


When do default nudges work?

Carl Bonander¹, Mats Ekman² and Niklas Jakobsson ^{2,*}

¹School of Public Health and Community Medicine, Institute of Medicine, University of Gothenburg, 405 30 Gothenburg, Sweden

²Karlstad Business School, Karlstad University, 651 88 Karlstad, Sweden

*Correspondence address. E-mail: niklas.jakobsson@kau.se

Abstract

Nudging is a burgeoning topic in science and in policy, but evidence on the effectiveness of nudges among differentially incentivized groups is lacking. This paper exploits regional variations in the rollout of the Covid-19 vaccine in Sweden to examine the effect of a nudge on groups whose intrinsic incentives are different: 16- to 17-year-olds, for whom Covid-19 is not dangerous, and 50- to 59-year-olds, who face a substantial risk of death or severe disease. We find a significantly stronger response in the younger group compared with the older (11.7 vs 3.6 percentage point increase in our study period), consistent with the theory that nudges are more effective for choices that are not meaningful to the individual.

Keywords: behavioral intervention, default, decision-making, nudge, pre-booked.

Introduction

Nudges have established themselves as attractive policy tools for influencing actions without removing options. Society sometimes benefits from people taking certain decisions, but individuals may lack incentives to take them, or they may prefer to make socially harmful choices. Numerous instances have been found where small changes in the choice architecture bring changes in decisions. These range from situations involving participation in retirement-savings plans (Madrian and Shea 2001; Atalay et al. 2014) to eating healthy (Cheung et al. 2019), eliciting organ donations (Li et al. 2013) or tipping behavior (Hoover 2022). Despite this burgeoning interest, a largely unresolved empirical issue is whether nudges sway those benefiting the most from pension savings, eating healthier food, etc., or those who benefit less (de Ridder et al., 2021; Szasz et al. 2018).

There are two extreme and opposing interpretations of the empirical literature on nudges. In one, individuals lack volition in general, so the tiniest nudge may cause individuals to choose something other than the option that they actually prefer. In the other, individuals have no reason to prefer one option over another in a particular setting, and therefore simply pick the default (cf. Bruns and Perino 2021; de Ridder et al., 2021). Comparing people with different intrinsic motivations to select a particular option while holding the choice situation constant is one way to find out which interpretation is most accurate. The nudge could be tested on two groups, each with its own control, and each with different intrinsic incentives to choose the nudge option. If both nudged groups differ from their controls, volition is weak. If the intrinsically motivated group makes the same choice with or without a nudge, volition is strong. During the rollout of the Covid-19 vaccine in Sweden, the Region of Uppsala was unique among Sweden's 21 regions in that its authorities summoned a large share of the region's inhabitants to pre-booked

appointments, rather than simply opening up telephone lines and web sites for people to make arrangements on their own. This is a nudge because anyone who receives a pre-booked appointment could ignore it and book an appointment (or not) as in the other 20 regions, where people were sent letters without a pre-booked appointment, but telling them how to make their own appointments (or appointments for their dependants), and stating that the vaccines were safe and effective and necessary to avoid the risk of severe disease or death. In this way, the authorities changed the default option to a suggested vaccination time compared to no suggested time.

For our purposes, the crucial aspect of these pre-booked appointments is that they were sent to individuals aged ≥ 50 years and to individuals between the ages of 16 and 17 (except for immunocompromised individuals and hospital staff, pre-booked appointments were not made for those aged between 18 and 49). Given the vast and well-known age gradient in the risk of death or severe disease from Covid-19 and the established connection between risk perception and protection behavior (Brewer et al. 2007), the pre-booked vaccinations in Uppsala provide an excellent opportunity to study the effectiveness of default nudges in age groups that can be expected to be affected differently by the nudge. Our main empirical strategy is to compare the Region of Uppsala and a synthetic control region (i.e. a weighted average of other Swedish regions constructed to mimic Uppsala; described in Section 3.2 below). We analyze and compare the impact of the vaccine rollout on vaccination shares for two groups: individuals aged 50–59, who are intrinsically motivated to get vaccinated due to the substantial risk, and those aged 16–17, who are not intrinsically motivated to get vaccinated. Our findings suggest that the pre-booked vaccinations had a small effect, if any, on the older group, and a considerable impact on the younger one. The difference in effect size is significant at the 10% level.

Received: Month 00, 2023. Revised: Month 00, 2023. Accepted: Month 00, 2023

© The Author(s) 2023. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Our results are consistent with the theory that nudges work best when the nudged individual has little reason to prefer one option over another. Some theoretical analyses incorporate intrinsic motivation into decision problems where the agent is sometimes nudgeable and sometimes not. Löfgren & Nordblom (2020) develop a model in which individuals who face an unimportant decision are more likely to take it inattentively and so are more nudgeable the less important they perceive the decision to be. A contrasting approach is taken by de Ridder *et al.* (2021), who piece together general propositions from empirical regularities. Like Löfgren and Nordblom, they stress the personal preferences of the nudged, arguing also that one cannot be nudged into something one does not want, and that nudges are likely most effective for those who are indifferent. In contrast to Löfgren and Nordblom, they place less emphasis on the relevance of nudges to 'System 1' cognitive processes, i.e. the fast and unconscious response of the brain (Kahneman 2011), arguing that deliberation ('System 2') might lower nudgeability, but not make it impossible. Our findings are consistent with Löfgren & Nordblom (2020) and de Ridder *et al.* (2021).

Another perspective is provided by Allcott & Kessler (2019), who develop a model that allows moral prices and moral utility (e.g. others' expectations) to influence decisions, in addition to inattention. These aspects may be related to perceptions of risk and age to the extent that individuals believe that getting vaccinated will stop transmission of the virus. Thus, there may be little personal gain but much pressure from others to get vaccinated for the young, and more personal gain for the old. Unfortunately, we cannot assess the relative importance of attention, importance, moral prices and moral utility in our data, and we must therefore remain silent on the exact reason for the lack of nudgeability of the 50- to 59-year-olds compared with younger persons.

Several findings of null effects of nudges can be explained with reference to incentives, skills or expertise of the nudged. In this vein, Löfgren *et al.* (2012) find no effect of pre-set defaults on decisions to offset carbon emissions for air travel among environmental economists at a conference (i.e. people with expertise on the issue of the nudge). Similarly, in influencing debtors to repay their debts and schoolchildren to apply for higher education, Holzmeister *et al.* (2022) and Ilie *et al.* (2022), respectively, find no effect of nudging messages, plausibly because the incentives are already clear to the recipients. In a study of German Internet users' willingness to contribute funds to environmental causes, Bruns & Perino (2021) ascertain the intrinsic motivation of their subjects by questionnaires and find no effect of nudges on the intrinsically unmotivated. Finally, in Bonander *et al.* (2022) we examine Region Uppsala's vaccination nudge on 16- to 17-year-olds using partly the same data as in the present paper, finding a large effect but noting the little personal concern of the nudged individuals. This paper extends our previous analysis by contrasting these results with a group with high personal concern, arguably allowing us to study the moderating effects of intrinsic motivation.

Despite receiving widespread attention in the literature, we are not aware of any previous attempts to estimate the impact of nudges on two groups with different intrinsic motivations in the field. The study that comes closest to doing this is the retirement-savings study by Madrian & Shea (2001), which finds that the effect of the pre-set default is lower (though still positive) for older persons. Age could imply a differential in intrinsic motivation because older employees are generally closer to retirement. In an extensive literature review, Szasz *et al.* (2018) note how several factors across different settings can moderate the effects

of nudges and highlight the importance of quantifying effect heterogeneity in future studies.

Our setting has at least two advantages for measuring the impact of a nudge and how it differs across groups with different intrinsic motivations. Firstly, the speed with which the vaccine was rolled out meant that the letters with pre-booked appointments were sent to both the quinquagenarians and the teenagers at roughly similar times, on 5 May 2021, for the 50- to 59-year-olds and on 15 July for the 16- to 17-year-olds. This time of the year is off-season for Covid-19 in Sweden and there was reason to believe it to be so from experiences in the previous year, meaning that the two groups made their vaccination decisions in roughly the same pathogenic environment. In early May, there was more virus in circulation than in July, but at the time of the vaccine rollout, there was talk of how the vaccine would require up to 2 weeks' time before it would fully protect the recipients. Thus, allowing some time for the vaccine to take effect, both groups could expect to be effectively vaccinated in the same low-risk environment.

Secondly, the age-stratified mortality and morbidity of Covid-19 create a natural differentiation in intrinsic motivation for the two groups to get vaccinated. While we cannot rule out that differences between young and old other than intrinsic motivation due to health status are behind our findings, the age gradient in health outcomes from a Covid-19 infection is so staggering that it plausibly dwarfs other potential factors. In one estimate, pre-vaccine seroprevalence data indicate a >400-fold difference in the infection fatality rate between persons aged ≤ 19 and those aged 50–59 (Pezzullo *et al.* 2023). Other estimates are less dramatic but nevertheless suggest a 100-fold difference (Ghisolfi *et al.* 2020; O'Driscoll *et al.* 2021). In any event, the risk difference between 50- to 59-year-olds and 16- to 17-year-olds is clear. Thus, we have a clean setting in which nature has differentiated intrinsic motivation, and the authorities have acted with sufficient celerity to ensure decision-makers in our study make their choices in roughly the same environment.

The remainder of the paper is organized as follows. Section 2 describes vaccination practices across Swedish regions. Section 3 presents data and methods used for estimating the effects of the nudge. Section 4 presents the results of the empirical analysis as well as several robustness checks. Section 5 compares the findings with earlier studies and discusses the implications for policies related to default nudges.

Vaccine nudges in Swedish regions

We contacted all 21 regions of Sweden to ask if the regional health authorities deviated from the standard practice of opening up telephone lines and web sites to let people make their own vaccination arrangements. It was not standard practice to send out letters with pre-booked appointments for vaccinations against Covid-19 in Sweden during the vaccine rollout. However, for particular age groups, authorities in a number of regions made special arrangements. The health authorities in the Region of Uppsala were the only ones to send letters with pre-booked appointments to everyone aged 16–17 and 50–59 years. Except for immunocompromised individuals and severely handicapped individuals who rely on assistance for almost everything, everyone up to the age of 50 made their own appointments in all of Sweden, including Uppsala. The health authorities in a total of five other regions made use of nudges in some way, but in every case, the use was much more limited compared with the Region of Uppsala.

The regions that come closest to using nudges in the same way as Uppsala are those of (1) Dalarna, where the health authorities

sent letters with pre-booked appointments to persons aged ≥ 80 , as well as to particular high-risk individuals of all ages, and (2) Västmanland, where individuals aged ≥ 75 were sent letters with pre-booked appointments. In (3) the Region of Jönköping, the health authorities made no use of nudges before 6 September 2021, but then sent letters with pre-booked appointments to 'severely immunocompromised individuals' (their words used in email correspondence with us) for a booster.

No other regions used pre-booked appointments. However, in (4) the Region of Västra Götaland, persons aged ≥ 85 were sent postcards informing them that they would receive a telephone call helping them to book an appointment. In addition, everyone born in 1946 or before (i.e. people at least 75 years old) was contacted by their local health providers so that they could find a suitable time for a vaccination appointment. In this region, immunocompromised individuals of all ages > 18 were also contacted to set vaccination appointments. This is a nudge in the sense that it changes the choice architecture and the contacted individuals could still say no or hang up and then make their own arrangements. In a similar vein, the health authorities of (5) the Region of Jämtland-Härjedalen contacted all 80-year-olds and over by telephone, but otherwise did not use any nudges.

Throughout the pandemic, normal in-person elementary education (ages 7–16) was maintained in all the regions of Sweden, while secondary education (up to age 18 or 19) moved partially online, but increasingly less so at the time that 16- to 17-year-olds became eligible for vaccination. It is possible that the 16- to 17-year-olds inferred from the open schools, as well as from the greater exposure to pathogens that normal life entails, that Covid-19 is not dangerous to them. Thus, there is some reason to believe that the risks of Covid-19 would have been generally understood among this age group. Apart from the pre-booked appointments, the letters that were sent out to inform this age group about the possibility to get vaccinated also did not substantively differ among regions.

Empirical methodology

Data

To estimate the effect on vaccinations of the letters with pre-booked vaccination appointments, we use weekly panel data on the share of vaccinated individuals in the 21 Swedish regions (NUTS3) obtained from the [Public Health Agency of Sweden \(2021\)](#). We have previously used some of this data to study the treatment effect on 16- to 17-year-olds only ([Bonander et al. 2022](#)). Our data contain the share of vaccinated individuals in the age groups 16–17, 18–29, 40–49 and 50–59 from weeks 1 to 46 in 2021. We choose the first week in 2021 as the starting point for our analysis since this includes the first vaccinations in the studied age groups. We choose week 46 (mid-November) as the end date because that was when the Swedish government announced mandatory vaccinations for participation in certain public events. Periods in which fewer than four individuals were vaccinated are treated as zero for reasons of personal integrity. Thus, some data are missing in weeks when very few individuals got vaccinated. However, this is mainly a problem in the early stages of the vaccine rollout, when mostly high-risk individuals (e.g. chronic lung disease, cancer and diabetes) were vaccinated. For descriptive purposes, we also have data on the share vaccinated in treated and nearby municipalities in week 49 (early December).

We include several covariates that may be relevant confounders following recent empirical findings on relevant predictors of attitudes towards COVID-19 vaccination ([Bonander](#)

[et al. 2022](#)): the share of foreign-born individuals living in the region and the share of the population with at least 3 years of higher education in 2020 ([Statistics Sweden 2021](#)); the share of COVID-19 deaths in 2020 ([Public Health Agency of Sweden 2021](#)); the share of the population that has received financial aid at any point during the past year in 2020; the number of adolescents aged 10–24 per 100 000 inhabitants who have received care due to alcohol addiction in 2019; the share of the population that has access to a fast broadband connection (100 Mbit/s) in 2020; the share of the population with high trust in how the Swedish health care system handled the COVID-19 pandemic in 2020; and NEETs (the share of 17- to 24-year-olds that neither study nor work) in 2020 ([Kolada 2021](#)). These variables do not necessarily have a direct bearing on attitudes towards vaccinations, but they are plausible measures of general satisfaction with authorities including health authorities, and attitudes to health more broadly. In sensitivity analyses, we also consider specifications without covariates.

The synthetic control method

We use the synthetic control method to construct a synthetic control that closely resembles Uppsala in terms of pre-intervention characteristics, from a weighted combination of all other Swedish regions ([Abadie et al. 2010](#); [Abadie 2021](#)). The synthetic control method is designed to estimate the impact of policy interventions affecting one unit (e.g. country, region or municipality) when a small number of control units are available. It is a data-driven approach for estimating counterfactuals—i.e. what would have happened without the nudge—which determines the weighted combination of untreated regions that provides the closest match to the treated region with regard to pre-intervention outcomes and covariates. The weighted average vaccination uptake from the synthetic control group then provides the counterfactual trend of the vaccination share for the treated region; i.e. it predicts how the vaccination shares would have turned out in the absence of the nudging intervention. The method offers greater robustness to non-parallel trends bias compared with difference-in-differences ([O'Neill et al. 2016](#)). [Abadie et al. \(2010\)](#) motivate the method assuming a factor model in which the observed outcomes are driven by both observed and unobserved factors. The key idea is that—absent structural breaks and concurrent interventions—if the synthetic control mirrors the treated region's pre-treatment outcomes, it can also forecast the post-treatment outcomes without the intervention. For a detailed presentation of the method, see [Abadie \(2021\)](#). Following [Abadie & Gardeazabal \(2003\)](#), we employ a nested optimization routine to determine the optimal set of unit weights (\mathbf{W}), which determine each untreated unit's contribution to a synthetic control, and variable importance weights (\mathbf{V}), which determine the likelihood of a good match on strong predictor outcomes.

Specifically, let $J + 1$ be the number of regions indexed by j , and let $j = 1$ denote Uppsala. The regions in the sample are observed for periods $t = 1, 2, \dots, T$, where T_0 is the number of pre-treatment periods. Next, we define two potential outcomes: Y_{jt}^I is the potential outcome if region j is exposed to the nudge intervention in time t , and Y_{jt}^N is the corresponding potential outcome without the nudge intervention, which is unobserved in Uppsala after T_0 . The goal of the analysis is to measure the post-treatment effect in Uppsala (region $j = 1$), defined as $\alpha_{1t} = Y_{1t}^I - Y_{1t}^N$. Since Y_{1t}^N is unobserved, we use the synthetic control method to impute it. The synthetic region is constructed as a weighted average of the control regions $j = 2, \dots, J + 1$ from the donor pool of control regions and represented by a vector of weights $\mathbf{W} = (w_2, \dots, w_{J+1})'$

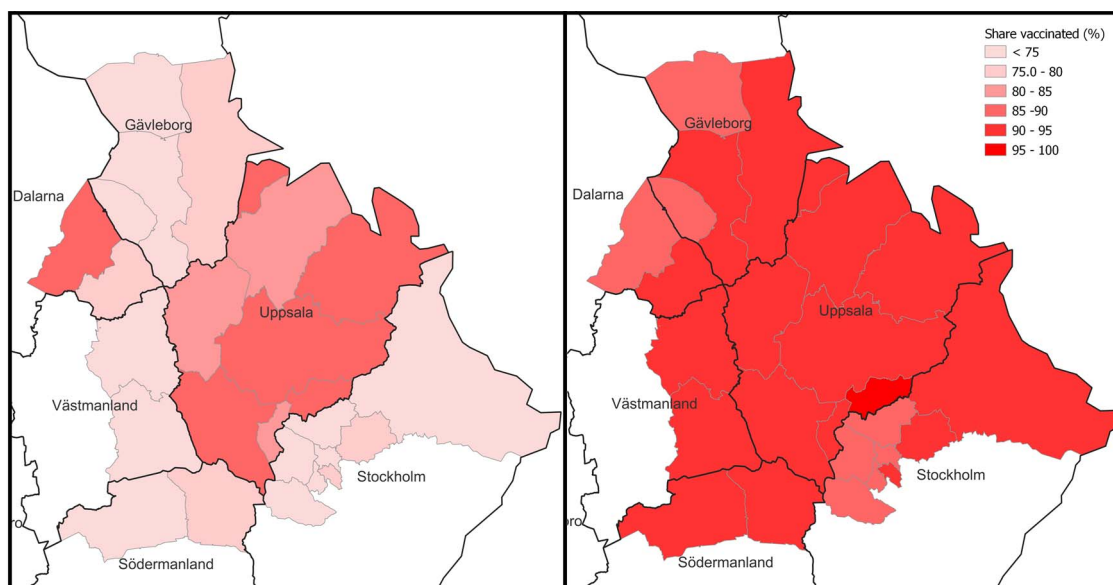


Figure 1. Share of vaccinated 16- to 17-year-olds (left) and 50- to 59-year-olds (right) in treated municipalities (Uppsala) and nearby (non-treated) municipalities, week 49

with $0 \leq w_j \leq 1$ and $w_2 + \dots + w_{j+1} = 1$. Each choice of \mathbf{W} gives a particular set of weights and characterizes a possible synthetic control. For each age group, we choose the \mathbf{W} that minimizes the pre-intervention mean squared error between Uppsala and its synthetic control on the eight predictors of the outcome variable presented in the data description and the outcome variable itself. The predictors are recorded once at the end of the year 2020, or 2019 (depending on data availability). Since there are zero vaccinated 16- to 17-year-olds in the weeks before intervention we match on the share of vaccinated 18- to 29-year-olds in the analyses of the younger age group. In the main analyses we use the average vaccination shares in all the included weeks before the intervention to construct the synthetic control, but we test 14 different ways of including lagged values of the outcome variable, as suggested by Ferman *et al.* (2020). To assess if the age-specific estimates differ from each other, we also estimate the difference in effect between the two age groups in the treated region by week relative to the start of the vaccinations.

Results

Descriptive municipal comparison

As a first step, we compare the vaccination shares in treated municipalities to neighboring municipalities. Figure 1 shows vaccination shares for 16- to 17-year-olds (left) and 50- to 59-year-olds (right) in treated municipalities (in region Uppsala) and nearby (non-treated) municipalities. The difference in vaccination shares between treated and untreated municipalities is bigger for the younger than for the older age group. A comparison between the eight treated municipalities and their eight bordering neighbors shows that the vaccination share was 12.9 percentage points (95% CI: 9.0, 16.8) for the younger age group, and 3.6 percentage points for the older age group (95% CI: 1.4, 5.9). This descriptive analysis indicates that there may be an effect of the vaccination nudge, especially in the younger age group.

Regional analyses

Figure 2 shows the trajectories of the share of vaccinated 16- to 17-year-olds (left panel) and 50- to 59-year-olds (right panel) in

the treated region (Uppsala) and the average of all other Swedish regions during the first 46 weeks of the year 2021. The vaccination shares follow similar paths at the beginning of the year. However, the trajectories diverge after week 30 (late July) for the 16- to 17-year-olds and from week 16 (mid- to late April) for the 50- to 59-year-olds.

Figure 3 shows that, prior to treatment, the trajectories of vaccinations in the treated region and its synthetic counterpart are very similar in both age groups. Furthermore, Table 1 compares the values of key predictors for the treated region before treatment, with the same values for the synthetic region and the average of all 20 control regions in the time periods before treatment. For all predictors, except share with financial aid, the treated region and its synthetic counterpart are more similar than the treated region and the unweighted average of all regions. The final column of Table 1 shows how the predictors are weighted. The \mathbf{W} weights that explain which (and to what extent) regions have been used to construct the synthetic control region were Stockholm (0.133), Östergötland (0.539), Kronoberg (0.224) and Västerbotten (0.014) for the younger age group, and Stockholm (0.199), Östergötland (0.301), Jönköping (0.119), Västra Götaland (0.232) and Jämtland (0.148) for the older age group.

Because the mean squared prediction error (MSPE) in the pre-intervention period is close (or equal) to zero in most panels, we use the post-intervention MSPE to rank the effect sizes across regions. The post-treatment differences between Uppsala and the synthetic controls in Fig. 3 show the increases in vaccination shares for 16- to 17-year-olds (left) and 50- to 59-year-olds (right). The letters with pre-booked appointments seem to have increased vaccinations in the younger age group, while the increase is less pronounced in the older age group. In the last week of the sample period, week 46 (mid-November), the vaccination share was 11.7 (84.5 vs 72.8) and 3.6 (93.3 vs 89.6) percentage points higher in the treated region than they would have been in the absence of treatment. The results are very similar if we exclude the five control regions with nudge-like interventions directed at much older age groups as described in Section 2 (see Fig. A1 in the Appendix).




Figure 2. Share of vaccinated 16- to 17-year-olds (left) and 50- to 59-year-olds (right) in the treated region (Uppsala) versus the average of all other regions, by week before and after intervention (vertical line)

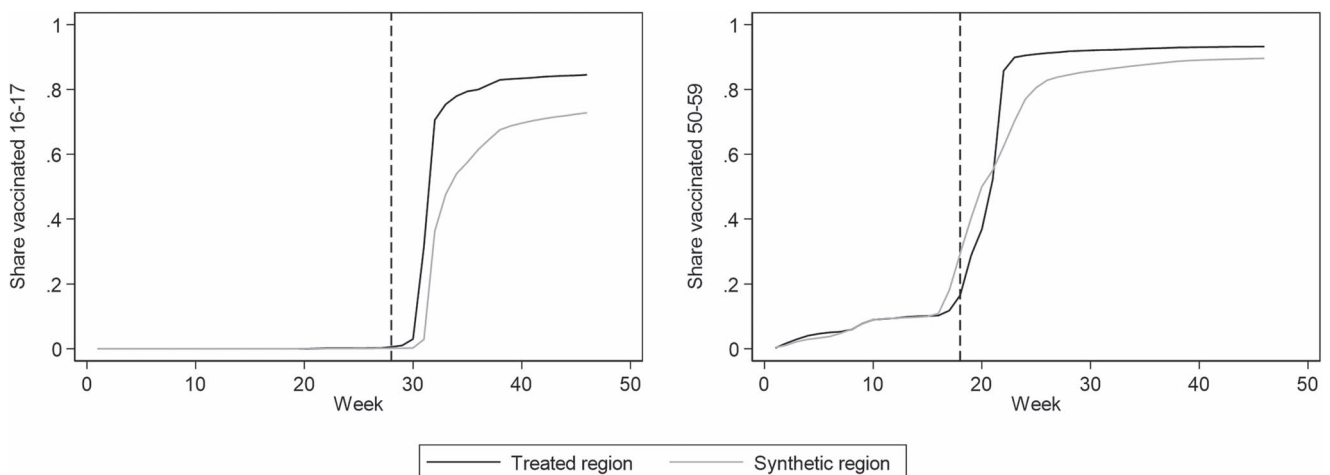


Figure 3. Path plot of the share of vaccinated 16- to 17-year-olds (left) and 50- to 59-year-olds (right): Treated region (Uppsala) versus the synthetic region

Robustness tests

We perform a series of tests to assess the robustness of our results. First, we conduct an in-space placebo test and a leave-one-out analysis, as suggested by [Abadie et al. \(2010\)](#) and [Abadie \(2021\)](#). More importantly, synthetic control analyses can be sensitive to how one chooses to input the pre-intervention outcomes into the optimization procedure. Following [Ferman et al. \(2020\)](#), we also estimate 14 different specifications where we vary how we input the pre-intervention outcomes in the synthetic control optimization method.

For the in-space placebo tests, we compare the effect estimated for the treated region with the effect of placebo interventions iteratively implemented in the non-treated 20 regions. This shows us if the results obtained for the treated region are large compared with the placebo results in the other regions, and it allows for measuring the fraction of regions with results as big as those of the treated region (placebo-based P -values). The results are presented in [Fig. 4](#). For the 16- to 17-year-olds (left panel), no other region has an effect estimate close to the one in the treated region; a ranking of the post-intervention effect sizes across all regions shows that Uppsala has by far the largest estimated effect, with a placebo-based P -value of $1/21 = 0.048$. Overall, the analysis indicates a large and persistent

effect of the intervention on 16- to 17-year-olds. For the 50- to 59-year-olds (right panel), the effect estimate in the treated region is not that different from the placebo effect estimates. A ranking of the effect sizes in this age group shows that the treated region does not clearly stand out compared with the untreated regions, with a placebo-based P -value of $3/21 = 0.143$ ([Fig. A2](#) in the Appendix). To assess if the two age-specific effect estimates are different from each other, we also perform an analysis where we compute the difference in effect estimates between 16- to 17-year-olds and 50- to 59-year-olds by week relative to the start of vaccinations in each age group. Again, we perform the analysis within Uppsala and compare the results with all other regions for placebo-based inference. We find that the difference in effect estimates between the two groups is among the largest in the Uppsala region, with a placebo-based P -value of $2/21 = 0.095$ ([Fig. A3](#) in the Appendix).

We conduct a leave-one-region-out robustness analysis, as suggested by [Abadie \(2021\)](#). One at a time, we take out each of the four/five regions that contribute to the synthetic control to ensure that the results are not entirely driven by the inclusion of a specific control region. All leave-one-out estimates closely track the findings from the main analysis. Thus, the main findings appear robust ([Fig. A4](#) in the Appendix).

Table 1. Vaccination share predictor means

	Uppsala	Synthetic Uppsala	Average of 20 control regions	V
Panel A: 16–17 y				
Share foreign-born	0.189	0.182	0.160	0.361
Share high education	0.180	0.157	0.138	0.033
Share of COVID-19 deaths	0.001	0.001	0.001	0.010
Share with financial aid	0.041	0.042	0.040	0.090
Alcohol addiction per 100 000	54.0	69.6	89.1	0.044
Share with fast internet	0.862	0.868	0.836	0.265
Share with high trust	0.752	0.750	0.737	0.076
Share NEETs	0.062	0.070	0.078	0.085
Share vaccinated (18–29 y)	0.077	0.075	0.074	0.035
Panel B: 50–59 y				
Share foreign-born	0.189	0.185	0.160	0.037
Share high education	0.180	0.161	0.138	0.014
Share of COVID-19 deaths	0.001	0.001	0.001	0.075
Share with financial aid	0.041	0.036	0.040	0.029
Alcohol addiction per 100 000	54.0	72.2	89.1	0.030
Share with fast internet	0.862	0.867	0.836	0.052
Share with high trust	0.752	0.746	0.737	0.072
Share NEETs	0.062	0.072	0.078	0.033
Share vaccinated (50–59 y)	0.069	0.070	0.077	0.658

Notes: The period for each predictor is 2020, except for Share vaccinated, which refers to the mean share for all pre-intervention weeks. V is the variable importance weights.

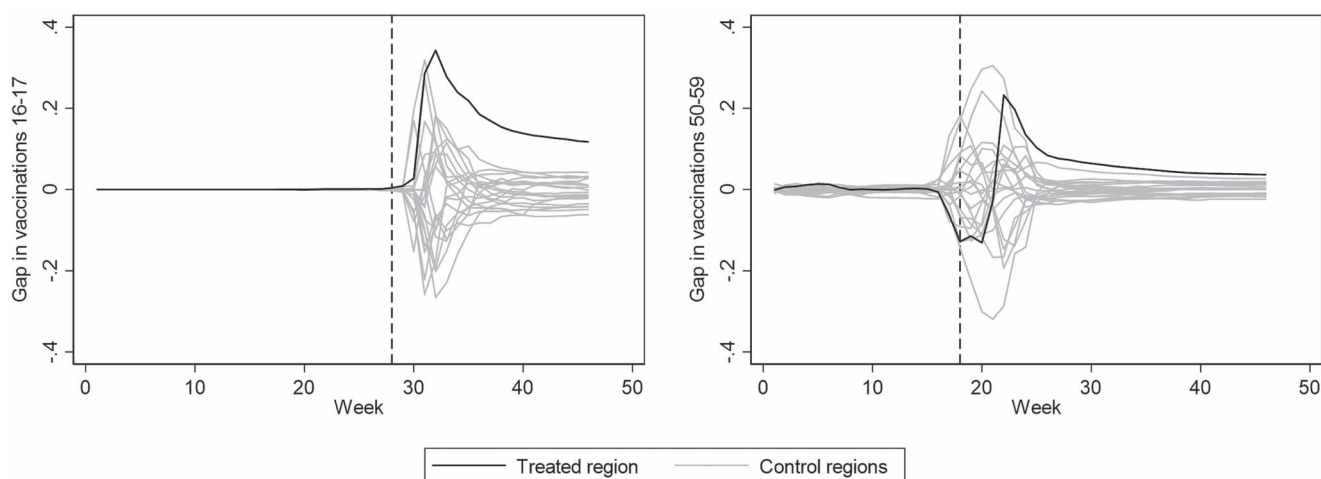


Figure 4. Permutation tests: Vaccination share gap in the treated region and placebo gaps for the control regions, 16- to 17-year-olds (left), 50- to 59-year-olds (right)

Following Ferman *et al.* (2020), we also estimate 14 different specifications in which we vary how we input the pre-intervention outcomes in the synthetic control optimization procedure and find that the effect size is largest in the treated region in all 14 specifications for 16- to 17-year-olds. For 50- to 59-year-olds, we find that the effect size is largest in the treated region in only one of the specifications. This implies that the effect estimate for the younger age group is robust to specification searching. However, for the older age group, the effect estimate is not robust (Table A2 in the Appendix).

Finally, we assess the vaccination share for other age groups. Since 18- to 29 and 40- to 49-year-olds in Uppsala were not sent letters with pre-booked appointments, we do not expect them to have a higher vaccination share. However, there may be spillovers from the treatment; increased vaccinations in the treated age group may increase vaccinations among friends and relatives in the untreated age group. From the effect plots, we see an

indication of an effect for both these age groups, but placebo estimations show that those effects are not only small but also statistically insignificant (Figs. A5 and A6 in the Appendix). In Uppsala, older age groups were sent letters with pre-booked appointments a few weeks before the 50- to 59-year-olds; the effects in these groups are even less pronounced than in the 50- to 59-year age group (Figs. A7-A9 in the Appendix).

Discussion and conclusion

Consistent with the theory that nudges are at their most effective the less the nudged individual cares about the decision, our results suggest that persons at low risk of adverse events due to Covid-19 are easily nudged into taking the vaccine, but those at higher risk are not. Our main innovation is to use two groups, each subjected to the same intervention and each with its own control group.

Our setting is well suited for measuring the impact of the nudge because the vaccine was rolled out quickly and the letters with pre-booked appointments were sent to both the quinquagenarians and the teenagers at roughly similar times between May and July 2021, a time of the year that is off-season for Covid-19 in Sweden, so the two groups made their vaccination decisions in a similar pathogenic environment. In addition, the age-stratified mortality and morbidity of Covid-19 create a natural differentiation in intrinsic motivation for the two groups to get vaccinated. Thus, natural factors and the auspicious timing of the vaccine rollout combine to create a clean setting for the impact of nudges to be estimated. In addition to shedding light on the effects of nudges more generally, our findings directly relate to the empirical literature on the effects of nudges and incentives to increase Covid-19 vaccinations (Barber & West, 2021; Bonander et al. 2022; Dai et al. 2021; Campos-Mercade et al. 2021).

Our study also has some noteworthy limitations. Firstly, our design is quasi-experimental, and while we believe our estimates are reasonably credible, a randomized trial would have been able to provide stronger causal evidence. Secondly, our interpretation of the findings relies on the assumption that pre-booked vaccinations influence behavior via nudging. However, there is at least one possible confounding factor unrelated to nudging that we cannot eliminate: the impact of who makes the bookings. If individuals make their own arrangements, popular times may fill up quickly and people may be reluctant to choose less suitable times. While time off work for Covid-19 vaccinations is widely accepted in Sweden, it may still look better if a worker leaves in the middle of the day because he has been called to a vaccination appointment rather than leaving because he has chosen it. Therefore, the very small effect we estimate among 50- to 59-year-olds may be due to a coordination factor unrelated to nudging, if any effect is present at all.

Nudges are known to have highly variable and context-specific effects (Szasz et al. 2018; Löfgren and Nordblom 2020; de Ridder et al. 2022). In light of the stark contrast in the effect of default nudges on vaccine take-up that we find, we believe that nudge theory would benefit from more efforts to highlight the effect moderation in heterogeneous groups of nudged individuals, to ascertain when nudges are powerful influences on action and when they are not.

Supplementary data

Supplementary data is available at *Oxford Open Economics* online.

Acknowledgements

We thank Siri Jakobsson Støre, Andreas Kotsadam, Katarina Nordblom and participants from various seminars and conferences for providing valuable comments and suggestions. Declarations of interest: None. We acknowledge funding from Jan Wallanders och Tom Hedelius stiftelse samt Tore Browaldhs stiftelse (grant number P22-0018), as well as from the Swedish Research Council for Health, Working Life, and Welfare (Forte; grant number 2020-00962).

References

- Abadie, A. (2021) 'Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects', *Journal of Economic Literature*, **59**: 391–425.
- Abadie, A., and Gardeazabal, J. (2003) 'The Economic Costs of Conflict: a Case Study of the Basque Country', *American Economic Review*, **93**: 113–32.
- Abadie, A., Diamond, A., and Hainmueller, J. (2010) 'Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program', *Journal of the American Statistical Association*, **105**: 493–505.
- Allcott, H., and Kessler, J. B. (2019) 'The Welfare Effects of Nudges', *American Economic Journal: Applied Economics*, **11**: 236–76.
- Atalay, K. et al. (2014) 'Savings and Prize-Linked Savings Accounts', *Journal of Economic Behavior & Organization*, **107**: 86–106.
- Bonander, C., Ekman, M., and Jakobsson, N. (2022) 'Vaccination Nudges: a Study of Pre-Booked COVID-19 Vaccinations in Sweden', *Social Science & Medicine*, **309**: 115248.
- Brewer, N. T. et al. (2007) 'Meta-Analysis of the Relationship between Risk Perception and Health Behavior: the Example of Vaccination', *Health Psychology*, **26**: 136–45.
- Bruns, H., and Perino, G. (2021) 'Point at, Nudge, or Push Private Provision of a Public Good', *Economic Inquiry*, **59**: 996–1007.
- Campos-Mercade, P. et al. (2021) 'Monetary Incentives Increase COVID-19 Vaccinations', *Science*, **374**: 879–82.
- Cheung, T. T. L. et al. (2019) 'Cueing Healthier Alternatives for Take-Away: a Field Experiment on the Effects of (Disclosing) Three Nudges on Food Choices', *BMC Public Health*, **19**: 974.
- Dai, H. et al. (2021) 'Behavioural Nudges Increase COVID-19 Vaccinations', *Nature*, **597**: 404–9.
- Ferman, B., Pinto, C., and Possebom, V. (2020) 'Cherry Picking with Synthetic Controls', *Journal of Policy Analysis and Management*, **39**: 510–32.
- Ghisolfi, S. et al. (2020) 'Predicted COVID-19 Fatality Rates Based on Age, Sex, Comorbidities and Health System Capacity', *BMJ Global Health*, **5**: 1–8.
- Holzmeister, F. et al. (2022) 'Nudging Debtors to Pay their Debt: Two Randomized Controlled Trials', *Journal of Economic Behavior & Organization*, **198**: 535–51.
- Hoover, H. (2022) 'Nudges as Norms: Evidence from the NYC Taxi Cab Industry', *Journal of Economic Psychology*, **92**: 102535.
- Ilie, S. et al. (2022) 'No Budge for any Nudge: Information Provision and Higher Education Application Outcomes', *Education Sciences*, **12**: 701.
- Kahneman, D. (2011) *Thinking, Fast and Slow*, New York, NY: Farrar, Straus and Giroux.
- Kolada. 2021. www.kolada.se.
- Li, D., Hawley, Z., and Schnier, K. (2013) 'Increasing Organ Donation Via Changes in the Default Choice or Allocation Rule', *Journal of Health Economics*, **32**: 1117–29.
- Löfgren, Å., and Nordblom, K. (2020) 'A Theoretical Framework of Decision Making Explaining the Mechanisms of Nudging', *Journal of Economic Behavior & Organization*, **174**: 1–12.
- Löfgren, Å. et al. (2012) 'Are Experienced People Affected by a Pre-Set Default Option – Results from a Field Experiment', *Journal of Environmental Economics and Management*, **63**: 66–72.
- Madrian, B., and Shea, D. F. (2001) 'The Power of Suggestion: Inertia in 401 (K) Participation and Savings Behavior', *The Quarterly Journal of Economics*, **116**: 1149–87.
- O'Driscoll, M. et al. (2021) 'Age-Specific Mortality and Immunity Patterns of SARS-CoV-2', *Nature*, **590**: 140–5.
- O'Neill, S. et al. (2016) 'Estimating Causal Effects: Considering Three Alternatives to Difference-in-Difference Estimation', *Health Services and Outcomes Research Methodology*, **16**: 1–21.
- Pezzullo, A. et al. (2023) 'Age-Stratified Infection Fatality Rate of COVID-19 in the Non-Elderly Population', *Environmental Research*, **216**: 114655.

- Public Health Agency of Sweden (2021) Vaccination Statistics Data delivered in November and December 2021. Public Health Agency of Sweden: Stockholm, Sweden.
- de Ridder, D., Kroese, F., and van Gestel, L. (2022) 'Nudgeability: Mapping Conditions of Susceptibility to Nudge Influence', *Perspectives on Psychological Science*, **17**: 346–59.
- Statistics Sweden. 2021. *Educational Attainment of the Population*. Örebro, Sweden.: Statistics Sweden. www.scb.se/hitta-statistik/statistik-efter-amne/utbildning-och-forskning/befolkningens-utbildning/befolkningens-utbildning/
- Szaszi, B. et al. (2018) 'A Systematic Scoping Review of the Choice Architecture Movement: Toward Understanding when and why Nudges Work', *Journal of Behavioral Decision Making*, **31**: 355–66.