|\Lambda} \subset \mathcal{P}$. The target $S_{TVA} \in \mathcal{P}_{|\Lambda}$ is defined to be the maximally admissible pooled and pruned set of policies (i.e., the coarsest partition).

⁸While sometimes what is “dosage” and “dosage ordering” is readily apparent from the arm, as in the SMS arm of our intervention with saturation levels 33% and 66%, in other cases the researcher has to decide this (of course this can be pre-specified). For example, in the seeds arm of our intervention, we decided that the information hub ambassador comes in two dosages, with those that are trusted for health advice as the higher dosage.

Letting Z_{STVA} denote the matrix of indicator variables for the pooled policies, our goal is to estimate the *pruned and pooled policy regression*:

$$(2.2) \quad y = Z_{STVA} \eta_{STVA}^0 + \epsilon.$$

Comparing (2.2) with (2.1), η^0 is the projection coefficients of $T\beta^0$ onto Z_{STVA} , that corresponds to grouping certain policies, and estimating the parameters for the grouped policies.

Let us make admissibility precise. For a treatment combination k , α_k^0 is the marginal effect of the dosages in k within its treatment profile relative to incrementally lower dosages. Formally, the marginal effect α^0 may be defined implicitly so that a policy's effect is the sum of marginal effects from increasing dosages up to its particular dosage profile:

$$(2.3) \quad \beta_k^0 = \sum_{k' \leq k; P(k')=P(k)} \alpha_{k'}^0.$$

Equation (2.3) can be inverted to recover α^0 in terms of β^0 . An explicit expression for its terms α_k^0 is more unwieldy in its full generality, but depending on the policy k , it can be a difference between two variants' effects or reflect a complementarity, i.e., the interaction effect from combining dosages in different arms. This is consistent with the interpretation that a policy's effect is the main effects of the highest dosages in each arm, considered separately, plus the relevant interaction effects.

For each policy k , consider the set of marginal effects on the right hand side of β_k^0 in (2.3) that are nonzero: these are the *active marginal effects for k* , $A(k)$. These are the set of marginal effects that “influence” the policy k , either as a main effect or as a complementarity.

ASSUMPTION 1. Π is an admissible pooling —i.e., $\Pi \in \mathcal{P}_\Lambda$ —if and only if $k, k' \in \pi$ implies $A(k) = A(k')$. That is, only policies influenced by the same set of active marginal effects may be pooled.

It is easy to see, through (2.3), that admissibility ensures that (i) only policies with equal treatment effects may be pooled, and (ii) only variants may be pooled (or if non-variant policies are being pooled, they must be pooled with the control – the null policy). Of course, given an idiosyncratic instance of unique policy effects β^0 , there may be other pools that satisfy (i) and (ii) besides the admissible pools. However, we restrict our attention to admissible pools since these work generally using only the sign of marginal effects.

Per Assumption 1, the more zeros there are in marginal effects, the more pooling choices become admissible. We can depict this in a Hasse diagram for a treatment profile. In a Hasse diagram, a line upwards from variants k to k' implies $k' > k$, and there is no variant k'' such that $k' > k'' > k$ (in the partial order). The running example is the case of a 2 arm treatment of 4 intensities (3 nonzero intensities “low,” “medium,” “high”); i.e., $M = 2, R = 4$ and the treatment profile where both arms are “on”. Figure 1 depicts the Hasse for this treatment

profile. Here unique policies are named per their intensity representations; i.e., $[r_1, r_2]$ where $r_i \in \{1, 2, 3\}$ is the (nonzero) dosage in arm i .

Zeros in the marginals make admissible certain “concatenations” in the Hasse diagram of policies. This is depicted in Figure 2, where the top panels (A-C) depict the zeros in marginal effects and the bottom panels (D-F) depict the maximal admissible policy concatenations they imply (of course subsets of these concatenations will also be admissible).

In Panel A, $\alpha_{[2,1]} = \alpha_{[3,1]} = 0$, meaning that keeping the intensity fixed as low in arm 2, there is no marginal contribution of increasing the intensity in arm 1. Panel D depicts how this makes admissible the concatenation of policies $\{[1, 1], [2, 1], [3, 1]\}$, and indeed $\beta_{[1,1]} = \beta_{[2,1]} = \beta_{[3,1]}$. The maximal admissible concatenated policy can be called $[1 : 3, 1]$.

In Panel B, $\alpha_{[2,2]} = 0$. $\alpha_{[2,2]} = (\beta_{[2,2]} - \beta_{[2,1]}) - (\beta_{[1,2]} - \beta_{[1,1]})$, there are only main effects in increasing dosages from low to medium intensities in both arms, and no further complementarity. Since furthermore $\alpha_{[1,2]} = 0$, there is no main effect on arm 2 from increasing low to medium intensity. That is, there is only a main effect in arm 1 from increasing low to medium intensities. This main effect is nonzero, since $\alpha_{[2,1]} \neq 0$. In Panel E, the maximal admissible concatenations reflect this: $\{[1, 1], [1, 2]\}$ into $\{1, 1 : 2\}$, and $\{[2, 1], [2, 2]\}$ into $[2, 1 : 2]$.

The main effect of arm 1 makes inadmissible these concatenated blocks from further concatenating. This changes in Panel C, where $\alpha_{[2,1]} = 0$ implies that there is no main effect in increasing arm 1 from low to medium either. This makes admissible the concatenation $\{1, 1 : 2\}$ and $[2, 1 : 2]$ into $[1 : 2, 1 : 2]$.

As illustrated through these examples, zeros in α^0 thus show up as admissible policy concatenations in Hasse diagrams. This motivates the *marginal effects regression*:

$$(2.4) \quad y = X\alpha^0 + \epsilon.$$

This is an invertible transformation of (2.1). X can be interpreted as indicators

$$X_{i\ell} := \mathbf{1}\{k(i) \geq \ell \cap P(k(i)) = P(\ell)\}.$$

In other words, X assigns for unit i a “1” for all policy variants that share $k(i)$ ’s treatment profile and are weakly dominated in intensity by $k(i)$ and a “0” otherwise.

The key object of interest is the support of (2.4):

$$S_\alpha := \{j \in [K] : |\alpha_j^0| \neq 0\}.$$

S_α is the set of all active marginals for any policy, i.e. $S_\alpha = \cup_k A(k)$. Since $S_{TVA} \in \mathcal{P}_\Lambda$ is the maximal admissible pooling, it is the tightest pooling using only S_α and nothing else. More precisely, it is the coarsest pooling uniformly consistent over α conditional on S_α .

The idea is to apply a model selection procedure to estimate S_α . In Online Appendix B, we show how to construct the unique maximally pooled and pruned set $S_{TVA} \in \mathcal{P}_\Lambda$ from S_α .⁹ The maximality ensures that no contiguous set of intensities thought to have the same treatment effects are left un-pooled.

2.2. Pooling and Pruning for Support Selection. The next step is to identify the support S_α . One natural place to start would be to apply LASSO directly to (2.4). However, this approach can fail to satisfy sign consistency because the marginal effects matrix X may fail an “irrepresentability criterion” which is necessary for consistent estimation (Zhao and Yu, 2006). Irrepresentability bounds the acceptable correlations in the design matrix. Intuitively, it requires that variables that are not in the support are not too strongly correlated with those that are. Otherwise, an irrelevant variable is “representable” by relevant variables, which makes LASSO erroneously select it with non-zero probability irrespective of sample size. We prove by construction that irrepresentability can fail to be satisfied in Online Appendix C, where we also show by simulation that irrepresentability failures can become dramatic with increasing R and M . The structure that we exploit in showing the failure is one in which higher dosage marginals are representable by lower dosage marginals, violating the condition.

A way out is provided by Jia and Rohe (2015). They show that, under some conditions, one can estimate the LASSO support by transforming the data to recover irrepresentability. They demonstrate that a simple left-multiplication (pre-conditioning) can de-correlate the data (at the expense of inflating variance in the error).

In Proposition 2.1 we demonstrate that in the specific instance of the crossed RCT design with ordered intensities, the pre-conditioning strategy of Jia and Rohe (2015) can be applied because the relevant sufficient conditions are met. Specifically, with an RCT, we can exactly characterize the design matrix and therefore the inflation factor. We can show that the variance inflation cost is tolerable, in the sense that we can consistently recover the support and the treatment effects.

The weighting is constructed as follows. Let us take the singular value decomposition of $X := UDV'$ where U is an $n \times K$ unitary matrix, D is a $K \times K$ diagonal matrix of singular values, and V is a $K \times K$ unitary matrix. The *Puffer transformation*—so named for the fish whose shape is suggested by the geometry of this transformation—is $F := UD^{-1}U'$. The regression of interest is now

$$(2.5) \quad Fy = FX\alpha + F\epsilon$$

⁹Following this same procedure with any estimate \hat{S}_α leads to an estimate \hat{S}_{TVA} of pooled and pruned policies.

where $F\epsilon \sim \mathcal{N}(0, UD^{-1}\Sigma D^{-1}U')$. As [Jia and Rohe \(2015\)](#) note, this satisfies irrepresentability since $(FX)'(FX) = I$, which is sufficient ([Jia and Rohe \(2015\)](#), [Bickel et al. \(2009\)](#)).

To understand why this works, recall that the matrices U and V' can be thought of as rotations and D as a rescaling of the principal components. So, the transformation F preserves the rotational elements of X without the rescaling by D and $FX = UV'$ as its singular value decomposition (with singular values of 1).

The reason this is useful is because when a matrix X has correlation, then the i th singular value of X captures the residual variance of X explained by the i th principal component after partialling out the variance explained by the first $i - 1$ principal components. So, when there is high correlation within X , less than K principal components effectively explain the variation in X and so the later (and therefore lower) singular values shrink to zero. F inflates the lowest singular values of X so that each of the principal components of the transformed FX explains the variance in FX equally. In that sense, FX is de-correlated and, for $K < n$, is mechanically irrepresentable. The cost is that this effective re-weighting of the data also amplifies the noise associated with the observations that would have had the lowest original singular values. Of course if the amplification is too strong, it can hinder efficiency of LASSO in finite sample and even prevent the sign consistency of LASSO, in the worst case.¹⁰

Our setting is particularly amenable to the Puffer transformation since the marginal effects design matrices are highly structured. In particular, the assignment probabilities to the various unique treatments are given, and as a result, the correlations with X are bounded away from 1. This has the implication that the minimum singular value is bounded below so that under standard assumptions on data generation, LASSO selection is sign consistent. While this is guaranteed for a sample size that grows in fixed K , the more important test is whether it works when K goes up with n ; we need to show that the Puffer transformation does not destroy the sign consistency of LASSO selection as the minimal singular value of X goes to zero as a function of K . In [Lemma A.1](#), we bound the rate at which the minimal singular value of X can go to zero as a function of K in a crossed RCT such as ours and [Proposition 2.2](#) below relies on this lemma to then prove that the Puffer transformation ensures irrepresentability and consistent estimation by LASSO in our context.

We make the following additional assumptions and discuss their restrictiveness below.

ASSUMPTION 2 (Design growth). $R \geq 3$, $K < n$, and $K = O(n^\gamma)$ for some $\gamma \in [0, 1/2)$.¹¹

¹⁰In $K > n$ cases—not studied here and not having a full characterization in the literature—even irrepresentability is not immediate and the theory developed is only for special cases (a uniform distribution on the Stiefel manifold) and a collection of empirically relevant simulations ([Jia and Rohe, 2015](#)).

¹¹This ensures support consistency ([Proposition 2.1](#)) at exponential rates. It also implies that $K^2 \log(K) = o(n)$, which allows for post-LASSO inference under a normal distribution ([Proposition 2.3](#)). The latter requirement stems from the growth rate of K needing to be tempered for the Central Limit Theorem to operate in this growing parameter regime.

ASSUMPTION 3 (Minimal marginal effect size). $|S_\alpha| < K$ and $\min_{k \in S_\alpha} |\alpha_k| > c > 0$ for c fixed in n .

ASSUMPTION 4 (Homoskedasticity). $\epsilon_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$, with $\sigma^2 > 0$ fixed in n .

ASSUMPTION 5 (Penalty sequence). Take a sequence $\lambda_n \geq 0$ such that $\lambda_n = \omega(n^{-\nu})$ where $0 < \nu < \frac{1}{2} - \gamma$ ¹²

Assumption 2 restricts the growth of the problem, preventing settings with too many treatments relative to observations. Without this assumption, the correct support may not be estimated with probability tending to one, and the post-estimators may not necessarily be asymptotically normally distributed. In practice, it means that the RCT cannot have cells in the fully saturated treatment design with very few units assigned to that unique treatment combination. Assumption 3 is the conventional LASSO-sparsity assumption applied to the marginal effects formulation. It imposes that *adjacent policy variants* are either appreciably different or have no difference (i.e., the so-called “beta-min” assumption in the literature). We do not handle the case of local alternatives among adjacent variants i.e., very small yet non-zero differences, but policies that are not variants of each other or are nowhere adjacent are allowed to be local alternatives as discussed in Section 2.5. Assumption 4 places our theory under homoskedastic errors following the literature on Puffer transformation. Extension to heteroskedasticity is left for future work. Finally, Assumption 5 imposes a restriction on the LASSO-penalties, standard in the regularization literature.

PROPOSITION 2.1. Assume 1-5. Let $\tilde{\alpha}$ be the estimator of (2.5) by LASSO:

$$\tilde{\alpha} := \operatorname{argmin}_{a \in \mathbb{R}^K} \|Fy - FXa\|_2^2 + \lambda_n \|a\|_1.$$

Then $P(\operatorname{sign}(\tilde{\alpha}) = \operatorname{sign}(\alpha^0)) = 1 - \exp(-\omega(n^{(1-2(\nu+\gamma))})) \rightarrow 1$.

In other words the correct support of (2.4) is selected with probability tending to 1 exponentially fast in n .

All proofs are in Appendix A unless otherwise noted.

2.3. Consistency of the TVA Estimator. Having constructed an estimator \hat{S}_α of the support S_α , the next step is to use Algorithm 2 in Online Appendix B to construct \hat{S}_{TVA} , the estimated set of pooled and pruned unique policies, and then estimate policy effects.¹³ The regression of interest is (2.2). We show this estimator is consistent.¹⁴

¹² $\omega(\cdot)$ (“small omega” notation, from the same family of notations as “big O” notation) denotes an asymptotically loose lower bound. Formally $f(n) = \omega(g(n))$ if and only if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = \infty$.

¹³Algorithm 2 constructing \hat{S}_{TVA} generalizes the Hasse concatenation examples in Section 2.1.1.

¹⁴We thank Adel Javanmard for a helpful discussion of the proof.

PROPOSITION 2.2. *Assume 1-5. Let $\hat{\eta}_{\hat{S}_{TVA}}$ be the post-Puffer LASSO OLS estimator of (2.2) on support \hat{S}_{TVA} . Then, with probability at least $1 - 2e^{-\frac{n^{1-2\gamma}\lambda^2}{2\sigma^2} + \gamma \log n} = 1 - e^{-\omega(n^{(1-2(\nu+\gamma)))})} \rightarrow 1$,*

$$\left\| \hat{\eta}_{\hat{S}_{TVA}} - \eta_{S_{TVA}}^0 \right\|_{\infty} \leq \sqrt{\frac{\log n}{n^{1-\gamma/2}}}.$$

2.4. Asymptotic Normality. The post-Puffer LASSO estimators are asymptotically normally distributed (pointwise) for the following reason. If the correct support, S_{TVA} were always selected, mechanically the estimators are asymptotically normal.

So, in practice, we need to worry about two errors: (a) the asymptotic distribution of the estimator with some incorrect support being selected and (b) the asymptotic distribution of the true estimator when the incorrect support is selected. We show in Appendix A that both of these terms are vanishing in our setup.

Intuitively, the second term can be ignored. After all, the true estimator itself is asymptotically normally distributed, so given the very unlikely event of incorrect selection, this term is asymptotically negligible. The first term requires more work. But again, one can show that the amount of potential bias accumulated due to selecting the wrong support is slow relative to the rate of actually estimating the wrong support.¹⁵

Given that these errors can be controlled, we show that the estimator is asymptotically normally distributed. Note that since the parameter vector is of increasing dimension, the asymptotic normality result must be stated slightly differently than in the usual way. The result states that any linear combination of any of the estimated parameters, when normalized properly, converges to the standard normal, which is the infinite dimensional analog to the Cramer-Wold device (He and Shao, 2000).

PROPOSITION 2.3. *Assume 1- 5. Then for $\hat{\eta}_{\hat{S}_{TVA}}$, the post-Puffer LASSO selection OLS estimator of (2.2) performed on support \hat{S}_{TVA} , we have*

$$\sqrt{nc'} \left(\hat{\eta}_{\hat{S}_{TVA}} - \eta_{\hat{S}_{TVA}}^0 \right) / (\sigma \|c\|) \rightsquigarrow \mathcal{N}(0, 1).$$

for any $c \in \mathbb{R}^{|\hat{S}_{TVA}|}$.

¹⁵An entirely different approach would be to use a recent focus in the literature on exact post-selection inference using the observation that the LASSO procedure to select a model generates a polyhedral conditioning set Lee et al. (2016). This generates a parameter estimator distribution that is a truncated, rather than complete, normal. In our special environment—a correctly specified linear model, sparse parameters, restrictions on shrinkage rate of minimal values of parameters on the support—the truncation points diverge when conditioning on the event that the true model is the estimated model. In the winner’s curse context an analogous point is made in Andrews et al. (2021), Proposition 3. This means that the distribution returns to the usual Gaussian. However, we provide a simpler, direct argument where we can calculate the distribution when the correct support is selected and bound the problematic terms in the event of poor selection.

The proof applies the central limit theorem of [He and Shao \(2000\)](#) for the growing number of parameters regime after controlling the events that the wrong support is selected as in [Javanmard and Montanari \(2013\)](#).¹⁶

It is well-known that one cannot uniformly (over the parameter space) build post-selection asymptotic distributions ([Leeb and Pötscher, 2005, 2008](#)). This is the subject of much discussion of a larger literature on post-selection inference—interpretations of the post-estimation procedures and its practical function ([Berk et al., 2013; Tibshirani et al., 2016; Lee et al., 2016](#)). In our context, several remarks are worth making. First, our claim is about pointwise inference, not uniformity over the parameter space. Second, we have nothing to say conditional on incorrect selection, hence the non-uniformity. Still, no matter what model is selected—even if an incorrect one—since in our setting the regressors are always orthogonal, there is some valid post-selection interpretation in the sense of [Berk et al. \(2013\)](#), but we do not characterize what occurs in the vanishing probability events. Third, as we recover the support with probability tending to one, and at an exponential rate, in a practical sense the non-uniformity occurs only for very small (local to zero) alternatives in the space of marginals, which are assumed away per Assumption 3.¹⁷ Loosely, recall that the non-uniformity comes up when the probability of correct selection does not go to one, or along the sequence is local to the event of failed selection. Given the very high rate of correct selection (tending to one exponentially fast in n), these unsupported local alternatives must be exceedingly close to the true parameter (the sequence of alternatives converging to the true parameter at very fast rate in n). See analogous discussion in [McCloskey \(2020\)](#) and the discussion of (A.1) in that paper.

Indeed, consistent with the theoretical results, as we will show in Online Appendix E.1, the estimators look normal in practice indicating that the non-uniformity concerns are likely to not be large in at least many practical cases, in our specific setting. Further, in our setting, since the elements with the highest effects tend to be selected first, and because of orthogonality, in practice the large parameter estimates almost always perform well.

2.5. The Effect of the Best Policy. The TVA procedure generates a set of pruned and pooled policies \hat{S}_{TVA} with post-LASSO estimates $\hat{\eta}_{\hat{S}_{TVA}}$. This full set of pooled estimates is of direct interest to the policymaker. We now propose an optional last step, in case a policymaker is particularly interested in a single “best” policy (based on the sample estimates) and its estimated performance. This policy is the one in \hat{S}_{TVA} with the highest post-LASSO estimate:

$$\hat{\kappa}^* = \operatorname{argmax}_{\kappa \in \hat{S}_{TVA}} \hat{\eta}_{\hat{S}_{TVA}, \kappa}$$

¹⁶We again thank Adel Javanmard for a helpful discussion of the proof.

¹⁷We are grateful to Adam McCloskey for pointing this out.

If the true best policy has a (population level) treatment effect that far exceeds that of next best policy, then with high probability $\hat{\kappa}^*$ is that policy and the post-LASSO estimator $\hat{\eta}_{\hat{S}_{TVA, \hat{\kappa}^*}}$ (and corresponding confidence interval) is its most efficient unbiased estimator. In that case, the policy that emerge is the best policy, and no further adjustment is needed.

However, if there are other non-pooled policies that perform similarly to the true best policy, there can be a bias in the estimated effectiveness of the “best” among them because of the winner’s curse (Andrews et al. (2021)). Specifically, when another policy $\kappa \in \hat{S}_{TVA}$ has a treatment effect close to $\hat{\kappa}^*$, in the sense that $|\eta_{\hat{S}_{TVA, \hat{\kappa}^*}}^0 - \eta_{\hat{S}_{TVA, \kappa}}^0|$ is small, there is a chance (due to sampling variation) that they may be incorrectly ranked in the sample: i.e., $\hat{\eta}_{\hat{S}_{TVA, \hat{\kappa}^*}} > \hat{\eta}_{\hat{S}_{TVA, \kappa}}$ even though $\eta_{\hat{S}_{TVA, \hat{\kappa}^*}}^0 < \eta_{\hat{S}_{TVA, \kappa}}^0$. In these close races sampling variation can end up determining the estimated best policy. Importantly, even if the ranking is correct, $\hat{\eta}_{\hat{S}_{TVA, \hat{\kappa}^*}}$ will be biased upwards relative to the true effect $\eta_{\hat{S}_{TVA, \hat{\kappa}^*}}^0$.

Given that such a winner’s curse arises only in close races, and our assumption of sparse marginal effects (Assumption 3) and steps of pruning and pooling, the winner’s curse arises primarily when *non-variant* policies have similar performance. For example, Assumption 3 doesn’t prevent the pooled policy (Any information hubs, No SMS, No incentives) from having a similar impact as (No seeds, Any reminders, Any slope incentives), as these emerge from marginals in different treatment profiles, and can’t be pooled. Were these policies to both have similar and near maximum impacts, then a winner’s curse can arise in their horserace.¹⁸

The relevant notion of “similar effect” is a *local alternative*:

DEFINITION 2.1. *Pooled policies κ, κ' are local alternatives if there is a constant $r_{\kappa\kappa'}$ fixed in n such that*

$$(2.6) \quad \eta_{S_{TVA, \kappa}}^0 = \eta_{S_{TVA, \kappa'}}^0 + \frac{r_{\kappa\kappa'}}{\sqrt{n}}.$$

This definition holds asymptotically, in that local alternatives have similar effects for *any* sample size. As we have shown, there can be several local alternatives when Assumption 3 holds. We assume a fixed number of local alternatives, and that these are the policies with the most (mutually) similar effects.

ASSUMPTION 6. *There are at most $q < \infty$, independent of n , pairs of local alternatives, i.e., pairs $\kappa, \kappa' \in S_{TVA}$ such that $\eta_{S_{TVA, \kappa}}^0 = \eta_{S_{TVA, \kappa'}}^0 + \frac{r_{\kappa\kappa'}}{\sqrt{n}}$ where $r_{\kappa\kappa'}$ fixed in n . All other policy pairs are further separated, i.e.. separated by at least $\omega(\frac{1}{\sqrt{n}})$.*

This nests $q = 0$, when there are no local alternatives. The best policy is then well separated from the next best policy and the post-LASSO estimates $\hat{\eta}_{\hat{S}_{TVA, \kappa^*}}$ don’t need any

¹⁸In fact, even within a treatment profile, two pooled policies can still have similar treatment effects, in spite of Assumption 3. It can occur if they are activated by marginals from sufficiently different parts of the Hasse diagram. All these cases are formalized in Proposition G.1 of the Online Appendix G.

adjustment. One is free to assume this or assess it ex-post. For $q > 0$, there is a chance that the best policy is local alternative to others, so there may be a winner’s curse.

We apply the hybrid estimator of Andrews et al. (2021) which attenuates the post-LASSO estimate. It balances performance with a small amount of median bias tolerance. Note that the main text of Andrews et al. (2021) requires that the estimators are *exactly* jointly normally distributed. They present two extensions, one for a conditioning event such as model-selection and the other for asymptotic normality. They do not work out the case with both issues present, as in our case. So we proceed with the assumption as in their main text, treating $\hat{\eta}_{\hat{S}_{TVA}}$ as normally distributed (which our simulations support as a reasonable approximation). Extending their work is beyond the scope of this paper.

We can apply Proposition 6 of Andrews et al. (2021) directly. This means we can pick two significance levels, $\alpha > \beta$ and use this to characterize confidence intervals and bias for the hybrid estimator. The hybrid estimator will be median unbiased (with a bias bounded by $\beta/2$). The (conditional) confidence interval has coverage $(1 - \alpha)/(1 - \beta)$. See Appendix G and Andrews et al. (2021) for details.

The winner’s curse-corrected estimators of Andrews et al. (2021) are not a cure-all. They cannot correct for ordinal misrankings of policies, and work for case in which the similar effect policies are $\Theta(\frac{1}{\sqrt{n}})$ apart.¹⁹ Nevertheless, we include discussion of such an adjustment because it helps debias estimates and lets us err on the side of conservative inference. The hybrid estimator has the appealing property that it rapidly converges to the post-LASSO estimate whenever the best policy is well separated from the next best policy in the sample. We can thus view it as insurance whenever we estimate the best policy, letting the data decide if the winner’s curse risk is appreciable and automatically correcting for it to the extent possible when it arises.

2.6. Summary of TVA. A summary of the overall procedure is presented in Algorithm 1.

3. SIMULATION PERFORMANCE

Here we run simulations in the environment described in Section 2 – namely, when a sparse set of policies have meaningful and meaningfully different impacts. In Sections 3.1 and 3.2 we outline the simulation setup and performance measures and show in Section 3.3 that TVA outperforms several other standard approaches. The relative deficiencies of other estimators also highlight the features that give TVA its edge.

3.1. Simulation Setup. The idea of the simulation setup is to generate simulated design matrices from marginal specifications (2.4) that resemble the data, score these on certain metrics, and aggregate these scores into measures of performance for sample size n .

¹⁹Note that conditional on correct support selection $\hat{S}_{TVA} = S_{TVA}$, the post-LASSO estimates of the best policy, even if it is selected incorrectly, must be $\Theta(1/\sqrt{n})$ of the true best policy effect.

Algorithm 1: Estimating Treatment Effects by Treatment Variant Aggregation

- (1) Given treatment assignment matrix T , calculate the treatment profile and marginal dosage intensity matrix X .
 - (2) Estimate $\widehat{S}_\alpha := \{j \in [K] : |\tilde{\alpha}_j| \neq 0\}$ by estimating (2.4) through a Puffer transformed LASSO.
 - (3) Calculate marginal effects support \widehat{S}_{TVA} from \widehat{S}_α using Algorithm 2 in Online Appendix B.
 - (4) Estimate pooled and pruned treatment effects of unique (relevant) policies, $\hat{\eta}_{\widehat{S}_{TVA}}$, using regression (2.2).
 - (5) Optional: To estimate the best policy in \widehat{S}_{TVA} , select $\hat{\kappa}^* = \operatorname{argmax}_{\kappa \in \widehat{S}_{TVA}} \hat{\eta}_{\widehat{S}_{TVA}, \kappa}$. Either report the OLS estimate $\hat{\eta}_{\widehat{S}_{TVA}, \hat{\kappa}^*}$, or, if anticipating a winner's curse, construct the hybrid estimator $\hat{\eta}_{\widehat{S}_{TVA}, \hat{\kappa}^*}^{\text{hyb}}$ with nominal size α and median bias tolerance $\beta/2$.
-

- (1) Fix $R = (5, 5, 3)$, $M = 3$ and $\sigma = 2.3$: parameters are chosen to loosely mimic our experiment where 3 treatment arms have asymmetric intensities leading to 75 unique policies and where σ is chosen such that the R^2 of the post-LASSO regression matches the experiment for a similar sample size.
- (2) The simulation results are plots of performance $\hat{m}(n)$ against sample size n where n ranges between 1,000 and 10,000.²⁰
- (3) These scores $\hat{m}(n)$ are generically computed as follows.
 - (a) A set \mathcal{C} of true supports of the marginal specification (2.4) is randomly chosen. Each member $S_\alpha^i \in \mathcal{C}$ is a particular support or “configuration” and each configuration has fixed support size $|S_\alpha^i| = M$. Specifically, each configuration \mathcal{C} is constructed by randomly sampling M covariates of X . Furthermore, if $S_\alpha^i = (k_1, k_2, \dots, k_M)$ in some given order, we assign coefficients $\alpha_{k_j} = 1 + 4 \cdot \frac{j-1}{M-1}$. That is, these nonzero coefficients are linearly spaced between 1 and 5. Thus each configuration fully specifies the set of coefficients α for (2.4).
 - (b) For each $S_\alpha^i \in \mathcal{C}$, a set $\mathcal{S}_{S_\alpha^i}(n)$ of simulations (design matrices) is generated based on the coefficients specified by the configuration, and the Gaussian noise, with sample size n . For each simulation $\hat{s}(n) \in \mathcal{S}_{S_\alpha^i}(n)$, it is scored by a metric $m(\hat{s}(n))$ that will be specified.
 - (c) These scores are aggregated over simulations $\mathcal{S}_{S_\alpha^i}(n)$, and then aggregated again over configurations \mathcal{C} , to produce the aggregated performance score $\hat{m}(n)$.

3.2. Performance Measures. Denoting by $\widehat{S}_\alpha^i(\hat{s}(n))$ the model selection estimator for S_α^i for simulation $\hat{s}(n)$, we use the following performance metrics throughout our simulations.

²⁰For some computationally intensive simulations n is logarithmically spaced.

Support selection accuracy:

$$m(\hat{s}(n)) := \frac{|\hat{S}_\alpha^i(\hat{s}(n)) \cap S_\alpha^i|}{|\hat{S}_\alpha^i(\hat{s}(n)) \cup S_\alpha^i|}.$$

This is a value between 0 and 1 that increases with support coverage, and is 1 if and only if the support is correctly selected. To construct the aggregated metric $\hat{m}(n)$ it is averaged over the simulations per configuration, and then averaged again over configurations.

“Some” best policy inclusion accuracy:

$$m(\hat{s}(n)) = \begin{cases} 1 & \text{if } \hat{\kappa}^{i*}(\hat{s}(n)) \cap \kappa^{i*} \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

where $\kappa^{i*} = \operatorname{argmax}_{\kappa \in S_{TVA}^i} \eta_{S_{TVA}, \kappa}$ denotes the true best pooled policy in the marginal effects support S_{TVA}^i (uniquely determined from S_α^i). This measure is again averaged over simulations per configuration, and then averaged over configurations. The final metric therefore gives the share of simulations per n where at least one true best policy was pooled into the estimated best pooled policy.

Minimum dosage best policy inclusion accuracy:

$$m(\hat{s}(n)) = \begin{cases} 1 & \text{if } k^{i*\min} \in \hat{\kappa}^{i*}(\hat{s}(n)) \\ 0 & \text{otherwise} \end{cases}$$

where $k^{i*\min}$ denotes the true minimum dosage best policy.²¹ Once aggregated this measure captures the share of simulations, per n where the minimum dosage best policy was included in the estimated best pool.

Mean squared error (of best policy effect): For each simulation $\hat{s}(n)$, the estimated best policy treatment effect is scored by its error with respect to the true treatment effect:

$$m(\hat{s}(n)) := \hat{\eta}_{\hat{S}_{TVA}, \hat{\kappa}^*}^{\text{hyb}} - \eta_{S_{TVA}, \kappa^*}.$$

And thus $\hat{m}(n)$ is simply the estimated MSE:

$$\hat{m}(n) := \frac{1}{|\mathcal{C}|} \sum_{\mathcal{C}} \frac{1}{|\mathcal{S}_{S_\alpha^i}(n)|} \sum_{\mathcal{S}_{S_\alpha^i}(n)} m^2(\hat{s}(n)).$$

²¹The “intersection” and “inclusion” operators for the best policy inclusion measures are to be understood in the following way: suppose the true best policy $\kappa^{i*} \in S_{TVA}^i$ pools together m policies as per S_α^i that we can organize into a set $S_1 = \{k_1^{i*}, \dots, k^{i*\min}, \dots, k_m^{i*}\}$. Equivalently we organize into S_2 the n policies composing the estimated best pool as per $\hat{S}_\alpha^i(n)$. Then $\hat{\kappa}^{i*}(\hat{s}(n))$ stands for S_2 and κ^{i*} for S_1 .

3.3. Performance of TVA. First of all, simulation performance of TVA attests to its main theoretical properties: support consistency, best policy estimation consistency, and normally distributed coefficient estimates. This is depicted in Figure 3 (panels A,B,D) and discussed in Online Appendix E.1. In what follows, we primarily make the case for the strong performance of TVA relative to its most straightforward alternative, a direct application of OLS. In our Online Appendix E.2 we provide comparisons for further LASSO-based alternative estimation strategies, the results of which are summarized below.

3.3.1. Direct OLS. An intuitive route for inference in this setting is estimating the unique policy specification (2.1) using only an application of OLS and nothing else (a strategy we call “direct OLS”). Since this is a fully saturated regression, this estimator has no theoretical issues with convergence nor with interpretation. Rather, this is about performance in the finite sample in the environment we describe. Specifically, it is worth noting that without additional restrictions on the parameter space, OLS yields uniformly most powerful tests. However, since our approach is different from deriving optimal tests and involves model selection, and because additional restrictions are imposed, OLS may not be the best possible method, as will be demonstrated in the simulations. Most obviously, there is a loss of power in separately estimating the impact of 75 distinct interventions.

Figure 3, panels B-D, documents three main inadequacies of direct OLS on the question of selecting and estimating best policy effects: (1) it fails at consistently identifying the minimum dosage best policy, (2) the estimates of best policy exhibit a stronger winner’s curse (3) the attenuations from applying Andrews et al. (2021) are large. More specifically, panel A shows that direct OLS (orange) does almost as well as TVA (blue) in estimating as best policy *some* policy that is part of the true best pooled policy, but it effectively picks the dosage at random; thus selection of the minimum dosage hovers at around 36.67%. Panel C exhibits a strong winner’s curse, which is to be expected in a situation with numerous candidates for “best policy”, since the odds that a particularly large shock was drawn and thrust one to the top is quite high. The fact that the resulting coefficients are over-attenuated can be inferred from Panel D, where the attenuated best policy MSE (yellow triangles) remains large (of 0.53) even for large n .²² In contrast, the winner’s curse attenuations for TVA are much more modest, because of reduction of the number of competing policies and therefore greater separation between them.

Besides the specific issue of best policy estimation, direct OLS has low power. This is depicted in Figure E.1, where simulated OLS estimates of all the unique policies (2.1) contrast with the pruned and pooled estimation (2.2) (refer to Online Appendix E).

²²Recall that while Andrews et al. (2021) estimators are consistent, this assumes some local separation of parameters, which is not guaranteed in these neck-and-neck competitions.

3.3.2. *Naive LASSO*. There are two ways we could “naively” apply LASSO. The first is to disregard pooling, and apply LASSO on the unique policy specification (2.1) because sparse dosages might also mean a sparse set of policies. While there is no theoretical issue with this procedure in terms of model consistency, using this for policy estimation leads to much the same performance limitations as direct OLS with regards to best policy estimation, namely a persistently high best policy MSE stemming from overly severe correction from Andrews et al. (2021)’s winner’s curse adjustment. Figure E.2 from our Online Appendix, contrasts this “No pooling, only pruning” version of LASSO to TVA on best policy estimation and documents these patterns in detail. The second way to “naively” apply LASSO is to consider both pooling and pruning as important, but adopt a sign inconsistent model selection procedure by applying LASSO directly on (2.4) without a Puffer transformation. As expected, simulations attest to inconsistent support selection though MSE on the best policy is comparable to TVA. Importantly, it fails to select the minimum dosage best policy with substantial probability relative to TVA (refer to Figure E.3 in Online Appendix E).

3.3.3. *Debiased LASSO*. Because we are interested in high dimensional inference²³ one alternative to a two step process of model selection and inference is the so-called “debiased LASSO” (Zhang and Zhang (2014), Javanmard and Montanari (2014), Javanmard and Montanari (2018), Van De Geer (2019)). The basic idea is that since the downward bias in LASSO is estimable, we can reverse it. A feature, however, is that these debiased coefficients are almost surely never exactly zero, so that there is no question of sparsity. We thus only need to consider applying debiased LASSO to (2.1). In Figure E.4 of our Online Appendix, we show that the debiased LASSO procedure suffers from the same limitations as direct OLS estimation, especially with regards to best policy estimation (high MSE due to over-attenuation of the winner’s curse).

3.3.4. *“Off the Shelf” Bayesian approaches: Spike and Slab LASSO (Nie and Ročková, 2022)*. The rules governing admissible pooling encodes the econometrician’s prior about the environment as does regularization. This raises the possibility of a Bayesian framework. Indeed, LASSO estimates have a Bayesian interpretation with Laplace priors. One can ask whether a more sophisticated, “explicitly” Bayesian approach can address our final objectives. This paradigmatically different route is the topic of future work. In Section E.2.4 of our Online Appendix, we just show that “off the shelf” Bayesian approaches are unlikely to help. In particular, we show that a direct application of spike and slab formulations – the most intuitively relevant method – underperforms relative to our TVA procedure with a performance pattern similar to that of applying Naive LASSO to the marginal specification (2.4).

²³Albeit, we are still in a $K < n$ regime (mechanically since the number of treatments cannot exceed the number of units), sometimes called low dimensional with diverging number of parameters.

Taken together, these simulation results make the case that TVA is both a powerful and robust candidate for our setting. Moreover, even under different sparsity and effect size relaxations, TVA is still strong as demonstrated in Online Appendix E.3. In particular it does better than the next best practical alternative of applying naive LASSO to the marginal specification (2.4).

4. INCREASING IMMUNIZATION IN HARYANA: CONTEXT, EXPERIMENTAL DESIGN, AND DATA

We now apply this method to a large-scale experiment conducted in collaboration with the government of Haryana to help them select the most effective policy bundle to stimulate demand for immunization. The objective of the experiment was explicitly to select the best policy to scale up, after one year-long experiment with 75 potentially distinct treatments, making it an excellent application. Indeed, we developed the method in order to analyze this data.

4.1. **Context.** This study took place in Haryana, a populous state in North India, bordering New Delhi. In India, a child between 12 and 23 months is considered to be fully immunized if he or she receives one dose of BCG, three doses of Oral Polio Vaccine (OPV), three doses of DPT, and at least one dose of a measles vaccination. India is one of the countries where immunization rates are puzzlingly low. According to the 2015-2016 National Family Health Survey, only 62% of children were fully immunized (NFHS, 2016). This is not due to lack of access to vaccines or health personnel. The Universal Immunization Program (UIP) provides all vaccines free of cost to beneficiaries, and vaccines are delivered in rural areas—even in the most remote villages. Immunization services have made considerable progress over the past few years and are much more reliably available than they used to be. During the course of our study we found that the monthly scheduled immunization session were almost always run in each village.

The central node of the UIP is the Primary Health Centre (PHC). PHCs are health facilities that provide health services to an average of 25 rural and semi-urban villages with about 500 households each. Under each PHC, there are approximately four sub-centres (SCs). Vaccines are stored and transported from the PHCs to either sub-centers or villages on an appointed day each month, where there is a mobile clinic where the Auxiliary Nurse Midwife (ANM) administers vaccines to all eligible children. A local health worker, the Accredited Social Health Activist (ASHA), is meant to help map eligible households, inform and motivate parents, and take them to the immunization session. She receives a small fee for each shot given to a child in her village.

Despite this elaborate infrastructure, immunization rates are particularly low in North India, especially in Haryana. According to the District Level Household and Facility Survey,

the full immunization coverage among 12-23 months-old children in Haryana fell from 60% in 2007-08 to 52.1% in 2012-13 (DLHS, 2013).

In the district where we carried out the study, a baseline study revealed even lower immunization rates (the seven districts that were selected were chosen because they have low immunization). About 86% of the children (aged 12-23 months) had received at least three vaccines. However, the share of children whose parents had reported they received the measles vaccine (the last in the sequence) was 39%, and only 19.4% had received the vaccine before the age of 15 months, while the full sequence is supposed to be completed in one year.

After several years focused on improving the supply of immunization services, the government of Haryana was interested in testing out strategies to improve household take-up of immunization, and in particular, their completion of the full immunization schedule. With support from USAID and the Gates Foundation, they entered into a partnership with J-PAL to test out different interventions. The final objective was to pick the best policy possibly scale up throughout the state.

Our study took place in seven districts where immunization was particularly low. In four districts, the full immunization rate in a cohort of children older than the ones we consider, was below 40%, as reported by parents (which is likely a large overestimate of the actual immunization rate, given that children get other kinds of shots and parents often find it hard to distinguish between them, as noted in Banerjee et al. (2021)). Together, the districts cover a population of more than 8 million (8,280,591) in more than 2360 villages, served by 140 PHCs and 755 SCs. The study covered all these PHCs and SCs, and are thus fully representative of the seven districts. Given the scale of the project, our first step was to build a platform to keep a record of all immunizations. Sana, an MIT-based health technology group, built a simple m-health application that the ANMs used to register and record information about every child who attended at least one camp in the sample villages. Children were given a unique ID that made it possible to track them across visits and centers. Overall, 295,038 unique children were recorded in the system, and 471,608 vaccines were administered. Data from this administrative database is our main source of information on immunization and we discuss its reliability below. More details on the implementation are provided in the publicly available progress report (Banerjee et al., 2021).

4.2. Interventions. The study evaluates the impact of several nudges on the demand for immunization: small incentives, targeted reminders, and local ambassadors.

4.2.1. Incentives. When households are indifferent or have a propensity to procrastinate, small incentives can offset any short term cost of getting to an immunization camp and lead to a large effect on immunization. Banerjee et al. (2010) shows that small incentives for immunization in Rajasthan (a bag of lentils for each shot and a set of plates for completing the course) led to a large increase in the rates of immunization. Similar results were subsequently

obtained in other countries, suggesting that incentives tend to be effective (Bassani et al., 2013; Gibson et al., 2017),(Chandir et al., 2022). In the Indian health system, households receive incentives for a number of health behavior, including hospital delivery, pre-natal care visits, and, in some states (like Tamil Nadu), immunization.

The Haryana government was interested in experimenting with incentives. The incentives that were chosen were mobile recharges for pre-paid phones, which can be done cheaply and reliably on a very large scale. Almost all families have at least one phone and the overwhelming majority of the phones are pre-paid. Mobile phone credits are of uniform quality and fixed price, which greatly simplify procurement and delivery.

A small value of mobile phone credit was given to the caregivers each time they brought their child to get immunized. Any child under the age of 12 months receiving one of the five eligible shots (i.e., BCG, Penta-1, Penta-2, Penta-3, or Measles-1), was considered eligible for the incentives intervention. Mobile recharges were delivered directly to the caregivers' phone number that they provided at the immunization camp. Seventy (out of the 140) PHCs were randomly selected to receive the incentives treatment.

In Banerjee et al. (2010), only one reward schedule was experimented with. It involved a flat reward for each shot plus a set of plates for completing the immunization program. This left many important policy questions pending: does the level of incentive make a difference? If not, cheaper incentives could be used. Should the level of rewards increase with each immunization to offset the propensity of the household to drop out later in the program?

To answer these questions, we varied the level of incentives and whether they increased over the course of the immunization program. The randomization was carried out within each PHC, at the subcenter level. Depending on which sub-center the caregiver fell under, she would receive one of the following.

- (1) Flat incentive, high: INR 90 (\$1.34 at the 2016 exchange rate, \$4.50 at PPP) per immunization (INR 450 total).
- (2) Sloped incentive, high: INR 50 for each of the first three immunizations, 100 for the fourth, 200 for the fifth (INR 450 total).
- (3) Flat incentive, low: INR 50 per payment (INR 250 total).
- (4) Sloped incentive, low: INR 10 for each of the first three immunizations, 60 for the fourth, 160 for the fifth (INR 250 total).

Even the high incentive levels here are small and therefore implementable at scale, but they still constitute a non-trivial amount for the households. The high incentive level was chosen to be roughly equivalent to the level of incentive chosen in the Rajasthan study: INR 90 was roughly the cost of a kilogram of lentils in Haryana during our study period. The low level was meant to be half of that (rounded to INR 50 since the vendor could not deliver recharges that were not multiple of 10). This was meaningful to the households: INR 50

corresponds to 100 minutes of talk time on average. The provision of incentives was linked to each vaccine. If a child missed a dose, for example Penta-1, but then came for the next vaccine (in this case, measles), they would receive both Penta-1 and measles and get the incentives for both at once, as per the schedule described above.

To diffuse the information on incentives, posters were provided to ANMs, who were asked to put them up when they set up for each immunization session. The village ASHAs and the ANMs were also supposed to inform potential beneficiaries of the incentive structure and amount in the relevant villages. However, there was no systematic large scale information campaign, and it is possible that not everybody was aware of the presence or the schedule of the incentives, particularly if they had never gone to a camp.

4.2.2. *Reminders.* Another frequently proposed method to increase immunization is to send text message reminders to parents. Busy parents have limited attention and reminders can put the immunization back at the “top of the mind.” Moreover, parents do not necessarily understand that the last immunization in the schedule (measles) is for a different disease and is at least as important as the previous ones. SMSs are also extremely cheap and easy to administer in a population with widespread access to cell phones. Even if not everyone gets the message, the diffusion may be reinforced by social learning, leading to faster adoption.²⁴

The potential for SMS reminders is recognized in India. The Indian Academy of Pediatrics rolled out a program in which parents could enroll to get reminders by providing their cell phone number and their child’s date of birth. Supported by the Government of India, the platform planned to enroll 20 million children by the end of 2020.

Indeed, text messages have already been shown to be effective to increase immunization in some contexts. For example, a systematic review of five RCTs finds that reminders for immunization increase take up on average (Mekonnen et al., 2019). However, it remains true that text messages could have no effect or even backfire if parents do not understand the information provided and feel they have no one to ask (Banerjee et al., 2018). Targeted text and voice call reminders were sent to the caregivers to remind them that their child was due to receive a specific shot. To identify any potential spillover to the rest of the network, this intervention followed a two step randomization. First, we randomized the study sub-centers into three groups: no reminders, 33% reminders, and 66% reminders. Second, after their first visit to that sub-center, children’s families were randomly assigned to either get the reminder or not, with a probability corresponding to the treatment group for their sub-centers. The children were assigned to receive/not receive reminders on a rolling basis.

²⁴See, e.g., Rogers (1995); Krackhardt (1996); Kempe, Kleinberg, and Tardos (2003); Jackson (2008); Iyengar, den Bulte, and Valente (2010); Hinz, Skiera, Barrot, and Becker (2011); Katona, Zubcsek, and Sarvary (2011); Jackson and Yariv (2011); Banerjee, Chandrasekhar, Duflo, and Jackson (2013); Bloch and Tebaldi (2016); Jackson (2017); Akbarpour, Malladi, and Saberi (2017).

The following text reminders were sent to the beneficiaries eligible to receive a reminder. In addition, to make sure that the message would reach illiterate parents, the same message was sent through an automated voice call.

(1) Reminders in incentive-treatment PHCs:

“Hello! It is time to get the «name of vaccine» vaccine administered for your child «name». Please visit your nearest immunization camp to get this vaccine and protect your child from diseases. You will receive mobile credit worth «range for slope or fixed amount for flat» as a reward for immunizing your child.”

(2) Reminders in incentive-control PHCs:

“Hello! It is time to get the «name of vaccine» vaccine administered for your child. Please visit your nearest immunization camp to get this vaccine and protect your child from diseases.”

4.2.3. *The Immunization Ambassador: Network-Based Seeding.* The goal of the ambassador intervention was to leverage the social network to spread information. The objective was to identify influential individuals who could relay to villagers both the information on the existence of the immunization camps, and, wherever relevant, the information that incentives were available. Existing evidence shows that people who have a high centrality in a network (e.g., they have many friends who themselves have many friends) are able to spread information more widely in the community (Katz and Lazarsfeld, 1955; Aral and Walker, 2012; Banerjee et al., 2013; Beaman et al., 2018; Banerjee et al., 2019). Further, members in the social network are able to easily identify individuals, whom we call information hubs, who are the best placed to diffuse information as a result of their centrality as well other personal characteristics (social mindedness, garrulousness, etc.) (Banerjee et al., 2019).

This intervention took place in a subset of 915 villages where we collected a full census of the population (see below for data sources). Seventeen respondents in each village were randomly sampled from the census to participate in the survey, and were asked to identify people with certain characteristics (more about those later). Within each village, the six people nominated most often by the group of 17 were recruited to be ambassadors for the program. If they agreed, a short survey was conducted to collect some demographic variables, and they were then formally asked to become program ambassadors. Specifically, they agreed to receive one text message and one voice call every month, and to relay it to their friends. In villages without incentives, the text message was a bland reminder of the value of immunization. In villages with incentives, the text message further reminded the ambassador (and hence potentially their contacts) that there was an incentive for immunization.

While our previous research had shown that villagers can reliably identify information hubs, a pertinent question for policy unanswered by previous work is whether the information

hubs can effectively transmit messages about health, where trust in the messengers may be more important than in the case of more commercial messages.

There were four groups of ambassador villages, which varied in the type of people that the 17 surveyed households were asked to identify. The full text is in Appendix L.

- (1) *Random* seeds: In this treatment arm, we did not survey villages. We picked six ambassadors randomly from the census.
- (2) *Information hub* seed: Respondents were asked to identify who is good at relaying information.
- (3) *Trusted* seed: Respondents were asked to identify those who are generally trusted to provide good advice about health or agricultural questions
- (4) *Trusted information hub* seed: Respondents were asked to identify who is both trusted and good at transmitting information

4.3. Experimental Design. The government was interested in selecting the best policy, or bundle of policies, for possible future scale up. We were agnostic as to the relative merits of the many available variants. However, we believed that there could be significant interactions between different policies. For example, our prior was that the ambassador intervention was going to work more effectively in villages with incentives, because the message to diffuse was clear. We therefore implemented a completely cross-randomized design, as illustrated in our Online Appendix Figure H.1.

We started with 2,360 villages, covered by 140 PHCs, and 755 sub-centers. The 140 PHCs were randomly divided into 70 incentives PHCs, and 70 no incentives PHCs (stratifying by district). Within the 70 incentives PHCs, we randomly selected the sub-centers to be allocated to each of the four incentive sub-treatment arms. Finally, we only had resources to conduct a census and a baseline exercise in about 900 villages. We selected about half of the villages from the coverage area of each subcenter, after excluding the smallest villages. Only among the 915 villages did we conduct the ambassador randomization: after stratifying by sub-center, we randomly allocated the 915 villages to the control group (no ambassador) or one of the four ambassador treatment groups.

In total, we had one control group, four types of incentives interventions, four types of ambassador interventions, and two types of SMS interventions. Since they were fully cross-randomized (in the sample of 915 villages), we had 75 potential policies, which is large even in relation to our relatively large sample size. Our goal is to identify the most effective and cost-effective policies and to provide externally valid estimates of the best policy’s impact, after accounting for the winner’s curse problem. Further, we want like to identify other effective policies and answer the question of whether different variants of the policy had the same or different impacts.

4.4. Data.

4.4.1. *Census and Baseline.* In the absence of a comprehensive sampling frame, we conducted a mapping and census exercise across 915 villages falling within the 140 sample PHCs. To conduct the census, we visited 328,058 households, of which 62,548 households satisfied our eligibility criterion (children aged 12 to 18 months). These exercises were carried out between May and November 2015. The data from the census was used to sample eligible households for a baseline survey. We also used the census to sample the respondent of the ambassador identification survey (and to sample the ambassadors in the “random seed” villages). Around 15 households per village were sampled, resulting in data on 14,760 households and 17,000 children. The baseline survey collected data on demographic characteristics, immunization history, attitudes and knowledge and was conducted between May and July 2016. A village-level summary of baseline survey data is given in Appendix K.

4.4.2. *Outcome Data.* Our outcomes of interest are the number of vaccines administered for each vaccine every month, and the number of fully immunized children every month. The main analysis of this paper focuses on the number of children who received the measles vaccines in each village every month. The measles vaccine is the last vaccine in the immunization schedule, and the ANMs check the immunization history and administer missing vaccines when a child is brought in for this vaccine. As a result, it is a good proxy for a child being fully immunized.

For our analysis, we use administrative data collected by the ANM using the e-health application on the tablet, stored on the server, to measure immunization. At the first visit, a child was registered using a government provided ID (or in its absence, a program-generated ID) and past immunization history, if any. In subsequent visits, the unique ID was used to pull-up the child’s details and update the data. Over the course of the program, about 295,038 children were registered, yielding a record of 471,608 immunizations. We use the data from December 2016 to November 2017. We do this because of a technical glitch in the system—the SMS intervention was discontinued from November 2017, although the incentives and information hub interventions were continued a little longer, through March 2018.

Since this data was also used to trigger SMS reminders and incentives, and for the government to evaluate the nurses’ performance,²⁵ it was important to assess its accuracy. Hence, we conducted a validation exercise, comparing the administrative data with random checks, as described in Online Appendix J. The data quality appears to be excellent. Finally, one concern (particularly with the incentive program) is that the intervention led to a pattern of substitution, with children who would have been immunized elsewhere (in the PHC or at the hospital) choosing to be immunized in the camp instead. To address this issue, we collected

²⁵Aggregated monthly reports generated from this data replaced the monthly reports previously compiled by hand by the nurses.

data immediately after the intervention on a sample of children who did not appear in the database (identified through a census exercise), to ascertain the status of their immunization. In Appendix I, we show that there does not appear to be a pattern of substitution, as these children were not more likely to be immunized elsewhere.

Below, the dependent variable is the number of measles shot given in a village in a month (each month, one immunization session is held at each site). On average, in the entire sample, 6.16 measles shot were delivered per village every month (5.29 in the villages with no intervention at all). In the sample at risk for the ambassador intervention (which is our sample for this study) 6.94 shots per village per month were delivered.

4.5. Interventions Average Effects. In this section, we present the average effects of the interventions using a standard regression without interactions.

We focus on the sample of census villages used throughout our analysis - which are the villages where the ambassador intervention was also administered - and run the following specification:

$$y_{dsvt} = \alpha + \beta' \text{Incentive}_s + \gamma' \text{SMS}_s + \delta' \text{Ambassador}_v + u_{dt} + \epsilon_{dsvt}.$$

We weight these village-level regressions by village population, and standard errors are clustered at the SC level.²⁶

The results (reported in Banerjee et al. (2019)) are depicted graphically in Figure 4 and show that, on average, using information hubs (“gossips” in that paper) as ambassadors has positive effects on immunization: 1.89 more children receive a measles vaccine on a base of 7.32 in control in this sample ($p = 0.04$). This is nearly identical to the effect of the high-powered, sloped incentive, though this intervention is considerably cheaper. In contrast, none of the other ambassador treatments—random seeding, seeding with trusted individuals, or seeding with trusted information hubs—have benefits statistically distinguishable from zero ($p = 0.42$, $p = 0.63$, and $p = 0.92$ respectively) and the point estimates are small, as well. To ensure that conclusions are not simply an artifact of this particular subsample, we show in Online Appendix H that these results are robust to running the analysis on the full sample.

The conclusion from this analysis is that financial incentives can be effective to boost demand for immunization, but only if they are large enough and increase with each immunization. Of the two cheaper interventions, the SMS interventions, promoted widely in India and elsewhere, seem disappointing. These results are similar to those found in Pakistan by Chandir et al. (2022) which also test various models of incentives and SMS, with small

²⁶This is the highest level at which a treatment is administered, so clustering at this level should yield the most conservative estimate of variance. In practice clustering at the village level or SC level does not make an appreciable difference.

differences. There, SMS alone has a statistically significant, but small impact on full immunization. Incentives are more effective, and the level of incentives matter (although in that setting, the slope does not matter).

In our setting, leveraging the community by enrolling local ambassadors, selected using the cheap procedure of asking a few villages who are good information hubs, seems to be as effective as using incentives. It leads to an increase of 26% in the number of children who complete the schedule of immunization every month. This alone could increase full immunization rate in those districts from 39% (our baseline full immunization rate, as reported by parents) to nearly 49%. This analysis does not fully answer the policymaker’s question, however. It could well be that the interventions have powerful interactions with each other, which has two implications. First, the main effect, as estimated, does not tell us what the impact of the policy would be in Haryana if implemented alone (because as it is, they are a weighted average of a complicated set of interacted treatments). Second, it is possible that the government could do better by combining two (or more) interventions. For example, our prior in designing the information hub ambassador intervention (described in our proposal for the project)²⁷ was that it would have a positive interaction effect with incentives, because it would be much easier for the information hubs to relay hard information (there are incentives) than a vaguer message that immunization is useful. The problem, however, is that there are a large number of interventions and interactions: we did not—nor was it feasible to—think through ex-ante all of the interactions that should or should not be included, which is why in Banerjee et al. (2019), we only reported the average effects of each different type of seeds in the entire sample, without interactions. In the next section, we adapt our disciplined approach to select which ones to include, and to then estimate the impact of the “best” policy.

5. RESULTS

5.1. Identifying effective policies.

5.1.1. *Method.* We adapt the TVA procedure for our case. We allow only some pooling within arms depending on the nature of the sub-treatment. In the incentive arms, slope and flat incentives are not allowed to pool, but the amount of money (high or low) is considered to be a dosage. In the ambassador arms, we do not allow pooling between random selection of ambassadors, trusted ambassador, and information hub. Within information hubs, however, we consider that the “trusted information hub” is an increased dosage of information hub, so these may pool with one another.

To summarize, interventions “information hubs,” “slope,” “flat,” and “SMS” are found in two intensities. The marginal specification (2.4) therefore looks like

²⁷<https://doi.org/10.1257/rct.1434-4.0>

$$\begin{aligned}
y_{dsvt} = & \alpha_0 + \alpha_{SMS}SMS_s + \alpha_{H,SMS}High\ SMS_s \\
& + \alpha_{Slope}Slope_s + \alpha_{H,Slope}High\ Slope_s + \alpha_{Flat}Flat_s + \alpha_{H,Flat}High\ Flat_s \\
& + \alpha_RRandom_v + \alpha_HInfo\ Hub\ (All)_v + \alpha_TTrust_v + \alpha_{TH}Trusted\ Info\ Hub_v \\
& + \alpha'_X X_{sv} + v_{dt} + \epsilon_{dsvt},
\end{aligned}$$

where we have explicitly listed the variables in “single arm” treatment profiles. X_{sv} is a vector of the remaining 64 marginal effects variables in “multiple arm” treatment profiles, and v_{dt} is a set of district-time dummies. Here SMS refers to “any SMS.”

Our Puffered LASSO model selection estimation follows the recommended implementation in Rohe (2014), which uses a sequential backward elimination version of LASSO (variables with p -values above some threshold are progressively deselected) on the Puffer $_N$ transformed variables (the N refers to a further right multiplication of Puffer transformed variables with a diagonal matrix that aids in correcting for the heteroskedasticity induced by the Puffer transformation). We select penalties λ for both regressions (number of immunizations and immunizations per dollar) to minimize a Type I error, which is particularly important to avoid in the case of policy implementation.²⁸ This makes sense because it is extremely problematic to have a government introduce a large policy based on a false positive. This reasoning is elaborated in Appendix D.

This gives \widehat{S}_α , an estimate of the true support set S_α of the marginal effects specification (2.4). We then generate a use of unique pooled policy set \widehat{S}_{TVA} (following the procedure we outline in Algorithm 2 in Online Appendix B). Next, we run the pooled specification (2.2) to obtain post-LASSO estimates $\widehat{\eta}_{\widehat{S}_{TVA}}$ of the pooled policies as well as $\widehat{\eta}_{\widehat{S}_{TVA}, \widehat{R}^*}^{hyb}$, the winner’s curse adjusted estimate of the best policy.

5.1.2. *Results.* The results are presented in Figure 5. Panel A presents the post-LASSO estimates where the outcome variable is the number of measles vaccines per month in the village. Panel B presents the post-LASSO estimates where the outcome variable is the number of measles vaccines per dollar spent. In each, a relatively small subset of policies

²⁸Rohe (2014), a supplementary note to Jia and Rohe (2015)), deduces a algorithmic equivalence between a backwards elimination procedure based on using Type I error thresholds and LASSO on Puffer $_N$ transformed variables. In particular, the variables deselected by a LASSO penalty level λ are exactly those variables with classical OLS p -value above the Type I error threshold. We take $\lambda = 0.48$ and $\lambda = 0.0014$ for the number of immunizations and immunizations per dollar outcomes, respectively. Both of these choices map to the same Type I error value ($p = 5 \times 10^{-13}$) used in the backwards elimination implementation of LASSO, a threshold selected to essentially eliminate false positives. Appendix D elaborates on this choice. The algorithmic equivalence thus permits a common interpretation of the two models on the two outcomes as admitting the same Type I error. Its other advantage is that we can implement Rohe’s recommendation of *sequentially* deselecting variables until this threshold rather than eliminating them in a single step, because, as Rohe notes, this better handles variable correlation.

is selected as part of \widehat{S}_{TVA} out of the universe of 75 granular policies (16% of the possible options in Panel A and 35% in Panel B).

In Figure 5, Panel A, two of the four selected pooled policies are estimated to do significantly better than control: information hubs seeding with sloped incentives (of both low and high intensities) and SMS reminders (of both 33% and 66% saturation) are estimated to increase the number of immunizations by 55% relative to control ($p = 0.001$), while trusted seeds with high-sloped incentives and SMS reminders (of both saturation levels) are estimated to increase immunizations by 44% relative to control ($p = 0.009$).

These two effective policies increase the number of immunizations, relative to the status quo, at the cost of a greater cost for each immunization (compared to standard policy). These policies induce 36.0 immunizations per village per month per \$1,000 allocation (as compared with 43.6 immunizations per village per month in control). The reason is that that the gains from having incentives in terms of immunization rates is smaller than the increase in costs (e.g., the incentives must be paid to all the infra-marginal parents).

Two things are worth noting to qualify those results, however. First, in [Chernozhukov et al. \(2018\)](#), we show that in the places where the full package treatment is predicted to be the most effective (which tends to be the places with low immunization), the number of immunizations per dollar spent is not statistically different in treatment and control villages. Second, immunization is so cost-effective that this relatively small increase in the cost of immunization may still mean a much more cost-effective use of funds than the next best use of dollars on policies to fight childhood disease ([Ozawa et al., 2012](#)).

Nevertheless, a government may be interested in the most cost-effective policy, if they have a given budget for immunization. We turn to policy cost effectiveness in Figure 5, Panel B. The most cost-effective policy (and the only policy that reduces per immunization cost) compared to control is the combination of information hub seeding (trusted or not) with SMS reminders (at both 33% or 66% saturation) and no incentives, which leads to a 9.1% increase in vaccinations per dollar ($p = 0.000$).

5.2. Estimating the Impact of the Best policy. To estimate the impact of the best policy, we first select the best policy from \widehat{S}_{TVA} based on the post-LASSO estimate. Then, we attenuate it using the hybrid estimator with $\alpha = 0.05$ and $\beta = \frac{\alpha}{10} = 0.005$, which this is the value used by [Andrews et al. \(2021\)](#) in their simulations. The hybrid confidence interval has the following interpretation: conditional on policy effects falling within a 99.5% simultaneous confidence interval, the hybrid confidence interval around the best policy has at least 95% coverage. It also has at least 95% coverage unconditionally.²⁹

²⁹Per Proposition 6 of [Andrews et al. \(2021\)](#), it has unconditional coverage between $1 - \alpha = 95\%$ and $\frac{1-\alpha}{1-\beta} = 95.58\%$.

Table 1 presents the results. In column 1, the outcome variable is the number of measles vaccines given every month in a given village. We find that for the best policy in the sample (information hub seeds with sloped incentives at any level and SMS reminders at any saturation) the hybrid estimated best policy effect relative to control is 3.26 with a 95% hybrid confidence interval of [0.032,6.25]. This is lower than the original post-LASSO estimated effect of 4.02. The attenuation is owing to a second best policy (trusted seeds with high sloped incentives with SMS reminders at any saturation), chasing the best policy estimate somewhat closely.³⁰ Nevertheless, even accounting for winner’s curse through the attenuated estimates and the adjusted confidence intervals, the hybrid estimates still reject the null. Thus, the conclusion is that accounting for winner’s curse, this policy increases immunizations by 44% relative to control.

While policymakers may chose this policy if they are willing to bear a higher cost to increase immunization, there may be settings where cost effectiveness is an important consideration. In column 2, the outcome variable is the number of vaccinations per dollar. Accounting for winner’s curse through hybrid estimation, for the best policy of information hubs (all variants) and SMS reminders (any saturation level), the hybrid estimated best policy effect relative to control is 0.004 with a 95% hybrid confidence interval of [0.003,0.004]. Notably, this appears almost unchanged from the naive post-LASSO. This is because no other pooled policy with positive effect is “chasing” the best policy in the sample; the second-best policy is the control (status quo), which is sufficiently separated from the best policy so as to have an insignificant adjustment for winner’s curse. Thus, adjusting for winner’s curse, this policy increases the immunizations per dollar by 9.1% relative to control.

One concern with these estimates may be that they are sensitive to the implied LASSO penalty λ chosen. To check the robustness of our results, we consider alternative values of λ . However, we also need a criteria for evaluating results under various λ since a marginal effects support will never be robust for the whole range of λ . Online Appendix D spells out this criteria, which amounts to formulating a set of “admissible” λ . In a nutshell, the criterion is to avoid including in \hat{S}_α first and second best policies that are very likely to be false positives. Including a false positive as the first best a serious error in the context of policy advising, but it also matters for the second best, since including these in the support may overly attenuate the best policy estimate for winner’s curse. In our case, we find that for both immunizations and immunizations per dollar, the winner’s curse adjusted estimates are robust for their respective sets of admissible λ . To exemplify this robustness, we can take the union of confidence intervals within their admissible sets. This is [0.32,6.25] for immunizations and [0.001,0.006] for immunizations per dollar. Neither is much wider than the single confidence interval for the choice of λ we highlight.

³⁰The increased attenuation from a more closely competing second-best policy emerges from the formulas for conditional inference given in Section 3 of [Andrews et al. \(2021\)](#).

Though admissible λ are on a different scale for the two outcomes, there is a sense in which the admissible set is larger for immunizations per dollar. This suggests a different kind of robustness concern which is more about the relative fragility of the TVA estimator for each of the outcomes. We can speak to this fragility using a bootstrapping analysis described in detail in our Online Appendix F. Intuitively, it captures stability of best policy estimation in terms of observation leverage, where conclusions driven by outliers will fare worse. In this analysis, the best policy for cost-effectiveness holds for 96% with highly concentrated estimates around the main one in actual data. Meanwhile, the best policy for immunization holds for 77% of bootstrapped samples with estimates more widely dispersed. This speaks to the relative stability of the best policy for cost effectiveness over that for immunizations.

6. CONCLUSION

Despite immunization being one of the most effective and cost-effective methods to prevent disease, disability, and mortality, millions of children each year go unvaccinated. The COVID-19 epidemic has made the situation worse: vaccine coverage has dipped to levels not seen since the 1990s (Bill and Melinda Gates, 2020). Swift policy action is critical to ensure that this dip is temporary and children who missed immunizations during the pandemic get covered soon.

In rural India, there was a priori reason to believe that nudges may work. After all, many children get their first vaccines but caregivers rarely follow through. This is consistent with the vast majority of caregivers reporting that vaccines are helpful. Yet, it was a priori unclear as to which nudge, let alone which policy bundle out of the 75 candidates, would be effective.

Respecting this genuine uncertainty was critical. If we had simply done parallel treatments of incentives, reminders, and ambassadors, we might have found no effects. Our key finding is that combined interventions work better than each in isolation. Though there is temptation of paring down the number of treatments a priori for power, there is a danger in not doing this in a data-driven way. The suggestion of avoiding all interactions in this setting (made in Muralidharan et al. (2019)) would have led to the conclusion that nothing is effective.

From the point of view of public health policy the interaction effects identified by TVA tells us that it is valuable to add network-based insights (information hubs), which are not in a typical policymaker's toolkit, to catalyze the effects of conventional instruments. From a basic research perspective, it also suggests that the information hubs, i.e. the person best placed in a village to circulate information, may be more effective when they have something concrete to talk about, such as incentives or something to explain such as SMSs. Such questions merit future research.

The method suggested here is applicable to many domains where policymakers have several arms with multiple potential doses, do not have the time or capacity to adaptively experiment, and have genuine uncertainty about which policy bundles should be effective. Rather than guessing, we suggest that policymakers consider a data-driven approach of treatment variant aggregation. The proposed method relies on strong assumptions that rule out some of the cases where model-selection leads to invalid inferences. Provided these assumptions are palatable, our findings show that TVA prunes and pools effectively and, this pays dividends when the policymaker wishes to adjust for the winner’s curse without falling into the trap of over-conservatism. The algorithm can be easily pre-specified and does not require the researcher to take a stance on the possible effects of myriad interactions which are likely difficult to predict in advance.

REFERENCES

- AKBARPOUR, M., S. MALLADI, AND A. SABERI (2017): “Diffusion, Seeding, and the Value of Network Information,” Available at SSRN: <https://ssrn.com/abstract=3062830>. 24
- ALATAS, V., A. G. CHANDRASEKHAR, M. MOBIUS, B. A. OLKEN, AND C. PALADINES (2019): “When Celebrities Speak: A Nationwide Twitter Experiment Promoting Vaccination In Indonesia,” Tech. rep., National Bureau of Economic Research. 7
- ANDREWS, I., T. KITAGAWA, AND A. MCCLOSKEY (2021): “Inference on winners,” Tech. rep. 1, 1, 15, 2.5, 2.5, 3.3.1, 22, 3.3.2, 5.2, 29, 30, ??, D, E.2.2, E.2, F, G.2
- ANGRIST, J. D. AND J.-S. PISCHKE (2009): *Mostly harmless econometrics: An empiricist’s companion*, Princeton university press. 38
- ARAL, S. AND D. WALKER (2012): “Creating social contagion through viral product design: A randomized trial of peer influence in networks,” *Management Science*. 4.2.3
- BANERJEE, A., E. BREZA, A. G. CHANDRASEKHAR, AND B. GOLUB (2018): “When Less is More: Experimental Evidence on Information Delivery During India’s Demonetization,” Tech. rep., National Bureau of Economic Research. 4.2.2
- BANERJEE, A., A. CHANDRASEKHAR, E. DUFLO, S. DALPATH, J. FLORETTA, M. JACKSON, H. KANNAN, A. SCHRIMPF, AND M. SHRESTHA (2021): “Evaluating the impact of interventions to improve full immunisation rates in Haryana, India,” . 4.1
- BANERJEE, A., A. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): “Diffusion of Microfinance,” *Science*, 341, DOI: 10.1126/science.1236498, July 26 2013. 24, 4.2.3
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2019): “Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials,” *The Review of Economic Studies*. 7, 4.2.3, 4.5
- BANERJEE, A. V., E. DUFLO, R. GLENNERSTER, AND D. KOTHARI (2010): “Improving immunisation coverage in rural India: clustered randomised controlled evaluation of

- immunisation campaigns with and without incentives,” *Bmj*, 340, c2220. 4, 4.2.1
- BASSANI, D. G., P. ARORA, K. WAZNY, M. F. GAFFEY, L. LENTERS, AND Z. A. BHUTTA (2013): “Financial incentives and coverage of child health interventions: a systematic review and meta-analysis,” *BMC Public Health*, 13, S30. 4, 4.2.1
- BEAMAN, L., A. BENYISHAY, J. MAGRUDER, AND A. M. MOBARAK (2018): “Can network theory-based targeting increase technology adoption?” Tech. rep., National Bureau of Economic Research. 4.2.3
- BERK, R., L. BROWN, A. BUJA, K. ZHANG, AND L. ZHAO (2013): “Valid post-selection inference,” *The Annals of Statistics*, 802–837. 2.4
- BICKEL, P. J., Y. RITOV, AND A. B. TSYBAKOV (2009): “Simultaneous analysis of Lasso and Dantzig selector,” *The Annals of statistics*, 37, 1705–1732. 2.2
- BILL AND F. MELINDA GATES (2020): “COVID-19: A Global Perspective: 2020 Goalkeepers Report,” . 6
- BLOCH, FRANCIS AND, M. O. AND P. TEBALDI (2016): “Centrality Measures in Networks,” <http://ssrn.com/abstract=2749124>. 24
- CHANDIR, S., D. A. SIDDIQI, S. ABDULLAH, E. DUFLO, A. J. KHAN, AND R. GLENNERSTER (2022): “Small mobile conditional cash transfers (mCCTs) of different amounts, schedules and design to improve routine childhood immunization coverage and timeliness of children aged 0-23 months in Pakistan: An open label multi-arm randomized controlled trial,” *EClinicalMedicine*, 50. 4.2.1, 4.5
- CHERNOZHUKOV, V., M. DEMIRER, E. DUFLO, AND I. FERNANDEZ-VAL (2018): “Generic machine learning inference on heterogenous treatment effects in randomized experiments,” Tech. rep., National Bureau of Economic Research. 5.1.2
- CHERNOZHUKOV, V., C. HANSEN, AND M. SPINDLER (2015): “Post-selection and post-regularization inference in linear models with many controls and instruments,” *American Economic Review*, 105, 486–90. E.3
- DLHS (2013): “District Level Household and Facility Survey-4,” . 4.1
- DOMEK, G. J., I. L. CONTRERAS-ROLDAN, S. T. O’LEARY, S. BULL, A. FURNISS, A. KEMPE, AND E. J. ASTURIAS (2016): “SMS text message reminders to improve infant vaccination coverage in Guatemala: A pilot randomized controlled trial,” *Vaccine*, 34, 2437–2443. 5
- GIBSON, D. G., B. OCHIENG, E. W. KAGUCIA, J. WERE, K. HAYFORD, L. H. MOULTON, O. S. LEVINE, F. ODHIAMBO, K. L. O’BRIEN, AND D. R. FEIKIN (2017): “Mobile phone-delivered reminders and incentives to improve childhood immunisation coverage and timeliness in Kenya (M-SIMU): a cluster randomised controlled trial,” *The Lancet Global Health*, 5, e428–e438. 4, 4.2.1

- HADAD, V., D. A. HIRSHBERG, R. ZHAN, S. WAGER, AND S. ATHEY (2021): “Confidence intervals for policy evaluation in adaptive experiments,” *Proceedings of the National Academy of Sciences*, 118, e2014602118. 1
- HAMMER, S. M., K. E. SQUIRES, M. D. HUGHES, J. M. GRIMES, L. M. DEMETER, J. S. CURRIER, J. J. ERON JR, J. E. FEINBERG, H. H. BALFOUR JR, L. R. DEYTON, ET AL. (1997): “A controlled trial of two nucleoside analogues plus indinavir in persons with human immunodeficiency virus infection and CD4 cell counts of 200 per cubic millimeter or less,” *New England Journal of Medicine*, 337, 725–733. 1
- HE, X. AND Q.-M. SHAO (2000): “On parameters of increasing dimensions,” *Journal of multivariate analysis*, 73, 120–135. 2.4, 2.4, A
- HINZ, O., B. SKIERA, C. BARROT, AND J. U. BECKER (2011): “Seeding Strategies for Viral Marketing: An Empirical Comparison,” *Journal of Marketing*, 75:6, 55–71. 24
- IYENGAR, R., C. V. DEN BULTE, AND T. W. VALENTE (2010): “Opinion Leadership and Social Contagion in New Product Diffusion,” *Marketing Science*, 30:2, 195–212. 24
- JACKSON, M. O. (2008): “Average Distance, Diameter, and Clustering in Social Networks with Homophily,” in *the Proceedings of the Workshop in Internet and Network Economics (WINE 2008), Lecture Notes in Computer Science, also: arXiv:0810.2603v1*, ed. by C. Papadimitriou and S. Zhang, Springer-Verlag, Berlin Heidelberg. 24
- (2017): “A Typology of Social Capital and Associated Network Measures,” *SSRN* <http://ssrn.com/abstract=3073496>. 24
- JACKSON, M. O. AND L. YARIV (2011): “Diffusion, strategic interaction, and social structure,” *Handbook of Social Economics, San Diego: North Holland, edited by Benhabib, J. and Bisin, A. and Jackson, M.O.* 24
- JAVANMARD, A. AND A. MONTANARI (2013): “Model selection for high-dimensional regression under the generalized irrepresentability condition,” in *Advances in neural information processing systems*, 3012–3020. 1, 2.4, A
- (2014): “Confidence intervals and hypothesis testing for high-dimensional regression,” *The Journal of Machine Learning Research*, 15, 2869–2909. 3.3.3, E.2.3, E.4
- (2018): “Debiasing the lasso: Optimal sample size for gaussian designs,” *The Annals of Statistics*, 46, 2593–2622. 3.3.3, E.2.3
- JIA, J. AND K. ROHE (2015): “Preconditioning the Lasso for sign consistency,” *Electronic Journal of Statistics*, 9, 1150–1172. 1, 2.2, 2.2, 10, 28, A, 37
- JOHRI, M., M. C. PÉREZ, C. ARSENAULT, J. K. SHARMA, N. P. PAI, S. PAHWA, AND M.-P. SYLVESTRE (2015): “Strategies to increase the demand for childhood vaccination in low-and middle-income countries: a systematic review and meta-analysis,” *Bulletin of the World Health Organization*, 93, 339–346. 4

- KARING, A. (2018): “Social Signaling and Childhood Immunization: A Field Experiment in Sierra Leone,” *University of California, Berkeley*. 6
- KASY, M. AND A. SAUTMANN (2021): “Adaptive treatment assignment in experiments for policy choice,” *Econometrica*, 89, 113–132. 1
- KATONA, Z., P. P. ZUBCSEK, AND M. SARVARY (2011): “Network Effects and Personal Influences: The Diffusion of an Online Social Network,” *Journal of Marketing Research*, 48:3, 425–443. 24
- KATZ, E. AND P. LAZARSFELD (1955): *Personal influence: The part played by people in the flow of mass communication*, Free Press, Glencoe, IL. 4.2.3
- KEMPE, D., J. KLEINBERG, AND E. TARDOS (2003): “Maximizing the Spread of Influence through a Social Network,” *Proc. 9th Intl. Conf. on Knowledge Discovery and Data Mining*, 137–146. 24
- KRACKHARDT, D. (1996): “Structural Leverage in Marketing,” in *Networks in Marketing*, ed. by D. Iacobucci, Sage, Thousand Oaks, 50–59. 24
- LEE, J. D., D. L. SUN, Y. SUN, AND J. E. TAYLOR (2016): “Exact post-selection inference, with application to the lasso,” *The Annals of Statistics*, 44, 907–927. 15, 2.4
- LEEB, H. AND B. M. PÖTSCHER (2005): “Model selection and inference: Facts and fiction,” *Econometric Theory*, 21, 21–59. 3, 2.4
- (2008): “Sparse estimators and the oracle property, or the return of Hodges’ estimator,” *Journal of Econometrics*, 142, 201–211. 2.4
- MCCLOSKEY, A. (2020): “Asymptotically uniform tests after consistent model selection in the linear regression model,” *Journal of Business & Economic Statistics*, 38, 810–825. 2.4
- MCKENZIE, D. (2019): “Be careful with inference from 2-by-2 experiments and their cross-cutting designs,” [Online; accessed October 19, 2019]. 1
- MEINSHAUSEN, N. AND B. YU (2009): “Lasso-type recovery of sparse representations for high-dimensional data,” *The annals of statistics*, 37, 246–270. 37, 44
- MEKONNEN, Z. A., K. A. GELAYE, M. C. WERE, K. D. GASHU, AND B. C. TILAHUN (2019): “Effect of mobile text message reminders on routine childhood vaccination: a systematic review and meta-analysis,” *Systematic reviews*, 8, 1–14. 4.2.2
- MURALIDHARAN, K., M. ROMERO, AND K. WUTHRICH (2019): “Factorial designs, model selection, and (incorrect) inference in randomized experiments,” . 1, 1, 1, 6
- NFHS (2016): “National Family Health Survey-4 State Fact Sheet for Haryana,” . 4.1
- NIE, L. AND V. ROČKOVÁ (2022): “Bayesian bootstrap spike-and-slab lasso,” *Journal of the American Statistical Association*, 1–35. 3.3.4, E.2.4, E.5
- OYO-ITA, A., C. S. WIYSONGE, C. ORINGANJE, C. E. NWACHUKWU, O. ODUWOLE, AND M. M. MEREMIKWU (2016): “Interventions for improving coverage of childhood

- immunisation in low-and middle-income countries,” *Cochrane Database of Systematic Reviews*. 4
- OZAWA, S., A. MIRELMAN, M. L. STACK, D. G. WALKER, AND O. S. LEVINE (2012): “Cost-effectiveness and economic benefits of vaccines in low-and middle-income countries: a systematic review,” *Vaccine*, 31, 96–108. 5.1.2
- REGAN, A. K., L. BLOOMFIELD, I. PETERS, AND P. V. EFFLER (2017): “Randomized controlled trial of text message reminders for increasing influenza vaccination,” *The Annals of Family Medicine*, 15, 507–514. 5
- ROČKOVÁ, V. AND E. I. GEORGE (2018): “The spike-and-slab lasso,” *Journal of the American Statistical Association*, 113, 431–444. E.2.4
- ROGERS, E. (1995): *Diffusion of Innovations*, Free Press. 24
- ROHE, K. (2014): “A note relating ridge regression and ols p-values to preconditioned sparse penalized regression,” *arXiv preprint arXiv:1411.7405*. 1, 5.1.1, 28, D
- TAYLOR, J. AND R. J. TIBSHIRANI (2015): “Statistical learning and selective inference,” *Proceedings of the National Academy of Sciences*, 112, 7629–7634. D
- TIBSHIRANI, R. J., J. TAYLOR, R. LOCKHART, AND R. TIBSHIRANI (2016): “Exact post-selection inference for sequential regression procedures,” *Journal of the American Statistical Association*, 111, 600–620. 2.4
- UDDIN, M. J., M. SHAMSUZZAMAN, L. HORNG, A. LABRIQUE, L. VASUDEVAN, K. ZELLER, M. CHOWDHURY, C. P. LARSON, D. BISHAI, AND N. ALAM (2016): “Use of mobile phones for improving vaccination coverage among children living in rural hard-to-reach areas and urban streets of Bangladesh,” *Vaccine*, 34, 276–283. 5
- UNICEF AND WHO (2019): “Progress and challenges with achieving universal immunization coverage: 2018 estimates of immunization coverage WHO,” . 1
- USER1551 ([HTTPS://MATH.STACKEXCHANGE.COM/USERS/1551/USER1551](https://math.stackexchange.com/users/1551/user1551)) (2017): “Find the uniform lower bound of the smallest eigenvalue of a certain matrix,” Mathematics Stack Exchange, <https://math.stackexchange.com/q/2438656> (version: 2017-09-21). 34
- VAN DE GEER, S. (2019): “On the asymptotic variance of the debiased Lasso,” *Electronic Journal of Statistics*, 13, 2970–3008. 3.3.3, E.2.3
- WAKADHA, H., S. CHANDIR, E. V. WERE, A. RUBIN, D. OBOR, O. S. LEVINE, D. G. GIBSON, F. ODHIAMBO, K. F. LASERSON, AND D. R. FEIKIN (2013): “The feasibility of using mobile-phone based SMS reminders and conditional cash transfers to improve timely immunization in rural Kenya,” *Vaccine*, 31, 987–993. 4, 5
- WHO (2019): “Situation Analysis of Immunization Expenditure,” . 1
- ZHAN, R., Z. REN, S. ATHEY, AND Z. ZHOU (2021): “Policy learning with adaptively collected data,” *arXiv preprint arXiv:2105.02344*. 1

- ZHANG, C.-H. AND S. S. ZHANG (2014): “Confidence intervals for low dimensional parameters in high dimensional linear models,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76, 217–242. 3.3.3
- ZHAO, P. AND B. YU (2006): “On model selection consistency of Lasso,” *Journal of Machine learning research*, 7, 2541–2563. 1, 2.2, 31, C.1, 36

FIGURES




FIGURE 1. Hasse diagram for $M = 2, R = 4$ for the treatment profile where both arms are active. A line upwards from treatment combinations $[r_1, r_2]$ to $[r'_1, r'_2]$ means that $[r_1, r_2] \leq [r'_1, r'_2]$ and $[r_1, r_2] \neq [r'_1, r'_2]$ in the intensity ordering.

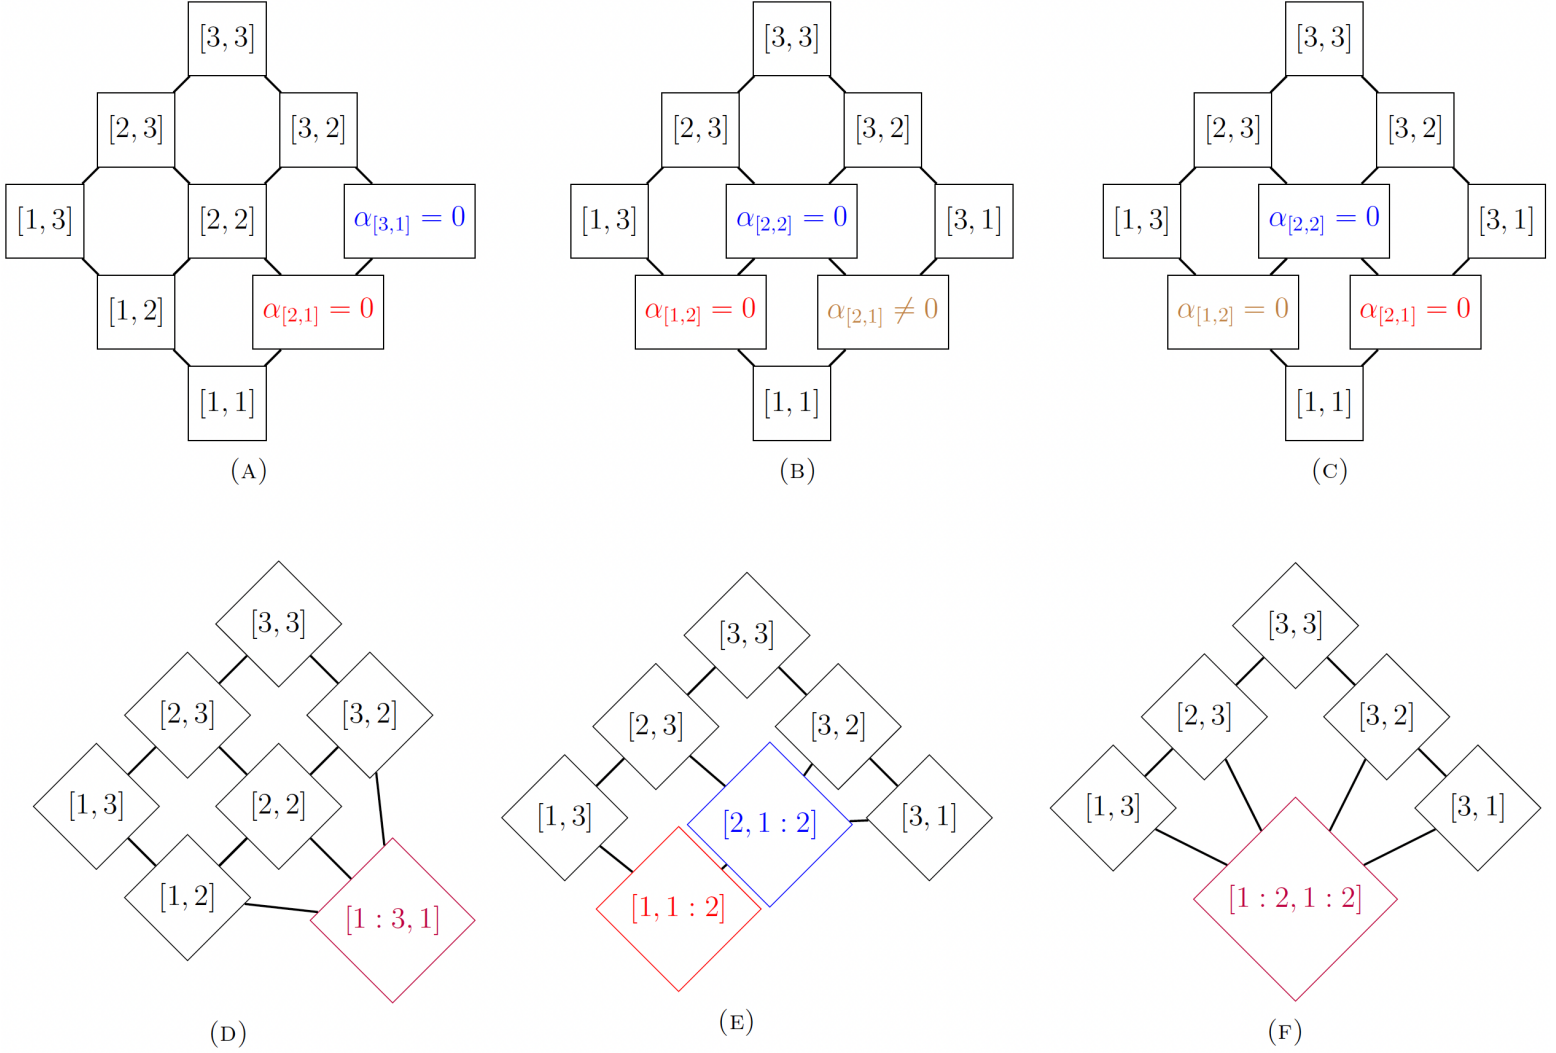


FIGURE 2. This figure shows how zeros in the marginals induce admissible "concatenations" in the Hasse diagram of policies. Panels A, B, C show examples of zeros in the marginal space and panels D, E, F show the corresponding policy-concatenations in the η -space. On panel A, $\alpha_{[2,1]} = \alpha_{[3,1]} = 0$. Since $\alpha_{[2,1]} = \beta_{[2,1]} - \beta_{[1,1]}$ and $\alpha_{[3,1]} = \beta_{[3,1]} - \beta_{[2,1]}$, $[1, 1], [2, 1], [3, 1]$ are concatenated on *panel D*. On panel B, $\alpha_{[2,1]} \neq 0 = \beta_{[2,1]} - \beta_{[1,1]}$, $\alpha_{[1,2]} = 0 = \beta_{[1,2]} - \beta_{[1,1]}$, and $\alpha_{[2,2]} = 0 = (\beta_{[2,2]} - \beta_{[2,1]}) - (\beta_{[1,2]} - \beta_{[1,1]})$. The implied poolings are shown on panel E where policies $\{[1, 1], [1, 2]\}$ and $\{[2, 1], [2, 2]\}$ are concatenated. In panel C note that the only change relative to panel B is that here $\alpha_{[2,1]} = 0$, but the implied concatenation on panel F yields $\beta_{[1,1]} = \beta_{[1,2]} = \beta_{[2,1]} = \beta_{[2,2]}$.

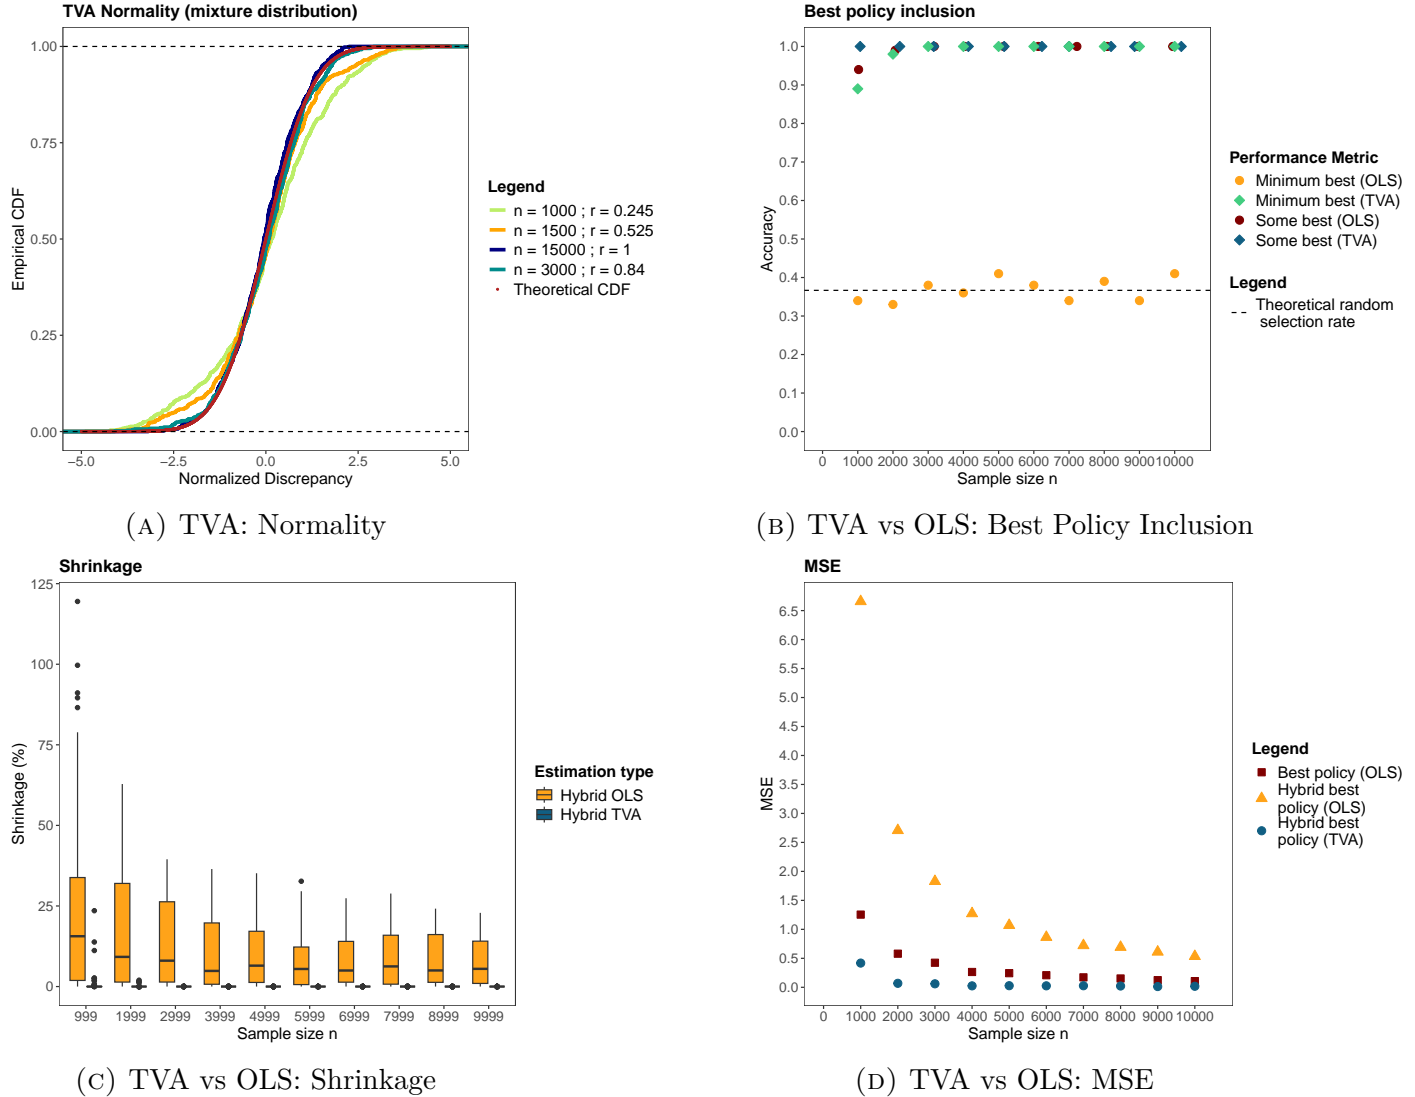


FIGURE 3. A plot comparing the performance of the TVA estimator to applying OLS on the unique policy specification (2.1) for a range of measures. On panel A, we first expose the normality of TVA estimates (r is correct support selection rate). Panel B then uses the best policy inclusion measures defined in subsection 3.2 and points are slightly jittered for better readability. For OLS, this measure is set to 1 whenever the highest treatment effect policy is part of the true best pool (some best) or equal to the minimum dosage best policy (min best). Panel C compares the amount of shrinkage imposed by the winner’s curse adjustment as percentage of the initial coefficient. Panel D compares the MSE of the best policy estimation, between the TVA and the OLS estimator of the unique policy specification (2.1) before and after adjusting for the winner’s curse. In all panels, there are 20 simulations per configuration and 5 configurations per n .




FIGURE 4. Effects on the number of measles vaccinations relative to control (7.32) by reminders, incentives, and seeding policies, restricted to the villages where the ambassador intervention was administered. The specification is weighted by village population, controls for district-time fixed effects, and clusters standard errors at the sub-center level.

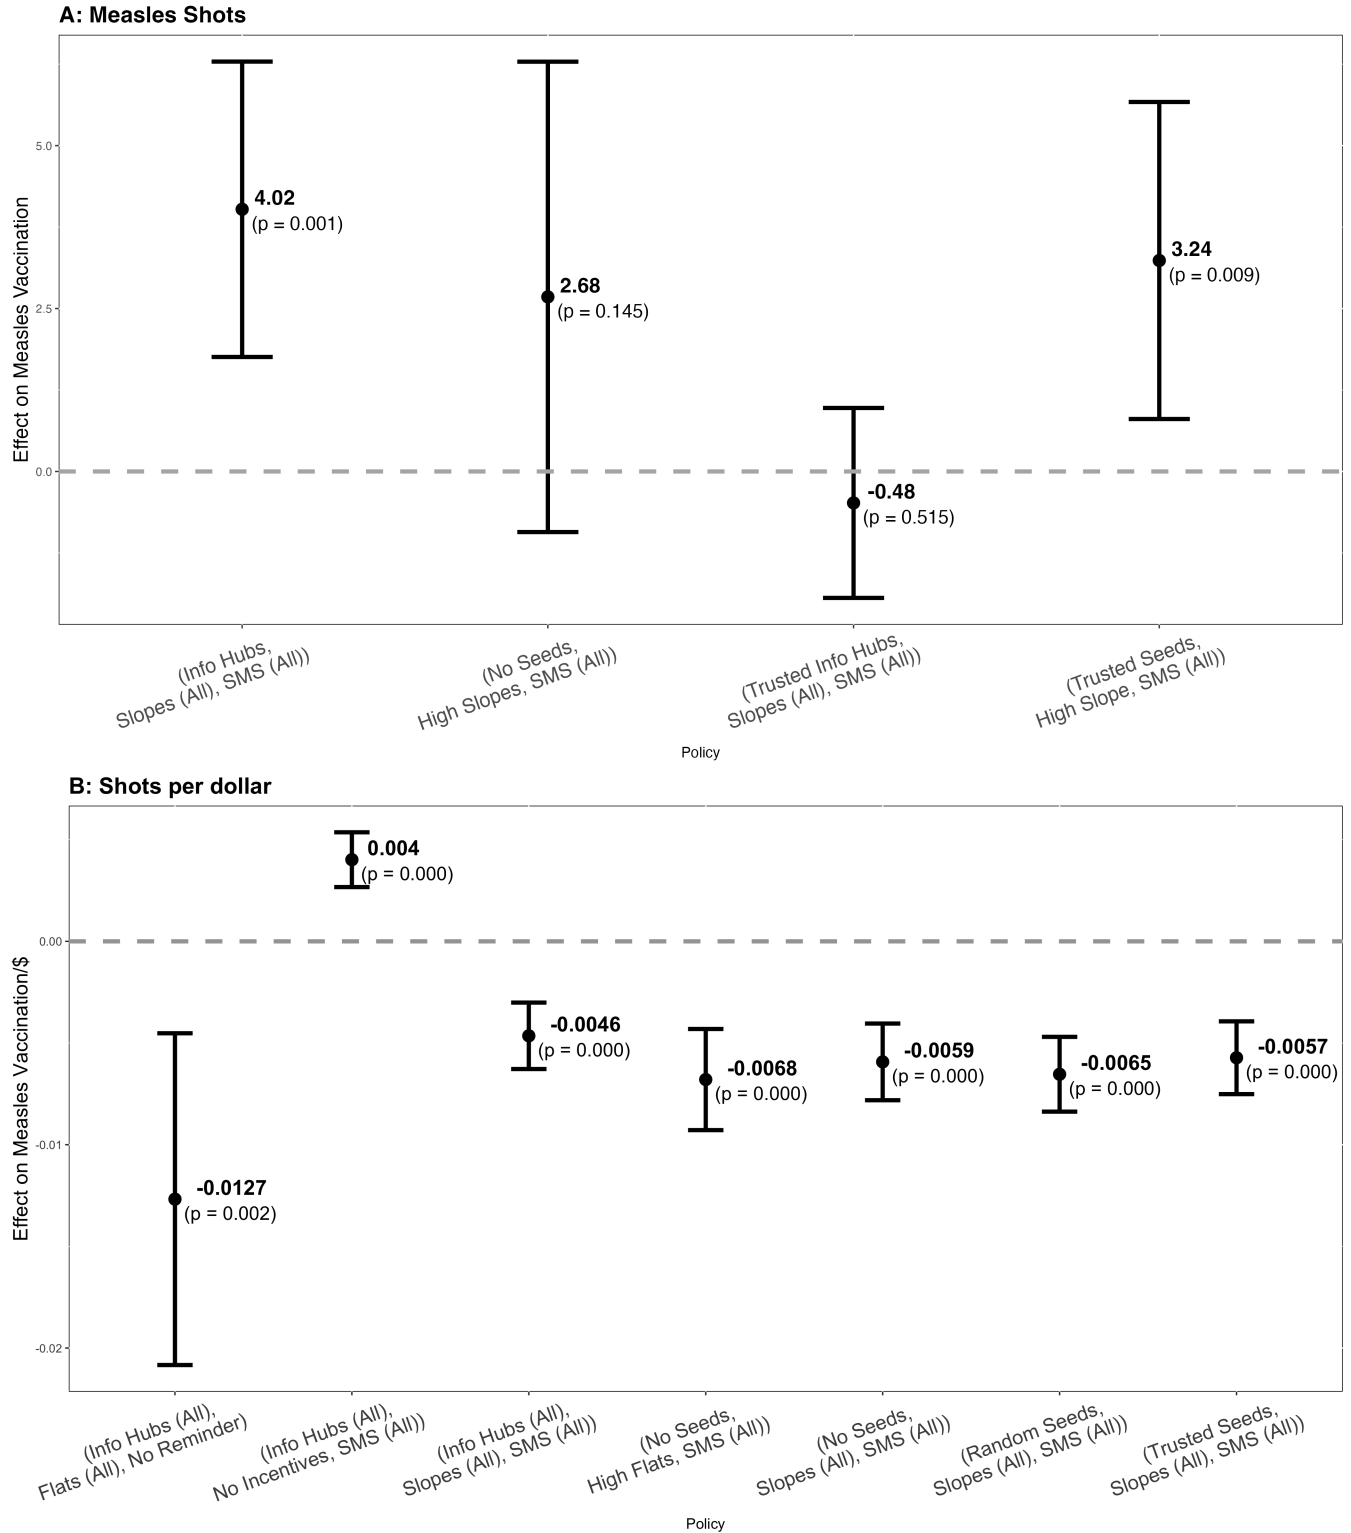


FIGURE 5. TVA estimates of combinations of reminders, incentives, and seeding policies on the number of measles vaccines (Panel A) and the number of measles vaccines per \$ (Panel B) relative to control (7.32 and 0.0436 respectively). The specifications are weighted by village population and include controls for district-time fixed effects. Standard errors are clustered at the sub-center level. 95% confidence intervals displayed.

TABLES

TABLE 1. Best Policies

	(1)	(2)
	# Measles Shots	# Measles Shots per \$1
WC Adjusted Treatment Effect	3.26	0.004
Confidence Interval (95%)	[0.32, 6.25]	[0.003, 0.005]
Control Mean	7.32	0.0436
Observations	204	814
Optimal Policy	(Information Hubs, SMS, Slope)	(Information Hubs POOLED, SMS)

Notes: Estimation using [Andrews et al. \(2021\)](#); hybrid estimation with $\alpha = 0.05, \beta = 0.005$. The specifications are weighted by village population and account for district-time fixed effects as well as variance clustered at the sub-center level.

APPENDIX A. PROOFS

Proof of Proposition 2.1. According to Theorem 1 of Jia and Rohe (2015), if $\min_{j \in S_\alpha} |\alpha_j| \geq 2\lambda_n$, then $\tilde{\alpha} =_s \alpha$ with probability greater than³¹

$$f(n) := 1 - 2K \exp\left(-\frac{n\lambda_n^2 \xi_{\min}^2}{2\sigma^2}\right),$$

where $\xi_{\min} = \xi_{\min}(\frac{X}{\sqrt{n}})$ is the minimum singular value of the \sqrt{n} -normalized design matrix. By Assumption 3, there is a uniform lower bound $c > 0$ on the magnitude of the non-zero $\{\alpha\}$. Since by Assumption 5, $\lambda_n \rightarrow 0$, for sufficiently high n $\min_{j \in S_\alpha} |\alpha_j| \geq 2\lambda_n$. Theorem 1 applies and $\text{sign}(\hat{\alpha}) = \text{sign}(\alpha)$ with probability greater than $f(n)$.

It will be convenient to re-express $f(n)$ as follows:

$$f(n) = 1 - 2 \exp\left(\log(K) - \frac{n\lambda_n^2 \xi_{\min}^2}{2\sigma^2}\right).$$

And applying Lemma A.1 and Assumption 2, for sufficiently high n :

$$f(n) \geq 1 - 2 \exp\left(\log(K) - \frac{n\lambda_n^2}{2\sigma^2 K^2}\right) \geq 1 - 2 \exp\left(\gamma \log(n) - \frac{n^{1-2\gamma} \lambda_n^2}{2\sigma^2}\right).$$

By Assumption 5, $\lambda_n^2 n^{1-2\gamma} = \omega(n^{1-2(\nu+\gamma)})$. Since $0 < \nu < \frac{1}{2} \implies 1 - 2\gamma > 1 - (2(\nu + \gamma)) > 0$, and logarithm growth is dominated by polynomial growth, it follows that $-\left(\gamma \log(n) - \frac{n^{1-2\gamma} \lambda_n^2}{2\sigma^2}\right)$ grows at polynomial rate $\omega(n^{1-2(\nu+\gamma)})$ and therefore $\lim_{n \rightarrow \infty} f(n) \geq 1$ at exponential rate. Since also $f(n) \leq 1$, it follows that $f(n) \rightarrow 1$ at exponential rates. ■

LEMMA A.1. *For the marginal effects design matrix X , for $R \geq 3$, wpa1 the lowest singular value of \sqrt{n} normalized design matrix, i.e., $\xi_{\min}(\frac{X}{\sqrt{n}})$, has the value $\xi_{\min}(\frac{X}{\sqrt{n}}) = \left(4R \sin^2\left(\frac{R-\frac{3}{2}\pi}{R-\frac{1}{2}\pi}\right)\right)^{-\frac{M}{2}}$. Thus, with probability approaching 1, $\xi_{\min}(\frac{X}{\sqrt{n}}) > (\frac{1}{K})$.³²*

Proof of Lemma A.1. It will be useful to index the design matrix X by R and M , i.e., $X = \mathbf{X}_{R,M}$. Let $\mathbf{C}_{R,M} = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbf{X}'_{R,M} \mathbf{X}_{R,M}$. Then $\lim_{n \rightarrow \infty} \xi_{\min}^2(\frac{\mathbf{X}_{R,M}}{\sqrt{n}}) = \lambda_{\min}(\mathbf{C}_{R,M})$, i.e., the lowest eigenvalue of $\mathbf{C}_{R,M}$. We will characterize this eigenvalue.

The combinatorics of the limiting frequencies of “1”s in marginal effects variables imply that $\mathbf{C}_{R,M}$ is a block diagonal matrix with structure $\mathbf{C}_{R,M} = K^{-1} \text{diag}(\mathbf{B}_{R,M}, \mathbf{B}_{R,M-1}, \dots, \mathbf{B}_{R,1})$ where $\mathbf{B}_{R,M-1}$ implies this block is found in $\mathbf{C}_{R,M-1}$ (pertaining to an RCT with one less

³¹The “ $=_s$ ” notation stands for equality in sign following Definition 1 in Zhao and Yu (2006).

³²This is a conservative bound; the optimal uniform lower bound is $(\frac{1}{K}(\frac{1}{4M}))^{\frac{1}{2}}$.

cross-treatment arm), etc. More than one block of $\mathbf{B}_{R,M-1}$, $\mathbf{B}_{R,M-2}$, ... $\mathbf{B}_{R,1}$ is found in $\mathbf{C}_{R,M}$, but only $\mathbf{B}_{R,M}$ determines the minimum eigenvalue.

The combinatorics of variable assignments also implies that

- (1) $\mathbf{B}_{R,M}$ is an $(R-1)^M \times (R-1)^M$ matrix with recursive structure $\mathbf{B}_{R,M} = \mathbf{B}_{R,1} \otimes \mathbf{B}_{R,M-1}$, where \otimes is the Kronecker product.³³
- (2) $\mathbf{B}_{R,1}$ is an $(R-1) \times (R-1)$ matrix with recursive structure

$$\mathbf{B}_{R,1} = \begin{bmatrix} R-1 & R-2 & \cdots & 1 \\ R-2 & & & \\ \vdots & & \mathbf{B}_{R-1,1} & \\ 1 & & & \end{bmatrix} \text{ and } \mathbf{B}_{2,1} = [1].$$

SUBLEMMA 1. $\lambda_{\min}(\mathbf{B}_{R,1}) = \left(4 \sin^2 \left(\frac{R-\frac{3}{2}\pi}{R-\frac{1}{2}}\right)\right)^{-1}$

Proof. The key insight³⁴ is that $\mathbf{B}_{R,1}^{-1}$ is the $(R-1) \times (R-1)$ tridiagonal matrix:

$$\mathbf{B}_{R,1}^{-1} = \begin{bmatrix} 1 & -1 & & & \\ -1 & 2 & -1 & & \\ & -1 & \ddots & \ddots & \\ & & \ddots & \ddots & -1 \\ & & & -1 & 2 \end{bmatrix}$$

which has known eigenvalues $\mu_j = 4 \sin^2 \left(\frac{j-\frac{1}{2}\pi}{R-\frac{1}{2}}\right)$ for $j = 1, 2, \dots, R-1$. Thus, given that the inverse of a matrix's eigenvalues are the inverse matrix's eigenvalues, $\lambda_{\min}(\mathbf{B}_{R,1}) = \left(4 \sin^2 \left(\frac{R-\frac{3}{2}\pi}{R-\frac{1}{2}}\right)\right)^{-1}$. ■

(Resuming the proof of Lemma A.1) Per the multiplicative property of the eigenvalues of a Kronecker product, together with the fact that all matrices in question are positive definite, it immediately follows that $\lambda_{\min}(\mathbf{B}_{R,M}) = \lambda_{\min}(\mathbf{B}_{R,1})\lambda_{\min}(\mathbf{B}_{R,M-1})$, which in turn implies $\lambda_{\min}(\mathbf{B}_{R,M}) = \lambda_{\min}(\mathbf{B}_{R,1})^M$. Since by Sublemma 1, $\lambda_{\min}(\mathbf{B}_{R,1}) < 1$, $\mathbf{B}_{R,M}$ is the block determining the rate with the smallest eigenvalue, and therefore, given that the eigenvalues of a block diagonal matrix are the eigenvalues of the blocks:

$$\lambda_{\min}(\mathbf{C}_{R,M}) = \frac{1}{K} \lambda_{\min}(\mathbf{B}_{R,M}) = \frac{1}{K} (\lambda_{\min}(\mathbf{B}_{R,1}))^M = \left(4R \sin^2 \left(\frac{R-\frac{3}{2}\pi}{R-\frac{1}{2}}\right)\right)^{-M}$$

where the last equality uses Sublemma 1. The Lemma follows. ■

³³Thanks to Nargiz Kalantarova for noticing this Kronecker product and its consequent implication for $\lambda_{\min}(\mathbf{B}_{R,M})$.

³⁴The argument is provided on Mathematics Stackexchange (user1551 (<https://math.stackexchange.com/users/1551/user1551>), 2017).

Proof of Proposition 2.2. The proof is found in Javanmard and Montanari (2013), proof of Theorem 2.7, with minor modifications, which we reproduce for completeness. Label the events $\mathcal{E} := \{\widehat{S}_{TV A} = S_{TV A}\}$ (that the treatment variants were aggregated correctly). Define the pseudo-true value $\eta_S^0 := \operatorname{argmin}_\eta \mathbb{E} [\|y - Z_S \eta\|_2^2]$, noting that $\eta_{S_{TV A}}^0$ satisfies this for $S = S_{TV A}$. Finally, let $\mathcal{F} := \left\{ \left\| \widehat{\eta}_{\widehat{S}_{TV A}} - \eta_{\widehat{S}_{TV A}}^0 \right\|_\infty > \epsilon \right\}$, so it is the event that the estimator exceeds the pseudo-true value on the estimated support by ϵ .

Then, we can write

$$\mathbb{P}(\mathcal{F}) = \mathbb{P}(\mathcal{F} \cap \mathcal{E}) + \mathbb{P}(\mathcal{F} \cap \mathcal{E}^c) \leq \mathbb{P}(\mathcal{F} \cap \mathcal{E}) + \mathbb{P}(\mathcal{E}^c).$$

By the proof of Proposition 2.1, we have

$$\mathbb{P}(\mathcal{E}^c) \leq 2K \exp\left(-\frac{n(\lambda/K)^2}{2\sigma^2}\right) = 2 \exp\left(\gamma \log n - \frac{n^{1-2\gamma}\lambda}{2\sigma^2}\right)$$

Turning to $\mathbb{P}(\mathcal{F} \cap \mathcal{E})$, on the event \mathcal{E} ,

$$\widehat{\eta}_{\widehat{S}_{TV A}} - \eta_{\widehat{S}_{TV A}}^0 = \left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}}\right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon,$$

since $\widehat{S}_{TV A} = S_{TV A}$. Therefore, for every $j \in \{1, \dots, K\}$, $\widehat{\eta}_{S_{TV A}, j} - \eta_{S_{TV A}, j}^0$ is normally distributed with variance order bounded above by $\frac{\sigma^2}{n \cdot C_{\min}}$ where $C_{\min} := \sigma_{\min}\left(n^{-1} Z'_{S_{TV A}} Z_{S_{TV A}}\right)$ is the minimum singular value of the design matrix. But $C_{\min} \geq \frac{1}{\sqrt{K}}$ by definition since each unit is assigned to a disjoint pooled policy, and each policy pools one or more variants. So

$$\mathbb{P}(\mathcal{F} \cap \mathcal{E}) = \mathbb{P}\left(\left\| \widehat{\eta}_{\widehat{S}_{TV A}} - \eta_{\widehat{S}_{TV A}}^0 \right\|_\infty > \epsilon \cap \mathcal{E}\right) \leq \mathbb{P}\left(\sup_j \left| \widehat{\eta}_{S_{TV A}, j} - \eta_{S_{TV A}, j}^0 \right| > \epsilon\right) \leq 2e^{-\frac{n \cdot \epsilon^2 \cdot K^{-1/2}}{2\sigma^2}},$$

using a Gaussian tail bound and a union bound for uniform control over $j \in \{1, \dots, |S_{TV A}|\}$.

Putting the pieces together we have

$$\mathbb{P}(\mathcal{F}) \leq 2 \exp\left(-\frac{n \cdot \epsilon^2 \cdot K^{-1/2}}{2\sigma^2}\right) + 2 \exp\left(\gamma \log n - \frac{n^{1-2\gamma}\lambda}{2\sigma^2}\right)$$

This establishes that $\mathbb{P}(\mathcal{F}) \rightarrow 0$ for every ϵ , i.e., the consistency of the estimator to the pseudo-true values. With probability $1 - 2 \exp\left(\gamma \log n - \frac{n^{1-2\gamma}\lambda}{2\sigma^2}\right) \rightarrow 1$, the event \mathcal{E} is active and the pseudo-true value will be the true value. In this case consider $\epsilon = \sqrt{\frac{\log n}{n^{1-\gamma/2}}}$. Then

$$\mathbb{P}(\mathcal{F} \cap \mathcal{E}) = 2e^{-\frac{n \epsilon^2 C_{\min}}{2\sigma^2}} \leq 2e^{-\frac{\log n}{n^{1-\gamma/2}} \cdot \frac{n K^{-1/2}}{2\sigma^2}} \leq 2e^{-\frac{\log n}{2\sigma^2}} = \frac{2}{n^{1/2\sigma^2}} \rightarrow 0,$$

i.e., $\left\| \widehat{\eta}_{\widehat{S}_{TV A}} - \eta_{\widehat{S}_{TV A}}^0 \right\|_\infty < \sqrt{\frac{\log n}{n^{1-\gamma/2}}}$ with high probability, completing the proof. ■

Proof of Proposition 2.3. As in the proof of Proposition 2.2, let $\mathcal{E} := \{\widehat{S}_{TV A} = S_{TV A}\}$. Let $c \in \mathbb{R}^{|\widehat{S}_{TV A}|}$ whose length depends on the estimated support (and is therefore random),

though we suppress this dependence. We can decompose³⁵ $\sqrt{n}c' \left(\hat{\eta}_{\hat{S}_{TV A}} - \eta_{\hat{S}_{TV A}}^0 \right)$, a real-valued random variable, as

$$\begin{aligned} \sqrt{n}c' \left(\hat{\eta}_{\hat{S}_{TV A}} - \eta_{\hat{S}_{TV A}}^0 \right) &= \mathbf{1} \{ \mathcal{E} \} \cdot \sqrt{n}c' \left(Z'_{S_{TV A}} Z_{S_{TV A}} \right)^{-1} Z'_{S_{TV A}} \epsilon + \mathbf{1} \{ \mathcal{E}^c \} \cdot \sqrt{n}c' \left(Z'_{\hat{S}_{TV A}} Z_{\hat{S}_{TV A}} \right)^{-1} Z'_{\hat{S}_{TV A}} \epsilon \\ &= \sqrt{n}c' \left(Z'_{S_{TV A}} Z_{S_{TV A}} \right)^{-1} Z'_{S_{TV A}} \epsilon - \mathbf{1} \{ \mathcal{E}^c \} \cdot \sqrt{n}c' \left(Z'_{S_{TV A}} Z_{S_{TV A}} \right)^{-1} Z'_{S_{TV A}} \epsilon \\ &\quad + \mathbf{1} \{ \mathcal{E}^c \} \cdot \sqrt{n}c' \left(Z'_{\hat{S}_{TV A}} Z_{\hat{S}_{TV A}} \right)^{-1} Z'_{\hat{S}_{TV A}} \epsilon. \end{aligned}$$

It is helpful to note that c is *any* vector of the same length as the size of the estimated support. So conditional on the event of the estimated support being the true support, for instance, it has the true length. Conditional on any realization of an alternative length support, one considers a conformal vector corresponding to the estimated support size.

The proof strategy is to see that the CLT ensures the asymptotic normality of the first term in this sum in the usual way, following [He and Shao \(2000\)](#) since we have a growing number of parameters, while the remaining terms will asymptotically vanish in probability.

Let us take the first term, so we are looking at the case *on* $S_{TV A}$. Here we show asymptotic normality. In order to show this, we need to show that every linear combination of the vector of regression coefficients (of which there are a growing number) is asymptotically normally distributed when properly normalized. We show the sufficient conditions for Corollary 2.1 in [He and Shao \(2000\)](#). First, by Assumption 2, note that $K^2 \log(K) = O(n^{2\gamma} \log(n)) = o(n)$, using that $2\gamma < 1$ and that $\log(n)$ grows more slowly than any polynomial in n . K thereby satisfies the hypothesized growth rate condition in Corollary 2.1 for smooth scores, which is our case since we study linear regression. Second, we check (D1)-(D3) after which the corollary applies. (D1) follows since $n^{-1} Z'_{S_{TV A}} Z_{S_{TV A}} = I$ by definition since every observation is assigned to one unique treatment dummy. (D2) follows since in the case of linear regression, the score and its derivative are bounded. As noted in [He and Shao \(2000\)](#), a sufficient condition for (D3) is that the regression vector, here $Z_{S_{TV A}, i}$ for observation i is such that $E |c' Z_{S_{TV A}, i}|^4$ for any c in the unit sphere of length $|S_{TV A}|$. But this is mechanically true since $\|Z_{S_{TV A}, i}\| = 1$ since it is a saturated vector of treatment bundle assignments and therefore has a single 1 corresponding to the entry for the treatment assigned.

Therefore, Corollary 2.1 of [He and Shao \(2000\)](#) applies. This means that

$$\sqrt{n}c' \left(\hat{\eta}_{\hat{S}_{TV A}} - \eta_{\hat{S}_{TV A}}^0 \right) / \sigma(c) \rightsquigarrow \mathcal{N}(0, 1).$$

for any $c \in \mathbb{R}^{|S_{TV A}|}$, where $\sigma^2(c) := \sigma^2 \|c\|^2$.

The second term can be handled by showing for any $c \in \mathbb{R}^{|S_{TV A}|}$, we have $\mathbf{1} \{ \mathcal{E}^c \} \cdot \sqrt{n}c' \left(Z'_{S_{TV A}} Z_{S_{TV A}} \right)^{-1} Z'_{S_{TV A}} \epsilon = o_p(1)$. We already showed that $\sqrt{n}c' \left(Z'_{S_{TV A}} Z_{S_{TV A}} \right)^{-1} Z'_{S_{TV A}} \epsilon$ is asymptotically normal and so $O_p(1)$. Since $\mathbf{1} \{ \mathcal{E}^c \}$ is $o_p(1)$, the whole term is $o_p(1)$.

³⁵We thank Adel Javanmard for a helpful discussion.

Regarding the third one: $\mathbf{1}\{\mathcal{E}^c\} \cdot \sqrt{n}c' \left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon$. So here we study $\mathbf{1}\{\mathcal{E}^c\} \cdot \sqrt{n}c' \left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon$ where $c \in \mathbb{R}^{|\widehat{S}_{TV A}|}$ and show that it is $o_p(1)$. The point is that $\left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon$, which potentially inherits omitted variable bias by including incorrect regressors, is nevertheless uniformly bounded in n and K . In fact, we can bound the $\|\cdot\|_\infty$ norm which implies bounding the size of the inner product with any aforementioned c . In detail, first note that $\left\| \left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon \right\|_\infty \leq \left\| Z'_{\widehat{S}_{TV A}} \epsilon \right\|_\infty$ because $\left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1}$ is a positive definite block diagonal matrix with every entry < 1 . Secondly $\left\| Z'_{\widehat{S}_{TV A}} \epsilon \right\|_\infty < Kn\epsilon$ since $Z_{\widehat{S}_{TV A}}$ is a binary matrix. Since ϵ is $O_p(1)$, $\left\| \left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon \right\|_\infty$ is uniformly $O_p(Kn)$ over all misspecifications $\widehat{S}_{TV A}$.

Thus, $\mathbf{1}\{\mathcal{E}^c\} \cdot \sqrt{n}c' \left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon$ is uniformly bounded in probability by $P(\mathcal{E}^c)O_p(Kn^{\frac{3}{2}})$. But since $P(\mathcal{E}^c) \leq 2K \exp\left(-\frac{n(\lambda/K)^2}{2\sigma^2}\right) = 2 \exp\left(\log(K) - \frac{n(\lambda/K)^2}{2\sigma^2}\right)$, recalling the proof of Proposition 2.2,

$$O_p(Kn^{\frac{3}{2}})P(\mathcal{E}^c) = O_p\left(e^{\left(\frac{3}{2}\log(n)+2\log(K)-\frac{n(\lambda/K)^2}{2\sigma^2}\right)}\right) = O_p\left(e^{\left(\frac{3}{2}\log(n)+2\gamma\log(n)-\frac{n^{1-2\gamma}\lambda^2}{2\sigma^2}\right)}\right) = o_p(1).$$

So $\mathbf{1}\{\mathcal{E}^c\} \cdot \sqrt{n}c' \left(Z'_{\widehat{S}_{TV A}} Z_{\widehat{S}_{TV A}} \right)^{-1} Z'_{\widehat{S}_{TV A}} \epsilon$ vanishes in probability and altogether,

$$\sqrt{n}c' \left(\widehat{\eta}_{\widehat{S}_{TV A}} - \eta_{\widehat{S}_{TV A}}^0 \right) / \sigma^2(c) \rightsquigarrow \mathcal{N}(0, 1) + o_p(1) + o_p(1) \rightsquigarrow \mathcal{N}(0, 1),$$

which completes the proof. ■

ONLINE APPENDIX : NOT FOR PUBLICATION

Selecting the Most Effective Nudge: Evidence from a Large-Scale
Experiment on Immunization

Abhijit Banerjee, Arun G. Chandrasekhar, Suresh Dalpath, Esther Duflo, John Floretta,
Matthew O. Jackson, Harini Kannan, Francine Loza, Anirudh Sankar, Anna Schrimpf, and
Maheshwor Shrestha

APPENDIX B. POOLING PROCEDURE AND PROPERTIES

Here we show how to construct a set of pooled policies \hat{S}_{TVA} from an estimated support of marginals \hat{S}_α . In case the marginals are correctly estimated $\hat{S}_\alpha = S_\alpha$, we will show that the implied pooling $\hat{S}_{TVA} = S_{TVA}$ has the properties of Λ -admissibility and maximality.

Given \hat{S}_α let $[\hat{S}_\alpha]$ denote its partition into sets of support vectors with the same treatment profile. Each $S \in [\hat{S}_\alpha]$ is thus a set of treatment combinations $\{k_1, \dots, k_n\}$. For each $k_i \in \{k_1, \dots, k_n\}$, define a set comprised of policies sharing k_i 's profile that weakly dominate it:

$$(B.1) \quad A_{k_i} = \{k \in \mathcal{K} | P(k) = P(k_i) \text{ and } k \geq k_i\}.$$

Now consider a simple operator that selects, for every set B , either the set or its complement. Let us write this as B^a , where $a \in \{1, c\}$. When $a = 1$, the operator selects B , while when $a = c$ the operator selects B^c . Following this, the pooled policies \hat{S}_{TVA} is defined as the collection of sets:

$$(B.2) \quad \mathcal{A} = A_{k_1}^{a_1} \cap \dots \cap A_{k_n}^{a_n}$$

which are nonempty, and where furthermore at least one $a_i = 1$. Algorithm 2 describes this construction of \hat{S}_{TVA} procedurally.

We will now lend intuition for this construction, which will clarify its properties. For this discussion, assume the marginal support is correctly estimated, i.e. $\hat{S}_\alpha = S_\alpha$. Considering the relationship between marginal effects α and policy effects β given by (2.3), A_k is the set of policies whose effects are partly determined or ‘‘influenced’’ by α_k .

It is useful to depict this on the Hasse with a simple example. Take $M = 2, R = 4$, and the treatment profile where both arms are on. The Hasse diagram of the unique policies within this treatment profile is shown on Figure C.1.

Consider for example the set $A_{[2,1]}$, stemming from $\alpha_{[2,1]}$, depicted as all those policies within the blue contour on this figure. These are all the policies that are influenced by $\alpha_{[2,1]}$. Motivated by this visual depiction of influence, let us say that in general A_k is the ‘‘sphere of influence’’ of α_k . A policy’s effect is entirely determined by the various ‘‘spheres’’ acting on it. We can restrict attention to those spheres stemming from $\alpha_k \neq 0$ because the other spheres are inactive (spheres of ‘‘non-influence’’). Each set \mathcal{A} constructed per (B.2), describes a set of policies subject to a unique set of active spheres. When $a_i = 1$, it means that the sphere A_{k_i} stemming from α_{k_i} is active; when $a_i = c$ the sphere is inactive. Thus in Figure C.1, $\mathcal{A}_1 := A_{[1,2]} \cap A_{[2,1]}$ are those policies influenced by both $\alpha_{[1,2]}$ and $\alpha_{[2,1]}$. $\mathcal{A}_2 := A_{[1,2]} \cap A_{[2,1]}^c$ are those policies influenced by *only* $\alpha_{[1,2]}$ (and not by $\alpha_{[2,1]}$); vice-versa for $\mathcal{A}_3 := A_{[1,2]}^c \cap A_{[2,1]}$. For any policy $k \in \mathcal{A}$, the $\alpha \leftrightarrow \beta$ parameter relationship (2.3) immediately implies $\beta_k = \sum_{i|a_i=1} \alpha_{k_i}$.

It is immediate by construction that this pooling is both admissible and maximal. Exactly those policies influenced by the same marginals are pooled, and nothing else is. Of course for a given realization of marginals α , one could have policies influenced k, k' influenced by different marginals (spheres of influence) which happen to have the same policy effects $\beta_k = \beta'_k$. But we don't want to pool such policies because jittering the nonzero α (keeping the support S_α fixed) would mean $\beta_k \neq \beta'_k$ almost surely. Since we want to consider pooling using only S_α irrespective of the exact value of α , exactly those policies influenced by the same spheres should be pooled.

OBSERVATION 1. S_{TVA} is the the coarsest pooling uniformly over α conditional on S_α




FIGURE C.1. Hasse Diagram with $A_{2,1}$ and $A_{1,2}$ with complements and intersections.

Algorithm 2: Pooling Procedure

input : Estimated support \widehat{S}_α from the marginal specification (2.4)
output: Estimated pooled policies \widehat{S}_{TVA} for pooled specification (2.2)
Partition \widehat{S}_α into $[\widehat{S}_\alpha]$ per the treatment profile mapping $P(\cdot)$;
Initialize $\widehat{S}_{TVA} \leftarrow \emptyset$;
for $S \in [\widehat{S}_\alpha]$ **do**
 discover $S = \{k_1, \dots, k_n\}$;
 generate $\{A_{k_1}, \dots, A_{k_n}\}$;
 for each $(a_1, \dots, a_n) | a_i \in \{1, c\}$ **do**
 generate $\mathcal{A} = A_{k_1}^{a_1} \cap \dots \cap A_{k_n}^{a_n}$;
 if $\mathcal{A} \neq \emptyset$ and $\mathcal{A} \neq A_{k_1}^c \cap \dots \cap A_{k_n}^c$ **then**
 | $\widehat{S}_{TVA} \leftarrow \widehat{S}_{TVA} \cup \mathcal{A}$;
 end
 end
end
return \widehat{S}_{TVA}

APPENDIX C. IRREPRESENTABILITY FAILURES

C.1. TVA without Puffering, (2.4), can fail irrerepresentability. Proof by construction. Irrepresentability says that no covariate is explained too much by others, in terms of the total magnitudes of regression coefficients. Formally, for any regression of any covariate on set of other covariates, the L_1 norm of the coefficients (excluding intercept) is not to exceed 1. Irrepresentability fails, therefore, when there is any covariate k^* and $K' \subset K, k^* \notin K'$ such that in an OLS regression:

$$X_{k^*}^n = \hat{\gamma}_0 + \sum_{k \in K'} \hat{\gamma}_k X_k^n.$$

$\sum_{k \in K'} |\hat{\gamma}_k| > 1$ for all n .³⁶ In that case X_{k^*} is considered “representable” by $X_{K'}$. The irrerepresentability criterion (Zhao and Yu, 2006) says that this is a problem when X_{k^*} is not in the support X_{S_α} of (2.4) but all $X_{K'}$ are, as with some nonvanishing probability LASSO will select X_{k^*} along with $X_{K'}$ on all points of the regularization path that include $X_{K'}$.³⁷ We will show that when $M \geq 2, R \geq 3$ (i.e. there are at least 2 arms and 2 nonzero dosages per arm, a case which our particular immunization cross randomized experiment nests), there can be an irrerepresentability failure for some S_α .

We provide a candidate k^* and K' for $M = 2, R = 3$, that can also be embedded in any higher number of arms or intensities.

PROPOSITION C.1. *Let $k^* = [2, 2]$ and $K' = \{[1, 2], [2, 1], [1, 1]\}$. The population level regression $X_{k^*} \sim X_{K'}$ is³⁸*

$$(C.1) \quad X_{[2,2]} = \frac{1}{2} (X_{[2,1]} + X_{[1,2]}) - \frac{1}{4} X_{[1,1]} + \epsilon \quad E[\epsilon X_{K'}] = 0$$

³⁶See Zhao and Yu (2006), p. 2544. We disprove weak irrerepresentability, which also disproves strong irrerepresentability. In the notation of that paper, the weak irrerepresentability condition is $\left| C_{21}^n (C_{11}^n)^{-1} \text{sign} \left(\beta_{(1)}^n \right) \right| < 1$, where the inequality holds element-wise. This condition says that the “signed sum” (per $\text{sign} \left(\beta_{(1)}^n \right)$) of the coefficients of each OLS regression of a nonsupport variable on all support variables, must be < 1 . Because the sign is unknown ex-ante, weak irrerepresentability must hold over all signs, which is equivalent to demanding that L_1 norm of each aforementioned regression must be < 1 . Consequently, to be representable, we just need to demonstrate one “representable” regression where this doesn’t hold.

³⁷This insight is clarified in Meinshausen and Yu (2009), who prove that LASSO solutions are L_2 -norm consistent regardless of whether design matrices satisfy irrerepresentability. Violations of irrerepresentability reduce the pressure on the LASSO objective function to yield sparse solutions, and with some nonvanishing probability LASSO will recover both support and nonsupport variables on all the points along regularization path which include the support (this is geometrically depicted in Figure 2 of Jia and Rohe (2015)). That is, with some nonvanishing probability LASSO will select only strict supersets of the correct support on the regularization path.

³⁸We borrow the “population level” terminology from Angrist and Pischke (2009). Formally, the proposition treats each X_k as a scalar random variable with appropriate unconditional and conditional distributions $P(X_{[2,2]} = 1) = \frac{n}{K}$, $P(X_{[1,1]} = 1 | X_{[2,1]} = 1) = 1$ etc. Letting $X_{K'}$ be the cartesian product $(X_{[2,1]}, X_{[1,2]}, X_{[1,1]})$, the proposition says that $\gamma := (\frac{1}{2}, \frac{1}{2}, \frac{1}{4})$ resolves $\arg \min_{b \in \mathbb{R}^3} E (X_{[2,2]} - X_{K'}' b)^2$. The necessary and sufficient first order condition is $E[(X_{[2,2]} - X_{K'}' \gamma) X_{K'}] = 0$

The L_1 norm of (C.1) is $\frac{1}{2} + \frac{1}{2} + \frac{1}{4} > 1$, so $X_{[2,2]}$ is representable by $\{X_{[2,1]}, X_{[1,2]}, X_{[1,1]}\}$. Thus if S_α includes $\{X_{[2,1]}, X_{[1,2]}, X_{[1,1]}\}$ but not $X_{[2,2]}$, then with nonvanishing probability LASSO will recover $X_{[2,2]}$ together with $\{X_{[2,1]}, X_{[1,2]}, X_{[1,1]}\}$ on all points of the regularization path that include $\{X_{[2,1]}, X_{[1,2]}, X_{[1,1]}\}$. Note that any higher number of arms or intensities has such two arm-two nonzero intensity subexamples, for example treatment profiles where only two arms ‘on’, and only considering the lowest two nonzero intensities. Thus this irrepresentability failure can be embedded in many such S_α for any $M \geq 2, R \geq 3$. For these S_α , with some nonvanishing probability LASSO will not retrieve S_α on the sample.

Proof of Proposition C.1. First observe that $E[X_{K'}X'_{K'}]$ is nonzero hence invertible, so the population regression is unique. To verify that (C.1) is the population regression, we just need to verify the first order condition $E[\epsilon X_{K'}] = 0$ for $\epsilon := X_{[2,2]} - \frac{1}{2}(X_{[2,1]} + X_{[1,2]}) + \frac{1}{4}X_{[1,1]}$. We will show $E[\epsilon|X_{K'}] = 0$ which is sufficient³⁹.

Note that $\epsilon X_{K'}$ is trivially 0 outside the treatment profile where both arms are ‘on’, since both ϵ and $X_{K'}$ are zero there. So we can narrow the analysis to just this one treatment profile.

Consider Table C.1 of assignments within the treatment profile where both arms are ‘on’.

TABLE C.1. Assignments within treatment profile where both arms are ‘on’

T_1	T_2	$X_{[1,1]}$	$X_{[1,2]}$	$X_{[2,1]}$	$X_{[2,2]}$	ϵ
1	1	1	0	0	0	0.25
2	1	1	0	1	0	-0.25
1	2	1	1	0	0	-0.25
2	2	1	1	1	1	0.25

Note that each row occurs with equal probability, since each row corresponds to a distinct unique policy that is assigned with equal propensity. It suffices to show that $E[\epsilon|X_k] = 0$ for each $X_k \in \{X_{[1,2]}, X_{[2,1]}, X_{[2,2]}\}$. Let us show that $E[\epsilon|X_{[2,1]}] = 0$, the other cases are demonstrated analogously. $X_{[2,1]} = 0$ in rows 1 and 3, in which case $\epsilon = 0.25$ or $\epsilon = -0.25$ with equal probability with expected value 0. Likewise $X_{[2,1]} = 1$ in rows 2 and 4, in which case $\epsilon = -0.25$ and $\epsilon = 0.25$ with equal probability with expected value 0. Thus $E[\epsilon|X_{[2,1]}] = 0$. Analogously it can be shown that $E[\epsilon|X_{[1,2]}] = E[\epsilon|X_{[1,1]}] = 0$. Having established the first order condition the proposition is proved. ■

C.2. TVA without Puffering, (2.4), can fail irrepresentability. Simulation results showing that these failures worsen in increasing R and M . Proposition C.1 finds an irrepresentability failure that applies to all $M \geq 2, R \geq 3$. Here we show via simulation that irrepresentability failures can worsen in increasing intensity and number of arms. Intuitively, increasing R and M increases correlations within the matrix X and makes these failures more

³⁹Using that $E[\epsilon X] = E[E[\epsilon X|X]] = E[XE[\epsilon|X]] = E[E[0|X]] = 0$

dramatic. Consider the marginal effect covariate where all M arms are “on” with highest intensity, i.e., X_{k^*} for $k^* = (R - 1, \dots, R - 1)$. Let $K' = K - \{k\}$, i.e. all the other covariates. We will show that X_{k^*} is “representable” by $X_{K'}$, and increasingly so in increasing R and M . Consider sample regressions $X_{k^*}^n \sim X_{K'}^n$

$$X_{k^*}^n = \hat{\gamma}_0 + \sum_{k \in K, k \neq k^*} \hat{\gamma}_k X_k^n.$$

Simulation results depict $\sum_{k \in K, k \neq k^*} |\hat{\gamma}_k|$.

In this simulation, we choose large $n = 10,000$ so that with high probability the same regression will exhibit representability on any other sample draw by Chebyshev’s inequality. We consider two kinds of regressions: an “unstandardized” regression where the raw marginal effects covariates are regressed, and a “standardized” regression where the marginal covariates are first standardized by the L_2 norm. The latter corresponds to a preprocessing step that LASSO packages typically apply before LASSO selection; we would like to know if there can be representability failures even in this case. Indeed, we see the L_1 norms are greater than 1 in both cases, and the higher dosage marginal effects indicators become increasingly representable (in terms of L_1 norm) by lower dosage marginal effects indicators in increasing R and M .

R	M	L_1 norm (standardized vars)	L_1 norm (unstandardized vars)
3	2	1.73	1.26
3	3	3.67	2.32
3	4	7.66	4.2
4	2	1.77	1.24
4	3	3.98	2.43
4	4	8.27	4.14
5	2	1.87	1.28
5	3	3.78	2.29
5	4	7.90	4.70

APPENDIX D. ROBUSTNESS

In this section we consider robustness of TVA conclusions on our data. As motivated in the body, we consider the plausible environment of Section 2 so that a data driven procedure can reveal a set of relevant pooled policies and estimate it together with an estimate of the single best policy.⁴⁰ The main issues are then the sensitivity of TVA results to (1) level of sparsity imposed and (2) the particular draw of the data. We discuss these below, and in doing so also elaborate on the “admissible” LASSO penalties λ for evaluating robustness as well as the choice of λ emphasized in the body. We intend for this to be helpful as a user’s guide for future practitioners.

Sparsity level

The LASSO penalty level λ , which determines the level of sparsity imposed by TVA, trades off Type I and Type II error. A higher λ implies less of the first at the price of the second. Our main recommendation is to err on the side of lower Type I error. For one, inclusion in \hat{S}_{TVA} of false positives might (misleadingly) over-attenuate the best policy for its winner’s curse if one of them emerges as the “second best”. Secondly, and more seriously, one of these false positives might itself be selected as the best policy. From a policy standpoint this is a particularly dangerous error in the context of government advice. So, adopting a conservative stance, the admissible λ will be higher than lower.

We determine the sufficiently high λ through the following exercise, which can be applied generally. Namely, we first consider raw sensitivity of best policy selection and winner’s curse adjusted estimates to a range of λ , as in Figure D.1 panel A.⁴¹ Within this range, we see that for immunizations/\$, the policy (Info Hubs (All), No Incentives, SMS (All)) is robustly selected across the range. On the other hand, for immunizations the policy (Info Hubs, Slopes (All), SMS (All)) is selected for $\lambda \geq 0.47$ while (No Seeds, High Slopes, Low SMS) is selected for $\lambda < 0.47$. Furthermore, for immunizations/\$ the winners curse adjusted estimates of (Info Hubs, Slopes (All), SMS (All)) reject 0 except at $\lambda \leq 0.00045$. On the other hand, for the outcome of immunizations the winners curse adjusted estimate of (Info Hubs, Slopes (All), SMS (All)) rejects 0 for the range $\lambda \geq 0.47$. For $\lambda < 0.47$ the other selected policy (No Seeds, High Slopes, Low SMS) always fails to reject 0.

Next, we interpret these preliminary findings in light on the aforementioned false positive risk. To do this, we plot the unadjusted post-LASSO estimates of the best policy together with the second best estimates, as in panel B of Figure D.1. This gives a sense of the the nominal estimates/confidence intervals and the ‘perpetrators’ of winners curse attenuations.

⁴⁰We emphasize again that exact sparsity of Assumption 2 can be practically relaxed to some of the approximate sparsity regimes explored in Section 3 subsection ??

⁴¹Note that for immunizations, the support coverage (of the 75 policies) ranges from 3% (right side of the diagram at the highest displayed λ) to 52% (left side of the diagram at the lowest displayed λ). For immunizations/\$ the support coverage ranges from 4% to 39%.

We evaluate these nominal values per the observation by Taylor and Tibshirani (2015) that a LASSO based procedure greedily favors false positives over true negatives, and so actual type I error rates and p -values in post-selection inference are larger than nominal ones. Thus nominally insignificant policies are even more likely to be true negatives, and nominally borderline significant policies are likely to be insignificant as well. Our conservative principle thus tells us dial up the λ to remove considerations of these. We demonstrate below.

Consider first the immunizations/\$ outcome. For $\lambda < 0.0008$ these second bests (policies in green and pink) have nominal OLS confidence intervals that fail to reject 0. Thus λ should be increased to the range $\lambda \geq 0.0008$. Since after $\lambda = 0.00160$ both this policy and the policy support is entirely deselected, the admissible λ is the interval $[0.0008, 0.00160)$. Per Figure D.1, panel A1, the winner’s curse estimates are robust in this admissible interval.

Proceeding to the immunizations outcome, for $\lambda \leq 0.19$ and $\lambda \in (0.415, 0.452]$, even nominal post-LASSO estimates of (No Seeds, High Slopes, Low SMS) fail to reject 0. For the range $0.19 < \lambda < 0.358$, the noisy and nominally insignificant policy (Info Hubs, High Slope, No Reminder) is responsible for the winner’s curse attenuations. We therefore dial up λ further. Doing so gives an extremely narrow range ($\lambda \in [0.358, 0.415]$) where (No Seeds, High Slopes, Low SMS) is significant, and that too barely ($p = 0.048$ at the displayed $\lambda = 0.039$). Keeping in mind the aforementioned observation by Taylor and Tibshirani (2015), this nominal significance is particularly unreliable at the threshold. Moreover, it is in stark contrast to the robust stability of the post-LASSO estimates of (Info Hubs, Slopes (All), SMS (All)) at $\lambda \geq 0.39$. Altogether this is strongly suggestive that (No Seeds, High Slopes, Low SMS) is a policy that we should disinclude as a false positive. This is achieved at $\lambda > 0.452$. Since after $\lambda \geq 0.53$ this and all policies are dropped, the admissible λ for this interval is $(0.452, 0.53)$. Per Figure D.1, panel A2, the winner’s curse estimates are robust in this admissible interval.

The two sets of admissible λ are for two entirely different scale of outcomes and cannot usually be directly compared. However, since our specific implementation of LASSO on Puffer_N is bijective with a p -value threshold in a backwards elimination (following Rohe (2014)), following this bijection the set of admissible λ for immunization/\$ is wider. So there is some sense in which TVA estimator is more fragile. Another approach that evaluates this fragility is a bootstrapping analysis, which is explained shortly. Also since we use this implementation with p value thresholds, for the result we highlight in the paper we use “bottleneck” $\lambda = 0.48$ for immunizations, namely $p = 5 \times 10^{-13}$. This maps to $\lambda = 0.0014$ for immunizations/\$. The corresponding Type I error level is thus constant.

Finally, we emphasize that checking a λ against a saturated regression of 75 raw coefficients for finely differentiated policies – a kind of “heatmap” test – is generally *not* a reliable



robustness check. In Figures F.1 (number of shots) and F.2 (shots/\$) of our Online Appendix F we show the fully unrolled saturated regression together with the pooling. But it is not clear how to interpret this with respect to pooling. The eye cannot sanity check the combinatorially large number of joint hypothesis tests for pooling that become relevant when more than one arm is on. Figure E.1, panel B from Section 3 considers a typical simulation where eye-balling fails to sanity check the right pooling choices.

Bootstrapping analysis



It is natural to ask about the fragility of TVA to the particular draw of the data; precisely this concern motivates, for example, our implementation of winner’s curse adjustments by Andrews et al. (2021), as well as simulations in Section 3 that directly speak to the variance of TVA. However, one might further wonder about just the observations in our dataset, with a concern akin to one about leverage of observations. An intuitive approach to address this is a bootstrapping analysis, where TVA is run on multiple bootstrapped samples. We can then speak to variation in both the set of supports selected as well the estimates of the pooled policies. Because this is a more exploratory analysis, its principal value lies in speaking to *relative* stability of conclusions between the two policies for the two outcomes.

⁴² Results are presented in Figures F.3 (immunizations) and F.4 (immunizations /\$) of our Online Appendix and demonstrate a robust stability for the cost-effective outcome, where the original support is selected in 96% of bootstrapped sample, against 77% for the immunizations outcome.

⁴²Note that we have slightly different goals here from the issue of bootstrapped standard errors.



(A) Best policy per LASSO penalty (after WC adjustment)



(B) Best and Second Best policies per LASSO penalty (before WC adjustment)

FIGURE D.1. Panel A shows the sensitivity of best-policy estimation per LASSO penalty λ for a sequential elimination version of LASSO on $\text{Puffer}_N(X, Y)$. Estimates are plotted after adjusting for the Winner's Curse. Panel B shows the sensitivity of both the best and second-best policy estimates. Here estimates are shown prior to winner's curse adjustment i.e. when best policy selection happens.

APPENDIX E. EXTENDED SIMULATIONS

This section serves the purpose of providing complementary evidence and extending the discussion on TVA’s performance (1) in the environment described in Section 2, (2) relative to alternative estimation strategies and (3) under relaxation sparsity and lower-bound assumptions.

E.1. Properties of TVA. Simulation performance of TVA attests to its main theoretical properties: support consistency, best policy estimation consistency, and normally distributed coefficient estimates.

Figure 3 in the main paper depicts these results. For TVA consistency, consider the blue and green performance points in panels B and D. Panel B shows that even for low n , TVA includes some of the best policies in \widehat{S}_α as well as rapidly (in n) and consistently including the *minimum* dosage best policy, that is of particular interest to the policymaker. Because TVA pools policies that comprise the best policy, redundancies to winner’s curse attenuations are minimized; Panel D shows that MSE of the best policy starts small and quickly falls to 0 (blue dots). Finally, besides best policy consistency, a more global concern is whether the TVA estimator is support consistent in the first place. This is explicitly verified in further simulations of our Online Appendix (e.g., Figure E.3). As elaborated in Section 3.3, this overall performance is in marked contrast to other estimators.

Besides being consistent, TVA estimates are also distributed asymptotically normally,, which permits reliable inference in the usual manner. This is demonstrated in Figure 3, Panel A which depicts a close match between the empirical CDF of standardized marginal policy estimates $\frac{\hat{\eta}_\kappa - \eta_\kappa}{se(\hat{\eta}_\kappa)}$ (for $\kappa \in \widehat{S}_{TVA}$, entered as a mixture distribution with equal weights) and the CDF of a standard normal distribution for a series of sample sizes. First, abstracting away from model misspecification complications, consider the case of a large sample size ($n = 15,000$) where the empirical support of TVA is always correct (blue empirical CDF). The almost perfect match to the theoretical CDF speaks to normally distributed estimates of all $M = 3$ components of the mixture distribution.⁴³ The cases where $n < 15,000$ relax insistence on correct model specification and demonstrate that this is not an artifact of a very large sample size. Here, $\eta_{\widehat{S}_{TVA,\kappa}}$ is the population pseudo-parameter in the (potentially) misspecified model. Even for rather moderate sample sizes ($n = 3000$), we see normally distributed estimates.

E.2. Performance relative to alternatives.

E.2.1. Direct OLS. Besides the specific issue of best policy estimation mentioned in Section 3.3.1, direct OLS has low power. This is depicted in Figure E.1, where simulated OLS

⁴³For this exercise we use a single randomly chosen configuration \mathcal{C} that determines the coefficients η_κ , $\kappa \in \{1, \dots, M\}$ against which to compare our estimates.

estimates of all the unique policies (2.1) contrast with the pruned and pooled estimation (2.2). As expected, the estimated effects of the pooled policies are less dispersed (panel A). In this visualization, we deliberately choose a configuration where the effects of different policies are similar, so that these histograms overlap. This makes the task of discovering the correct way to pool an interesting challenge (further exemplified in the direct OLS estimates of a single simulated draw of data in Panel B, where 95%-confidence intervals of policies from distinct pools strongly overlap), and it highlights the need of a disciplined procedure.

E.2.2. *Naive LASSO*. As outlined in section 3.3, there are two ways one could “naively” apply LASSO. The first is to disregard pooling, and, because sparse dosages might also mean a sparse set of policies, apply LASSO on the unique policy specification (2.1). While there is no theoretical issue with this procedure in terms of model consistency, using this for policy estimation leads to much the same performance limitations as direct OLS with regards to best policy estimation, namely a persistently high best policy MSE stemming from overly severe correction from Andrews et al. (2021)’s winner’s curse adjustment.

Figure E.2 which contrasts this “No pooling, only pruning” version of LASSO to TVA on best policy estimation, documents this. Panel A shows that while the MSE for our TVA estimator quickly converges to 0, the one for the LASSO estimator persistently lies above 0.1 regardless of n . Panels B and C verify that this is driven by winner’s curse attenuations, and not model selection issues. Panel B tracks the MSE conditional on both procedures selecting the right model of their respective specifications as the oracle; it is the same pattern as in Panel A. Panel C explicitly shows the shrinkage imposed by the best policy estimator; it is much higher when the model selection doesn’t pool.

The second way to “naively” apply LASSO is to consider both pooling and pruning as important, but adopt a sign inconsistent model selection procedure by applying LASSO directly on 2.4 without a Puffer transformation. Figure E.3 establishes the contrast with TVA. As expected, simulations attest to inconsistent support selection (Panel A). On the question of best policy estimation, it is more subtle. This naive LASSO does manage to identify at least some of the best policy, furthermore the actual MSE of the best policy is comparable to TVA (Panel C). However, a key deficiency from a policymaking perspective is that it fails to select the minimum dosage best policy with substantial probability relative to TVA (Panel D).⁴⁴

⁴⁴The reason why the naive LASSO will select some best policy stems from the fact that though it is not sign consistent, it is l_2 consistent, which means that it will select a strict superset of the correct support Meinshausen and Yu (2009). Basically, this will prune too little and pool too finely. In particular, as seen in our simulations, it will often cleave the best policy. The minimum dosage best policy is then at substantial risk of not being in the pool selected as the best in the data

E.2.3. *Debiased LASSO*. Because at a higher level we are interested in high dimensional inference ⁴⁵ one alternative to a two step process of model selection and inference is the so-called debiased LASSO [Javanmard and Montanari \(2014\)](#), [Javanmard and Montanari \(2018\)](#), [Van De Geer \(2019\)](#). The basic idea is that since LASSO permits high dimensional estimation but at the price of downwardly biased coefficients, but also since this bias is estimable, we can reverse the bias. In particular, to the LASSO coefficients we can add a debiasing term proportional to the subgradient at the ℓ_1 norm of the LASSO solution $\hat{\theta}^n$

$$\hat{\theta}^u = \hat{\theta}^n + (1/n)M\mathbf{X}^T(Y - \mathbf{X}\hat{\theta}^n)$$

It is also possible to supply standard errors for these debiased coefficients. Note, however, that these debiased coefficients are almost surely never exactly zero, so that there is no question of sparsity. We thus only need to consider applying debiased LASSO to (2.1).

In Figure E.4 we show that the debiased LASSO procedure suffers from the similar limitations as direct OLS estimation, especially with regards to best policy estimation. That it does better with winner’s curse attenuations relative to the same adjustments on direct OLS possibly indicates that it might interact better with those adjustments despite the fact that Gauss-Markov theorem guarantees that the unadjusted direct OLS must dominate the unadjusted debiased LASSO. Nevertheless, TVA sharply dominates both alternatives because it can pool.

E.2.4. *“Off the Shelf” Bayesian approaches: Spike and Slab LASSO*. Because we envision a world in which there are many potential treatment variant aggregations to be done through pooling and pruning, this attests to our prior about the environment. Indeed, LASSO estimates have a Bayesian interpretation in terms of Laplace priors. One can ask whether a more sophisticated, “explicitly” Bayesian approach can address our final objectives. This paradigmatically different route is the topic of future work. Below, we just show that “off the shelf” Bayesian approaches are unlikely to help. In particular, we show that a direct application of spike and slab formulations – the most intuitively relevant method – underperforms relative to our TVA procedure. The Spike and Slab LASSO uses a prior of the form

$$\pi(\beta|\gamma) = \prod_{i=1}^p [\gamma_i \underbrace{\psi_1(\beta_i)}_{\text{Slab}} + (1 - \gamma_i) \underbrace{\psi_0(\beta_i)}_{\text{Spike}}], \gamma \sim \pi(\gamma)$$

with γ_0 and γ_1 , two Laplace distributions with very high (λ_0) and a very low (λ_1) scale parameters respectively (i.e. $\lambda_0 \gg \lambda_1$). This allows solving for the posterior of both the model parameters as well as the model itself ([Ročková and George \(2018\)](#)).

In Figure E.5 we contrast a Puffer-transformed Bayesian Bootstrap Spike and Slab LASSO (BBSSL - [Nie and Ročková \(2022\)](#)) to TVA. The performance of BBSSL is very similar to

⁴⁵albeit still in a $K < n$ regime

that when applying Naive LASSO to the marginal specification (2.4). BBSSL is support inconsistent (Panel A) and is clearly outperformed on minimum dosage best policy selection (Panel D). It does however identify at least one best policy most of the time (Panel C) and has a best policy MSE close to that of TVA (Panel B).

E.3. Performance Under Five Sparsity Relaxation Regimes. Our main theoretical guarantees in Section 2 hold in an environment with exact sparsity and with marginal effect sizes uniformly bounded away from 0. Here we explore *practical* performance relaxing this in several plausible ways. Although performance of TVA suffers, it is still strong; moreover, TVA does better than the next best practical alternative of applying naive LASSO to the marginal specification (2.4).

We consider five regimes of sparsity and effect size relaxations. Although the support configurations are no longer necessarily of cardinality M , they are still randomly chosen as in Section 3.1. In Regimes 1 and 2 we relax exact zeros in the non-primary marginals to small effect sizes; these are either rapidly diminishing as $\Theta(\frac{1}{n})$ (Regime 1) or moderately diminishing as $\Theta(\frac{1}{\sqrt{n}})$ (Regime 2).⁴⁶ In Regime 3 we further relax sparsity in the first regime by expanding the true support to include marginal effects of both large and medium sizes. In Regime 4 we relax the lower bound on marginal effect sizes for the primary marginals, diminishing at a rate $\Theta(\frac{1}{n^{0.2}})$ between moderate and rapid. In the fifth regime we further relax sparsity in Regime 4 by expanding the support with rapidly diminishing coefficients. A summary of the regime configurations is described below:

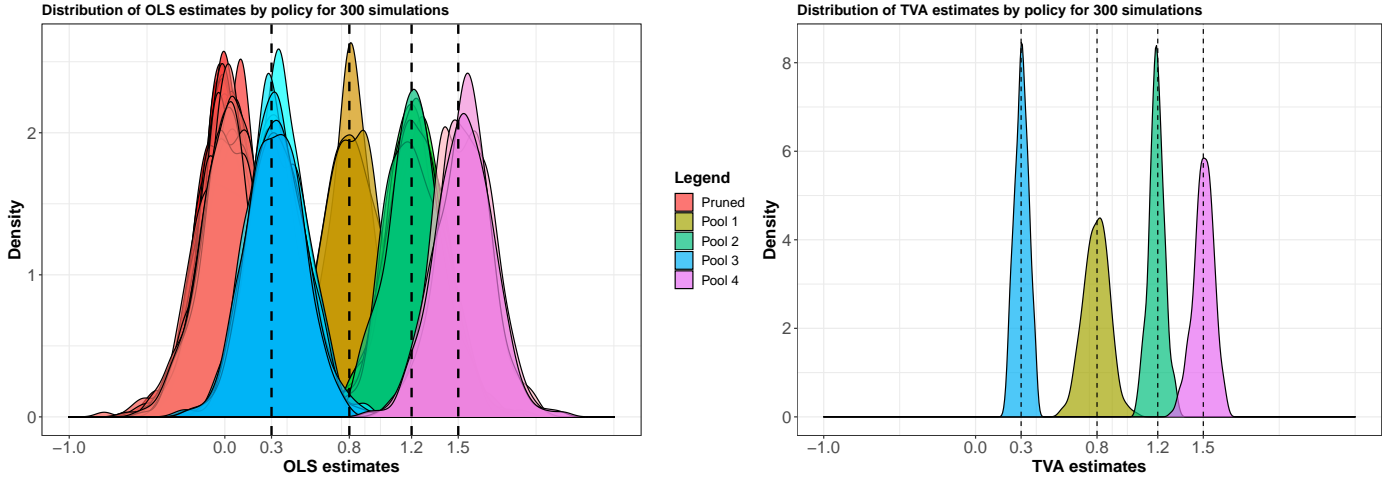
- (1) **Regime 1:** M constant marginal effect sizes in $[1, 5]$ & M rapidly diminishing remaining marginal effect sizes in $[1, 5]/n$.
- (2) **Regime 2:** M constant marginals in $[1, 5]$ & M moderately diminishing remaining marginals $([1, 5]/\sqrt{n})$.
- (3) **Regime 3:** M large constant marginals in $[5, 10]$, M medium marginals in $[1, 2]$ & M rapidly diminishing remaining marginals in $[1, 5]/n$.
- (4) **Regime 4:** M decreasing marginals in $[1, 5]/n^{0.2}$ (and zero marginals everywhere else).
- (5) **Regime 5:** M decreasing marginals in $[1, 5]/n^{0.2}$ & M moderately diminishing remaining marginals in $[1, 5]/\sqrt{n}$.

The main finding is that support accuracy of TVA is generally strong. Even in the case of model misspecification, the MSE of the best policy is still low for moderate sample size. Furthermore its distinct advantage relative to alternatives with regards to best policy estimation — the much more reliable selection of the minimum dosage best policy — remains equally strong.


⁴⁶Note that only a diminishing rate retains a threat of misspecified model selection at any n .

Figure E.6, panel A speaks to the first pattern: even if support accuracy suffers from relaxing sparsity requirements performance remains generally high. For regimes 1-3 for example, support accuracy converges to 100% albeit more slowly than in the exactly sparse environment (and convergence in R2 is, as expected, slower than in R1). For regimes 4 and 5, even though TVA is support inconsistent, this does not imply a higher MSE of the best policy demonstrating that model misspecification is not generally threatening with regards to this final objective. Panel B shows this very clearly where MSE is steadily decreasing and comparable across all regimes. Performance on some best policy selection — the easier task — seems also particularly unaffected by the sparsity relaxations, with accuracy ranging between 80%-90% (panel C). Finally TVA’s distinctive advantage over all alternatives explored in Section 3.3 is the reliable selection of the minimum dosage best policy. Performance on minimum best policy selection remains generally strong across regimes (panel D) with a steady increase as sample size grows. Notably the minimum best inclusion rate is above 90% across all regimes for even moderate sample sizes ($n = 3000$).

Finally, Figures E.7-E.11 demonstrate that TVA strongly outperforms naive LASSO across regimes for support accuracy and minimum dosage best policy selection and sometimes significantly so (e.g. Regime 3, Figure E.9). This suggests that naive LASSO is an unhelpful way to proceed particularly if one were interested in the post-LASSO set in addition to the best policy. Both methods then, are equally strong at identifying at least one best policy and with respect to the final MSE of the best policy with performance mostly ranging between 80%-100% for the former and never higher than 0.75 (R1, R3) for the latter. However only TVA helps with the joint best policy objectives of low MSE while reliably choosing the minimum dosage best policy. This is compelling in regimes 2 and 3 where LASSO’s minimum best selection rate hovers around 60-70% against TVA reaching 90% accuracy for a moderate sample size ($n = 3000$) already.



(A) Densities (300 simulations)



(B) Estimated Coefficients

FIGURE E.1. This figure presents results from a simulation setting with $n = 4000$ and where 4 pooled policies (each composed of 3-8 unique policies) are non-zero with effects 0.3, 0.8, 1.2 and 1.5 respectively. For OLS applied to (2.1) and TVA, we show both the distribution of policy estimates across 300 simulations (panel A) and the estimated policy coefficients for one representative simulated draw (panel B). Note that the color labeling for OLS plots corresponds to the true underlying pooled and pruned policies. For the TVA density plots, we condition on the event that the TVA estimator has selected the correct support (mean support accuracy is 92.2% across simulations).




FIGURE E.2. A plot comparing the performance of the TVA estimator to applying LASSO on the unique policy specification (2.1). Panel A compares the MSE of the hybrid estimators of Andrews et al. (2021) for winner's curse-adjusted best policy estimation for both methods. Panel B is exactly the same but conditional on selecting the true support in (2.4) and (2.1). Panel C compares the amount of shrinkage imposed by the winner's curse adjustment for both methods, in percentage of the initial coefficient. In all simulations, there are 20 simulations per n

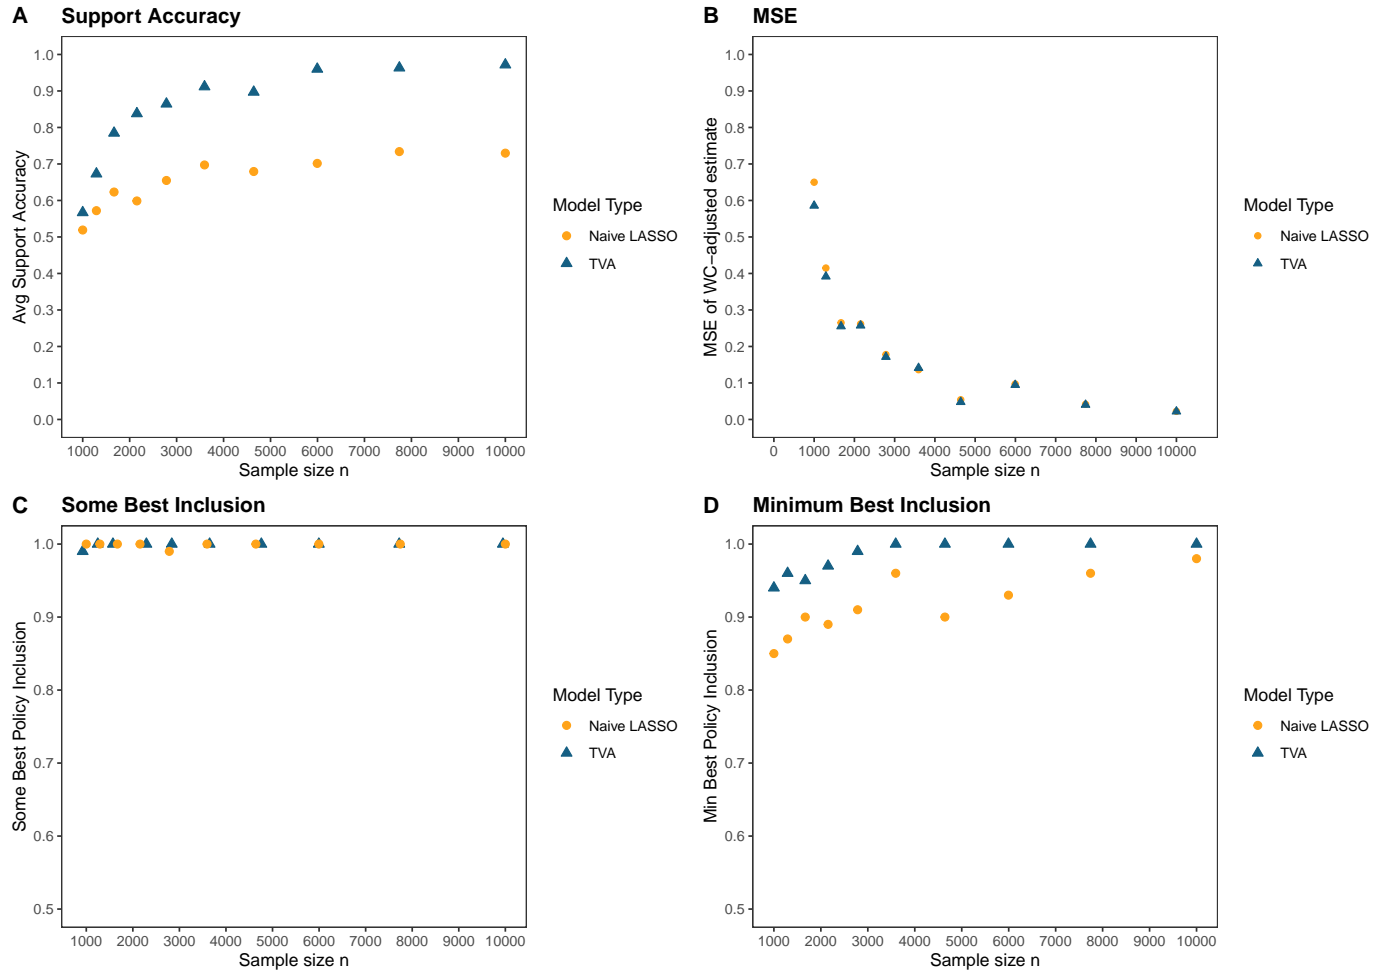


FIGURE E.3. A plot comparing the performance of the TVA estimator to applying LASSO on (2.4) using Chernozhukov et al. (2015). Panel A compares average support accuracy and Panel B compares MSE of the best policy treatment effect as a function of sample size n . Panels C and D look at best policy selection accuracies. On panel C points are slightly jittered for better readability. There are 20 simulations per support configuration per n , for five support configurations.

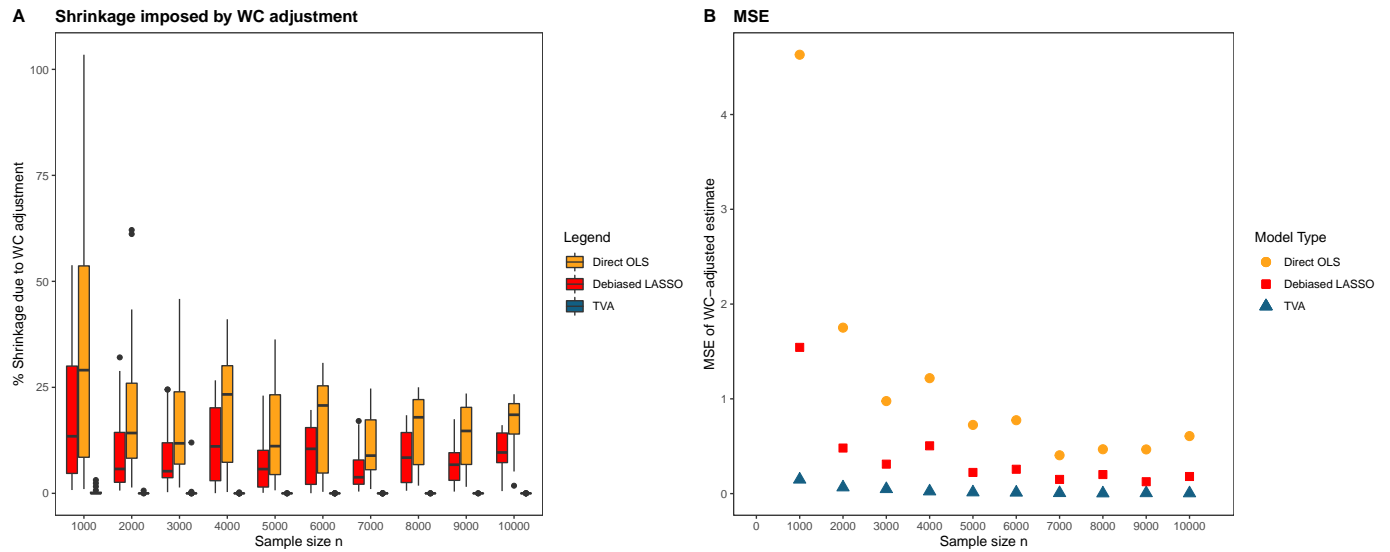


FIGURE E.4. A plot comparing the performance of the TVA estimator to applying OLS on (2.1) and a debiased LASSO on (2.1) following [Javanmard and Montanari \(2014\)](#). Panel A compares the amount of shrinkage imposed by the winner's curse adjustment as percentage of the initial coefficient. Panel B compares the mean squared error of the best policy estimate as a function of sample size n . There are 20 simulations per n .

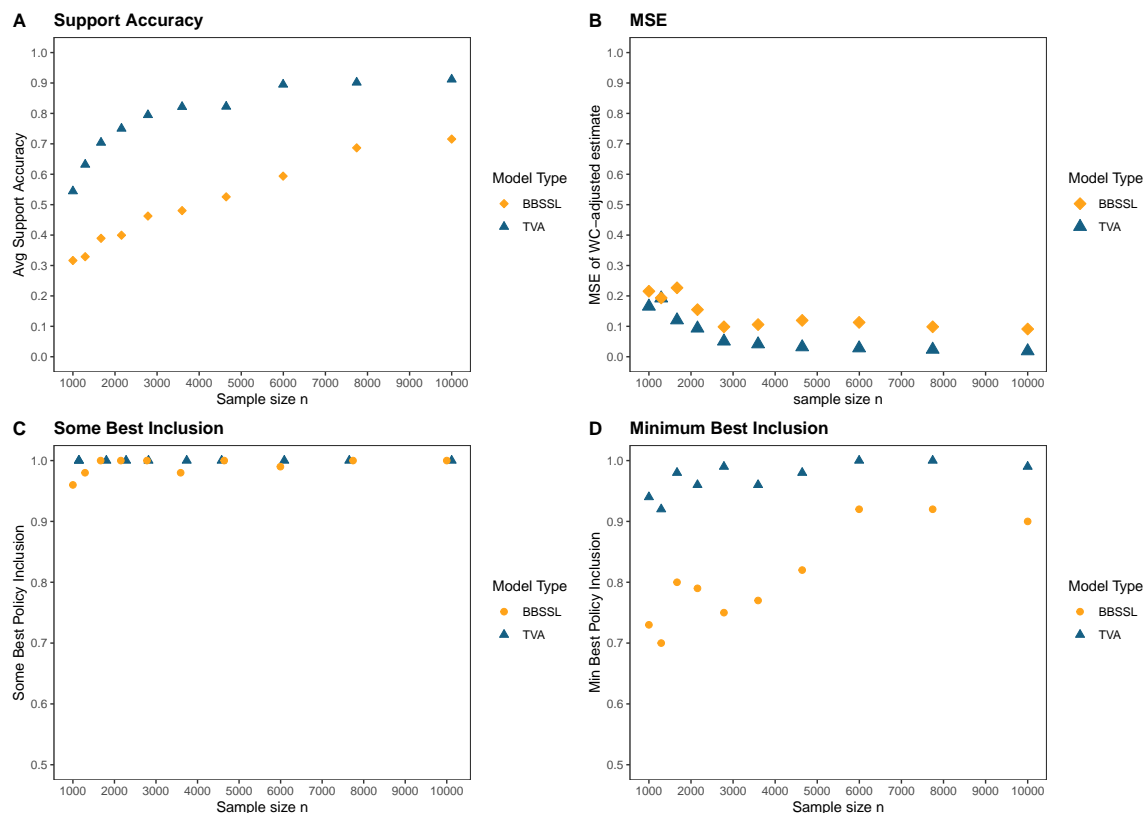





FIGURE E.5. A plot comparing the performance of the TVA estimator to applying Bayesian Bootstrap Spike and Slab LASSO (Nie and Ročková (2022)) on (2.4). Panel A compares support accuracy and panel B compares MSE on the final WC adjusted estimate. Panels C and D compare best policy inclusion measures as a function of n . For BBSSL the Laplace parameters are chosen as $\lambda_1 = 10^{-n/100}$ and λ_0 ranging from 1 to 10^5 (100 steps). This yields a solution path as a function of λ_0 and policies selected at least 95% of times along the solution path compose the final support. There are 20 simulations per support configuration per n for five support configurations.




(A) Support Accuracy



(B) MSE of Best Policy



(C) Some Best Policy Inclusion



(D) Min Best Policy Inclusion

FIGURE E.6. This figure shows the performance of the TVA estimator under relaxations of the sparsity assumptions. We show support accuracy (panel A), MSE of the best policy (panel B), Some Best policy inclusion rate (panel C) and Minimum Best policy inclusion rate (panel D) as a function of n for five different levels of violating sparsity requirements (R1-R5). Per regime, there are 20 simulations per support configuration per n , for five support configurations.

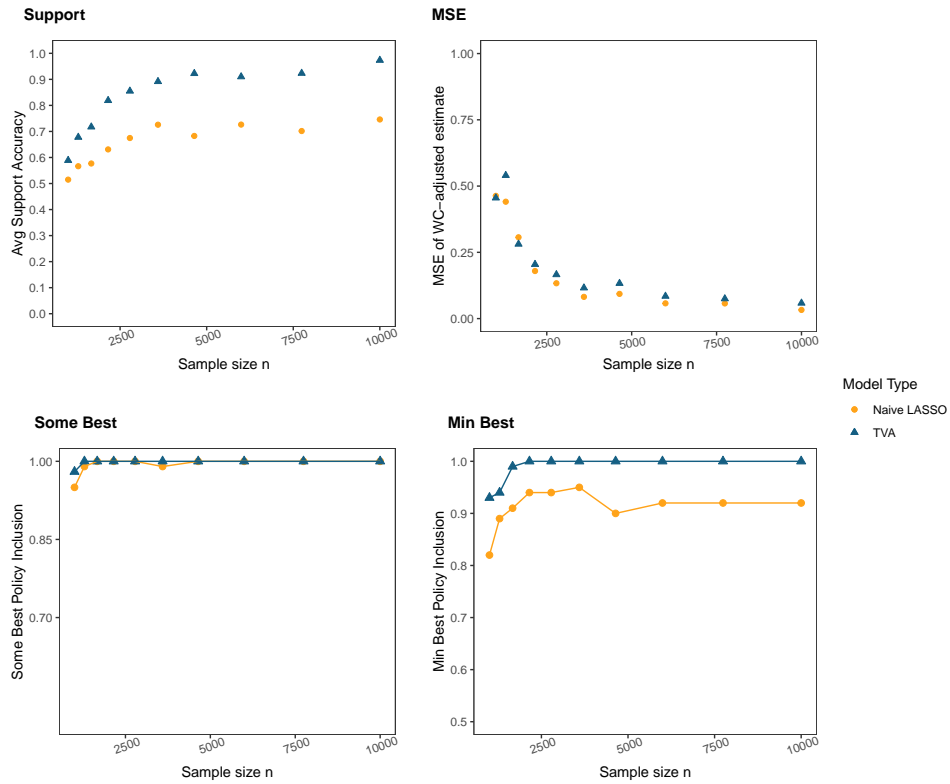


FIGURE E.7. A comparison between the TVA estimator and a direct implementation of LASSO on (2.4) for support accuracy, MSE and best policy inclusion measures, under **regime 1**. There are 20 simulations per support configuration per n , for five support configurations.




FIGURE E.8. A comparison between the TVA estimator and a direct implementation of LASSO on (2.4) for support accuracy, MSE and best policy inclusion measures, under **regime 2**. There are 20 simulations per support configuration per n , for five support configurations.




FIGURE E.9. A comparison between the TVA estimator and a direct implementation of LASSO on (2.4) for support accuracy, MSE and best policy inclusion measures, under **regime 3**. There are 20 simulations per support configuration per n , for five support configurations.

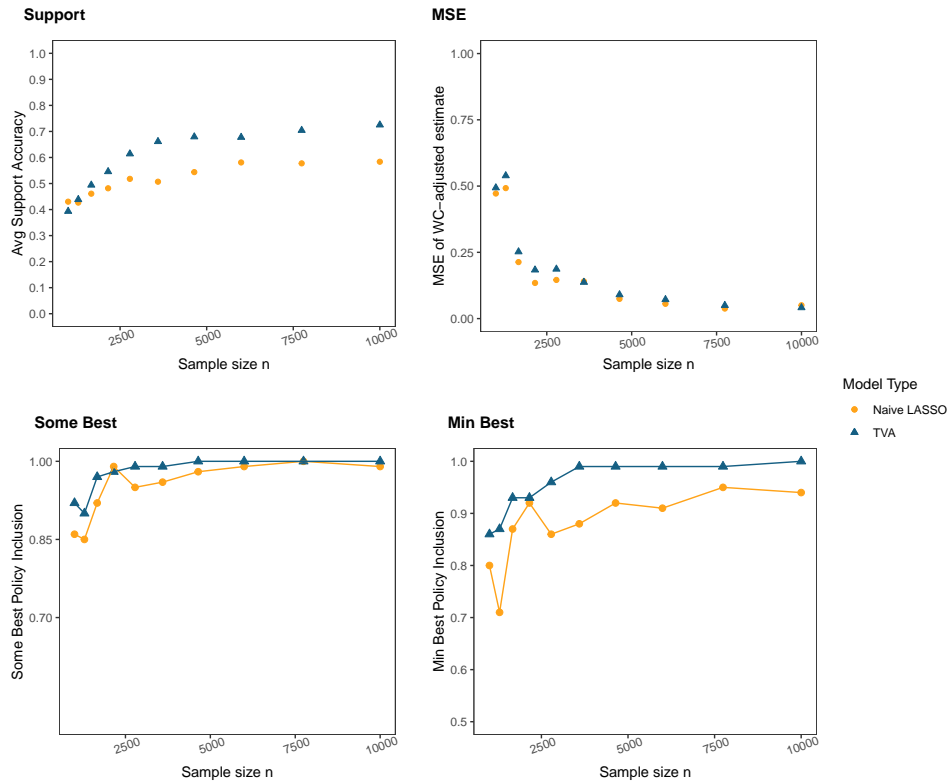


FIGURE E.10. A comparison between the TVA estimator and a direct implementation of LASSO on (2.4) for support accuracy, MSE and best policy inclusion measures, under **regime 4**. There are 20 simulations per support configuration per n , for five support configurations.




FIGURE E.11. A comparison between the TVA estimator and a direct implementation of LASSO on (2.4) for support accuracy, MSE and best policy inclusion measures, under **regime 5**. There are 20 simulations per support configuration per n , for five support configurations.

APPENDIX F. EXTENDED ROBUSTNESS

This section complements the discussion of Appendix D. We report the results from a fully saturated regression on the 75 unique policies to emphasize that some pooling choices made by TVA can be non-trivial making eyeballing a generally *non*-reliable sanity check. We also elaborate more on the bootstrapping analysis suggested in Appendix D to explore TVA performance in the context of our experiment.

Fully saturated regression Below we report the results from running a saturated regression of 75 raw coefficients for finely differentiated policies. We contrast this with the pooling and pruning choices made by the TVA estimator highlighted in different colors (red policies are pruned while green / blue policies are pooled together). This makes the point very clearly that a simple eye-balling of these results would have been misleading in inferring the choices made by TVA. Two examples are worth noting:

- Some seemingly efficient policies at standard confidence levels are pruned while others are pooled together: this is striking when comparing the policies in Figure F.2 (Shots/\$), Panel A (No Seed) in the last pooling profile (all were pruned) with those in Panel B (Trusted Seed) in the last pooling profile (all were pooled and kept in the support).
- Some policies that are individually underpowered (though positive) were pooled together (and kept in the support) suggesting that the pool ended up being powered and highly effective. An example can be seen in Figure F.1 (Measles Shots), Panel C (Gossip Seed) in the last pooling profile.

Bootstrapping analysis It is natural to ask about the fragility of TVA to the particular draw of the data; precisely this concern motivates, for example, our implementation of winner’s curse adjustments by Andrews et al. (2021), as well as simulations in Section 3 that directly speak to the variance of TVA. However, one might further wonder about just the observations in our dataset, with a concern akin to one about leverage of observations. An intuitive approach to address this is a bootstrapping analysis, where TVA is run on multiple bootstrapped samples. We can then speak to variation in both the set of supports selected as well the estimates of the pooled policies. Because this is a more exploratory analysis, its principal value lies in speaking to *relative* stability of conclusions between the two policies for the two outcomes. ⁴⁷

⁴⁷Note that we have slightly different goals here from the issue of bootstrapped standard errors

The 200 bootstrapped samples we run for each outcome are stratified at the policy level and results for each sample are displayed in Figures F.3 (immunizations) and F.4 (immunizations /\$). For immunizations/\$, the support from the original sample was LASSO selected in 96% of bootstrapped samples and the estimated coefficients are concentrated around our main estimate, including the best policy (Info Hubs (All), No Incentives, SMS (All)). Notably there is almost no winner's curse adjustment since the best policy is consistently well separated from the second best. For the immunizations outcome we again observe little variation in the support, and the most effective policy (Info Hubs, Slopes (All), SMS (All)) is identified as such in 77% of the bootstrapped samples. However, there is considerably more variation in the winner's curse estimates, with some bootstrap samples sharply attenuating the best policy estimate. Taken altogether, this speaks to tighter competition and more sensitivity to leverage of certain observations for immunizations than immunizations/\$.




FIGURE F.1. A plot showing the coefficients from the OLS regression on all 75 unique policies for the Measles shot outcome. Panels are organized by seeds and within each panel the pooling profiles are delimited by dashed lines. Policies shown in red are policies pruned by the TVA Estimator while policies shown in other colors were pooled together.




FIGURE F.2. A plot showing the coefficients from the OLS regression on all 75 unique policies for the Shots/\$ outcome. Panels are organized by seeds and within each panel the pooling profiles are delimited by dashed lines. Policies shown in red are policies pruned by the TVA Estimator while policies in blue were pooled together.




FIGURE F.3. Results of post-LASSO OLS estimates for the Measles Shot outcome are shown for 200 bootstrapped samples stratified at the policy level. The raw data marginal effects support is shown in white and “Bootstrap Supports” 1 and 2 are the two unique supports that were selected across all bootstrapped samples. We also report winner’s curse adjusted coefficients for the raw data (green) and the bootstrapped samples (red). On the x-axis, the raw data best pooled policy is highlighted in red and was selected in 77% of the bootstrapped samples. Finally the minimum dosage best policy (Info Hubs, Low Slopes, Low SMS) is only selected when the TVA estimator actually selects the best pooled policy, hence 77% of the time.




FIGURE F.4. Results of post-LASSO estimates for the Shots/\$ outcome are shown for 200 bootstrapped samples stratified at the policy level. The raw data marginal effects support is shown in white and “Bootstrap Supports” 1-3 are the three unique supports that were selected across all bootstrapped samples. We also report winner’s curse adjusted coefficients for the raw data (green) and the bootstrapped samples (red). On the x-axis, the raw data best pooled policy is highlighted in red and was selected in 96% of the bootstrapped samples. Finally, the minimum dosage best policy (Info Hubs, No Incentives, Low SMS) is only selected when the TVA estimator actually selects the best pooled policy, hence 96% of the time.

APPENDIX G. WINNER'S CURSE EXTENDED

G.1. Which policies can be local alternatives? As we said, nothing prevents policies in different treatment profiles from being local alternatives. But in fact, there can be other non pooled policies within a treatment profile that can be local alternatives. Proposition G.1 fully characterizes the local alternatives respecting the assumptions on the environment, particularly Assumption 3.

PROPOSITION G.1. *Let $\kappa, \kappa' \in S_{TVA}$ be local alternatives, i.e., $\eta_{S_{TVA}, \kappa}^0 = \eta_{S_{TVA}, \kappa'}^0 + \frac{r_{\kappa\kappa'}}{\sqrt{n}}$ for some $r_{\kappa\kappa'}$ fixed in n . Then, for α^0 to respect Assumption 3, one of the following has to hold:*

- (1) $P(\kappa) \neq P(\kappa')$ (i.e., κ and κ' have different treatment profiles ⁴⁸), or
- (2) $P(\kappa) = P(\kappa')$, and κ and κ' are nowhere adjacent in the Hasse diagram.⁴⁹

These conditions are also sufficient to allow for local alternatives, in that for any $S_{TVA} \in \mathcal{P}_{|\Delta}$, all pairs $\kappa, \kappa' \in S_{TVA}$ meeting condition (1) or (2) can be simultaneously made local alternatives by some choice of α^0 satisfying Assumption 3.

Proof. The only non-trivial case is for κ and κ' that are policy variants, i.e. have the same treatment profile. Here the proof follows the basic intuition from the Hasse diagram. If κ and κ' were adjacent anywhere in the diagram (say for some treatment combinations k, k' such that κ pools k and κ' pools k') then $\alpha_{\min\{k, k'\}} = \frac{r_{\kappa\kappa'}}{\sqrt{n}}$, violating Assumption 3. On the other hand, if κ and κ' are not adjacent anywhere, one can consider the policies "in between" them, i.e. the policies $\zeta_1, \dots, \zeta_n \in S_{TVA}$ such that κ is adjacent to ζ_1 , κ' is adjacent to ζ_n , ζ_i is adjacent to ζ_{i+1} for $i \notin \{1, n\}$, and the union of these ζ_i pools all z such that $\min\{k, k'\} < z < \max\{k, k'\}$ and which are neither pooled by κ nor κ' . Then κ and κ' can be made local alternatives with the relevant "in between" marginals from α^0 satisfying Assumption 3, by making the $|\eta_{S_{TVA}, \zeta_i}^0|$ sufficiently large. Following this construction where applicable (in each stage of the iterative procedure where a new pair is made into local alternatives, to ensure that prior local alternatives remain local alternatives, the absolute magnitude of the policy effects of the prior alternatives may have to be increased), an entire set α^0 satisfying Assumption 3 ensues. ■

G.2. Does Andrews et al. (2021) extend to asymptotic normality and model selection? The main text of Andrews et al. (2021) focuses on the case where the estimators in question, $X(\kappa)$ are *exactly* jointly normally distributed. While two extensions are presented, one for a conditioning event such as model-selection (addressed in their Appendix A) and

⁴⁸This is a slight abuse of notation, since $P(\cdot)$ was defined originally over treatment combinations, not pooled policies. So, $P(\cdot)$ here is simply the well defined extension to the latter.

⁴⁹Formally, this is the condition that for any treatment combinations k, k' such that κ pools k and κ' pools k' , either k and k' are incomparable or there is a treatment combination z pooled by neither κ nor κ' such that $\min\{k, k'\} < z < \max\{k, k'\}$.

another for the case of asymptotic normality which is required for practical settings such as regression (in their Appendix D), the paper does not formally work out the case with both issues present.

We have both in our setting, under Assumptions 1-6. We have a conditioning event ($\hat{S}_{TVA} = S_{TVA}$) occurring with probability tending to one and we are in a regression setting with asymptotic normality. So while the theoretical properties of the estimator in our setting are highly plausible and coherent with our simulations below, the extension to the nested case of model selection remains to be proven. It is beyond the scope of our present paper to nest both of their extensions, and we leave it for future work.

So, in Section 2.5 of the main text, we assume that the distribution is exact. This allows us to focus on how local alternatives in our Hasse diagram may impact the problem. Therefore, we assume

$$X \sim \mathcal{N}(\mu, \Omega).$$

Then we can exactly apply the results of Proposition 6 in Andrews et al. (2021), and build the hybrid confidence set.

APPENDIX H. APPENDIX FIGURES

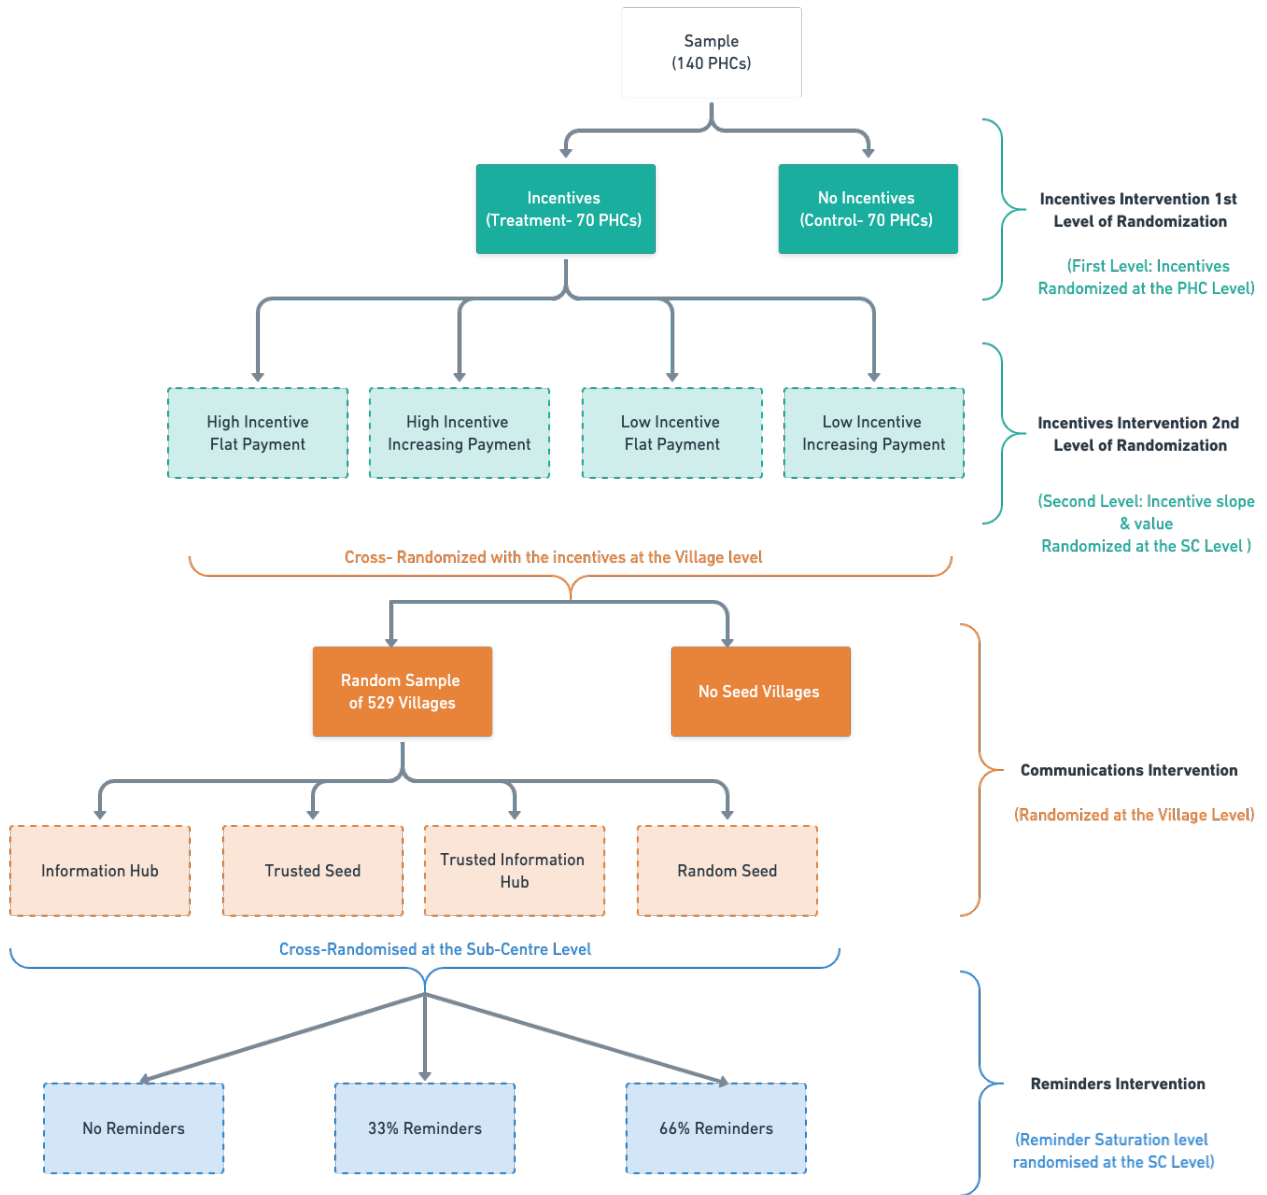


FIGURE H.1. Experimental Design




FIGURE H.2. Effects on the number of measles vaccinations relative to control (5.29) by reminders, incentives, and seeding policies, for the full sample. The specification includes a dummy for being in the Ambassador Sample, is weighted by village population, controls for district-time fixed effects, and clusters standard errors at the sub-center level.

National Immunization Schedule (NIS) for Infants, Children and Pregnant Women

Vaccine	When to give	Dose	Route	Site
For Pregnant Women				
TT-1	Early in pregnancy	0.5 ml	Intra-muscular	Upper Arm
TT-2	4 weeks after TT-1*	0.5 ml	Intra-muscular	Upper Arm
TT- Booster	If received 2 TT doses in a pregnancy within the last 3 yrs*	0.5 ml	Intra-muscular	Upper Arm
For Infants				
BCG	At birth or as early as possible till one year of age	0.1ml (0.05ml until 1 month age)	Intra-dermal	Left Upper Arm
Hepatitis B - Birth dose	At birth or as early as possible within 24 hours	0.5 ml	Intra-muscular	Antero-lateral side of mid-thigh
OPV-0	At birth or as early as possible within the first 15 days	2 drops	Oral	Oral
OPV 1, 2 & 3	At 6 weeks, 10 weeks & 14 weeks (OPV can be given till 5 years of age)	2 drops	Oral	Oral
Pentavalent 1, 2 & 3	At 6 weeks, 10 weeks & 14 weeks (can be given till one year of age)	0.5 ml	Intra-muscular	Antero-lateral side of mid-thigh
Rotavirus#	At 6 weeks, 10 weeks & 14 weeks (can be given till one year of age)	5 drops	Oral	Oral
IPV	Two fractional dose at 6 and 14 weeks of age	0.1 ml	Intra dermal two fractional dose	Intra-dermal: Right upper arm
Measles /MR 1 st Dose§	9 completed months-12 months. (can be given till 5 years of age)	0.5 ml	Sub-cutaneous	Right upper Arm
JE - 1**	9 completed months-12 months.	0.5 ml	Sub-cutaneous	Left upper Arm
Vitamin A (1 st dose)	At 9 completed months with measles-Rubella	1 ml (1 lakh IU)	Oral	Oral
For Children				
DPT booster-1	16-24 months	0.5 ml	Intra-muscular	Antero-lateral side of mid-thigh
Measles/ MR 2 nd dose §	16-24 months	0.5 ml	Sub-cutaneous	Right upper Arm
OPV Booster	16-24 months	2 drops	Oral	Oral
JE-2	16-24 months	0.5 ml	Sub-cutaneous	Left Upper Arm
Vitamin A*** (2 nd to 9 th dose)	16-18 months. Then one dose every 6 months up to the age of 5 years.	2 ml (2 lakh IU)	Oral	Oral
DPT Booster-2	5-6 years	0.5 ml.	Intra-muscular	Upper Arm
TT	10 years & 16 years	0.5 ml	Intra-muscular	Upper Arm

- *Give TT-2 or Booster doses before 36 weeks of pregnancy. However, give these even if more than 36 weeks have passed. Give TT to a woman in labour, if she has not previously received TT.
- **JE Vaccine is introduced in select endemic districts after the campaign.
- *** The 2nd to 9th doses of Vitamin A can be administered to children 1-5 years old during biannual rounds, in collaboration with ICDS.
- #Phased introduction, at present in Andhra Pradesh, Haryana, Himachal Pradesh and Orissa from 2016 & expanded in Madhya Pradesh, Assam, Rajasthan, and Tripura in February 2017 and planned in Tamil Nadu & Uttar Pradesh in 2017.
- § Phased introduction, at present in five states namely Karnataka, Tamil Nadu, Goa, Lakshadweep and Puducherry. (As of Feb' 2017)

FIGURE H.3. National Immunization Schedule for Infants, Children, and Pregnant Women.

Data collection activity	Survey Sample	Timeline	Purpose
PHC Survey	161 PHCs across 7 districts	March – April 2015	Sampling PHCs for the study
Census for Baseline	328,058 households in 912 sample villages	September – November 2015	Identifying sampling frame for baseline survey
Baseline Survey	17,000 children in 912 sample villages	May – July 2016	Collecting baseline data
Nominations Survey	15,504 households in 912 sample villages	December 2015 – February 2016	Identifying and sampling seeds
Seeds Survey	2,117 nominated respondents	June – August 2016	Obtaining consent and information on seeds
Child verification survey		June 2017 – March 2018	Verify information entered by ANMs on tablets
New Census	~120,000 households, in 200 sample villages	February – April 2018	Identifying sampling frame for endline survey
Revised Endline Survey	4000 children in 200 sample villages	May – June 2018	Collecting outcome data

FIGURE H.4. Overview of Survey Data Collection Activities.

APPENDIX I. SUBSTITUTION PATTERNS

TABLE I.1. Incentive Treatment Effects for Non-Tablet Children from Endline Data

	<i>Dependent variable:</i>						
	At Least 2	At Least 3	At Least 4	At Least 5	At Least 6	At Least 7	Measles 1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
High Slope	-0.158 (0.062)	-0.052 (0.072)	-0.076 (0.093)	-0.196 (0.106)	-0.187 (0.101)	-0.027 (0.105)	-0.135 (0.108)
High Flat	-0.021 (0.088)	-0.024 (0.063)	-0.091 (0.078)	-0.078 (0.155)	-0.053 (0.152)	0.102 (0.167)	0.185 (0.143)
Low Slope	0.090 (0.064)	0.175 (0.060)	0.104 (0.085)	-0.026 (0.099)	-0.152 (0.080)	-0.079 (0.077)	0.051 (0.100)
Low Flat	0.004 (0.076)	0.069 (0.096)	-0.010 (0.122)	-0.110 (0.176)	-0.005 (0.160)	-0.079 (0.173)	-0.102 (0.148)
Control Mean	0.69	0.54	0.4	0.31	0.17	0.11	0.39
Total Obs.	1179	1165	1165	1042	1042	706	613
Zeros Replaced	0	0	0	0	0	0	0

Note: Specification includes District Fixed Effects, and a set of controls for seeds and reminders. Control mean shown in levels, and standard errors are clustered at the SC Level

TABLE I.2. Seeds Treatment Effects for Non-Tablet Children from Endline Data

	<i>Dependent variable:</i>						
	At Least 2	At Least 3	At Least 4	At Least 5	At Least 6	At Least 7	Measles 1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Random	-0.101 (0.059)	-0.045 (0.066)	0.0003 (0.088)	-0.143 (0.122)	0.015 (0.102)	0.058 (0.102)	-0.017 (0.101)
Information Hub	-0.040 (0.082)	-0.121 (0.080)	-0.025 (0.113)	-0.112 (0.135)	0.018 (0.123)	-0.105 (0.073)	-0.092 (0.119)
Trusted	0.034 (0.070)	-0.033 (0.073)	0.111 (0.100)	0.027 (0.128)	0.147 (0.118)	-0.011 (0.107)	0.106 (0.111)
Trusted Information Hub	-0.103 (0.075)	-0.082 (0.079)	-0.031 (0.100)	-0.209 (0.113)	-0.099 (0.086)	-0.057 (0.082)	-0.344 (0.099)
Control Mean	0.78	0.66	0.5	0.58	0.3	0.22	0.62
Total Obs.	469	461	461	389	389	251	231
Zeros Replaced	0	0	0	0	0	0	0

Note: Specification includes District Fixed Effects, and a set of controls for incentives and reminders. Control mean shown in levels, and standard errors are clustered at the SC Level

TABLE I.3. Reminders Treatment Effects for Non-Tablet Children from End-line Data

	<i>Dependent variable:</i>						Measles 1
	At Least 2	At Least 3	At Least 4	At Least 5	At Least 6	At Least 7	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
33%	0.079 (0.058)	0.093 (0.066)	0.070 (0.085)	0.033 (0.104)	0.019 (0.094)	-0.138 (0.080)	0.011 (0.086)
66%	0.031 (0.053)	0.096 (0.055)	0.042 (0.069)	-0.074 (0.091)	-0.073 (0.079)	-0.044 (0.067)	-0.097 (0.084)
Control Mean	0.64	0.48	0.34	0.28	0.19	0.13	0.46
Total Obs.	1179	1165	1165	1042	1042	706	613
Zeros Replaced	0	0	0	0	0	0	0

Note:

Specification includes District Fixed Effects, and a set of controls for seeds and incentives. Control mean shown in levels, and standard errors are clustered at the SC Level

APPENDIX J. DATA VALIDATION

A household survey was conducted to monitor program implementation at the child-level—whether the record entered in the tablet corresponded to an actual child, and whether the data entered for this child was correct. This novel child verification exercise involved J-PAL field staff going to villages to find the households of a set of randomly selected children which, according to the tablet data, visited a session camp in the previous four weeks. Child verification was continuous throughout the program implementation, and the findings indicate high accuracy of the tablet data. We sampled children every week to ensure no additional vaccine was administered in the lag between them visiting the session camp and the monitoring team visiting them. Data entered in the tablets was generally of high quality. There were almost no incidences of fake child records, and the child’s name and date of birth were accurate over 80% of the time. For 71% of children the vaccines overlapped completely (for all main vaccines under age of 12 months). Vaccine-wise, on average, 88% of the cases had matching immunization records. Errors seem genuine, rather than coming from fraud: they show no systematic pattern of inclusion or exclusion and are no different in any of the treatment groups.

APPENDIX K. BASELINE STATISTICS

TABLE J.1. Selected Baseline Statistics of Haryana Immunization

	Population-Weighted Average
Baseline Covariates–Demographic Variables	
<i>(Village Level Averages)</i>	
Fraction participating in Employment Generating Schemes	0.045
Fraction Below Poverty Line (BPL)	0.187
Household Financial Status (on 1-10 scale)	3.243
Fraction Scheduled Caste-Scheduled Tribes (SC/ST)	0.232
Fraction Other Backward Caste (OBC)	0.21
Fraction Hindu	0.872
Fraction Muslim	0.101
Fraction Christian	0.001
Fraction Buddhist	0
Fraction Literate	0.771
Fraction Unmarried	0.05
Fraction of Adults Married (living with spouse)	0.504
Fraction of Adults Married (not living with spouse)	0.002
Fraction of Adults Divorced or Separated	0.001
Fraction Widow or Widower	0.039
Fraction who Received Nursery level Education or Less	0.17
Fraction who Received Class 4 level Education	0.086
Fraction who Received Class 9 level Education	0.158
Fraction who Received Class 12 level Education	0.223
Fraction who Received Graduate or Other Diploma level Education	0.081
Baseline Covariates–Immunization History of Older Cohort	
<i>(Village Level Averages)</i>	
Number of Vaccines Administered to Pregnant Mother	2.271
Number of Vaccines Administered to Child Since Birth	4.23
Fraction of Children who Received Polio Drops	0.998
Number of Polio Drops Administered to Child	2.989
Fraction of Children who Received an Immunized Card	0.877
Number of Observations	
Villages	903

APPENDIX L. INFORMATION HUB QUESTIONS

- (1) Random seeds: In this treatment arm, we did not survey villages. We picked six ambassadors randomly from the census.
- (2) Information hub seed: Respondents were asked to identify who is good at relaying information.

We used the following script to ask the question to the 17 households:

“Who are the people in this village, who when they share information, many people in the village get to know about it. For example, if they share information about a music festival, street play, fair in this village, or movie shooting many people would learn about it. This is because they have a wide network of friends, contacts in the village and they can use that to actively spread information to many villagers. Could you name four such individuals, male or female, that live in the village (within OR outside your neighbourhood in the village) who when they say something many people get to know?”

- (3) “Trust” seed: Respondents were asked to identify those who are generally trusted to provide good advice about health or agricultural questions (see appendix for script)

We used the following script to elicit who they were:

“Who are the people in this village that you and many villagers trust, both within and outside this neighbourhood? When I say trust I mean that when they give advice on something, many people believe that it is correct and tend to follow it. This could be advice on anything like choosing the right fertilizer for your crops, or keeping your child healthy. Could you name four such individuals, male or female, who live in the village (within OR outside your neighbourhood in the village) and are trusted?”

- (4) “Trusted information hub” seed: Respondents were asked to identify who is both trusted and good at transmitting information

“Who are the people in this village, both within and outside this neighbourhood, who when they share information, many people in the village get to know about it. For example, if they share information about a music festival, street play, fair in this village, or movie shooting many people would learn about it. This is because they have a wide network of friends/contacts in the village and they can use that to actively spread information to many villagers. Among these people, who are the people that you and many villagers trust? When I say trust I mean that when they give advice on something, many people believe that it is correct and tend to follow it. This could be advice on anything like choosing the right fertilizer for your

crops, or keeping your child healthy. Could you name four such individuals, male or female, that live in the village (within OR outside your neighbourhood in the village) who when they say something many people get to know and are trusted by you and other villagers?”