Journal of Health Economics 74 (2020) 102386

Contents lists available at ScienceDirect

Journal of Health Economics

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

It runs in the family – Influenza vaccination and spillover effects*

Nicolas Bouckaert [a], Anne C. Gielen [b,c], Tom Van Ourti [d,*]

[a] Belgian Health Care Knowledge Centre, Brussels, Belgium; and KU Leuven, Leuven, Belgium
[b] Erasmus School of Economics and Tinbergen Institute, Erasmus University Rotterdam, The Netherlands
[c] IZA, Germany
[d] Erasmus School of Health Policy and Management, Erasmus School of Economics, Tinbergen Institute and Erasmus Centre for Health Economics Rotterdam, Erasmus University Rotterdam, The Netherlands

ARTICLE INFO

Article history:
Received 7 June 2019
Received in revised form
23 September 2020
Accepted 25 September 2020
Available online 18 October 2020

Keywords:
Influenza vaccination
Family spillover effects
Regression discontinuity

ABSTRACT

We study a population-based influenza vaccination program in the Netherlands, and the spillovers it has within families. Individuals aged 65 years and over qualify for the program and receive a personal invitation for a free flu shot, while ineligible individuals have to pay out-of-pocket and face additional barriers to getting vaccinated. The quasi-random variation at age 65 is exploited to analyse program impact on vaccination behavior of cohabiting partners and adult children. We find that the program induced a 10 percentage points increase in vaccination coverage among individuals at age 65. The program further led to a similar effect on vaccination take-up by cohabiting younger partners, but spillovers on children were negative. These asymmetric patterns of vaccination uptake are consistent with partners and children learning about influenza mortality risk, target group membership, and cost and benefits of vaccination, as well as salience. We conclude that public health campaigns should pay attention to the effects on voluntary preventive care participation as within-family spillovers impact the program's overall public health impact.

* This project has used data provided by Statistics Netherlands via a remote access facility. As stipulated in the data agreement, Statistics Netherlands pre-viewed the findings of this project prior to publication to ensure that privacy sensitive, individual-specific information was not revealed. The data from this study can only be applied for through a government data sharing portal of Statistics Netherlands (https://www.cbs.nl/en-gb/our-services/customised-services-microdata/microdata-conducting-your-own-research). Part of the work was undertaken while Nicolas Bouckaert was a PhD fellow of the FWO Flanders (1174913N); Anne C. Gielen a Marie Curie Intra-European Fellow (PIEF-GA-2011-299133); and Tom Van Ourti a visiting researcher at the Milken Institute School of Public Health of the George Washington University. The authors would like to thank the editor Mathias Kifmann, two anonymous reviewers, Philippe Beutels, Adrian Bruhin, Geert Dhaene, Lorens Helmchen, Hale Koç, Jürgen Maurer, Ali Moghtaderi, Magne Mogstad, Owen O'Donnell, Erwin Ooghe, Francesco Principe, Samantha Rawlings, Erik Schokkaert, Erdal Tekin, Joost Timmermans, Hans Van Brabandt, Ellen Van de Poel, Hans van Kippersluis, GP private practice Van Eerd; and seminar participants at several universities and conferences for useful comments and suggestions. We have no conflicts of interest to disclose, and all errors are our own.

* Corresponding author.
E-mail addresses: nicolas.bouckaert@kce.fgov.be (N. Bouckaert), gielen@ese.eur.nl (A.C. Gielen), vanourti@ese.eur.nl (T. Van Ourti).

https://doi.org/10.1016/j.jhealeco.2020.102386

1. Introduction

Like with many other preventive care measures for communicable diseases, vaccination coverage for seasonal influenza is considered to be too low, and encouraging influenza vaccination uptake is a key public health strategy in many countries. Although studies have shown that patient reminders, financial incentives and information provision about the costs and benefits of influenza vaccination are effective in increasing vaccination uptake among targeted groups in the population (Szilagyi et al., 2000; Bronchetti et al., 2015; Nuscheler and Roeder, 2016), little is known about the extent to which these measures impact vaccination uptake by individuals that were not targeted. If vaccination behavior of other individuals is also affected, this has important health and economic implications such as the costs associated with work productivity, absenteeism, and health care. Hence, credible identification of vaccination spillover effects is vital to determine overall program (cost-)effectiveness of preventive care policies for infectious diseases.

This paper studies spillover effects of a Dutch influenza vaccination program. Every year, between 2 and 20 percent of the Dutch population is affected by influenza (Vrieze et al., 2016). Although most people recover within 1–2 weeks from influenza without requiring medical attention, influenza can cause severe illness and even death among vulnerable people, including the elderly, pregnant women and people with an underlying health condition (WHO, 2019). Vaccination is the leading preventive strategy for reducing infection risk, is protective after about 2 weeks, and needs to be taken yearly. The Netherlands has a vaccination policy in place aimed at directly protecting individuals who are most at risk of influenza mortality.[1] Chronically ill individuals with specific disorders and individuals aged 65 (60 as of 2008) or above on May 1st of the next calendar year receive a personal invitation letter in September/October of the current calendar year to collect and obtain their flu shot from the GP at zero cost. Individuals not targeted by the program can also get vaccinated, but they do not receive an invitation, have to pay out-of-pocket costs, and first have to collect the vaccine from the pharmacy before going to their GP to have the actual flu shot.

We exploit the quasi-random variation at age 65 to identify within-family spillovers in vaccination behavior, both within the household (cohabiting partners) and between households (adult children). The household is an important transmitter in the spread of influenza as within-household influenza infection has been shown to contribute more to influenza incidence than any other source of infection (Welliver et al., 2001; Ferguson et al., 2006), while adult children are affected by other important infection sources outside the immediate household such as indirect exposure via their own children attending childcare or school, direct exposure via workplace environments, or via the local neighborhood (Longini et al., 2004, 2005; Chao et al., 2010). Vaccination uptake among groups thus has the potential of strongly containing the spread of influenza, especially in a small and densely populated country, such as the Netherlands, where parents and adult children live in close proximity.

Using a regression discontinuity design (RDD) and 11 years of cross-sectional Health Interview Surveys (HIS) for influenza seasons 1997–98 to 2007–08 which is linked to Dutch administrative data from Statistics Netherlands, we find that spillovers in the family context are almost as important as the direct impact on vaccination uptake of individuals upon turning 65. We find that vaccination incidence among partners increases with 10 percentage points, from 25 to 35 percent, but only when the partner is younger than 65 and not yet qualifies for the vaccination program based on age. This spillover effect is as large as the direct program impact which increases vaccination coverage from 30 to 40 percent when individuals turn 65. We further find evidence for substantial spillovers on adult children, but only upon their older parent turning 65. Interestingly, the child spillovers are negative, and show a reduction in vaccination coverage among the children from 9 to 4 percent. Our results suggest that these asymmetric spillover effects towards partners and children are in line with them updating their influenza mortality risk, partners learning about costs and benefits of influenza vaccination, and children updating their beliefs about target group membership, but also with the salience of the partner's/paternal receipt of the invitation for free influenza vaccination. While our sample size precludes credible identification of the health (care) impacts of these spillovers, a companion paper indicates that the health (care) impacts might be substantial. Van Ourti and Bouckaert (2020) find that the Dutch vaccination program reduced GP visits, prescription drug use and mortality among individuals that cross the 65 age threshold. This suggests potential program-associated health (care) effects of the spillovers on partners and children, although the overall program effect ultimately depends on the positive health (care) impacts among partners outweighing the negative impact on children. Various sensitivity, specification and placebo tests underscore the robustnesst of our results.

This paper adds to the literature on influenza vaccination in high income countries by focusing on spillovers in individual vaccination behavior within the family context, and by providing suggestive evidence on the underlying mechanisms. As far as we know, this has not been studied before in the economics literature, but there are two related studies estimating health externality benefits arising from influenza vaccination.[2] Ward (2014) documents

[1] Most countries in Europe target the old and sick, but in other countries, including the United States and the United Kingdom, influenza vaccination is also recommended for younger age groups (except newborns) with the aim to prevent disease transmission and indirectly protect individuals at risk (CDC, 2010; MacDonald, 2016).

[2] Carpenter and Lawler (2019) find that US state mandates for a tetanus, diphteria and pertussis (Tdap) booster prior to middle school entry increased vaccination rates for meningococcal disease and human papillomavirus (HPV) among the same adolescents, but did not affect influenza vaccination rates. Hoffmann et al. (2019) find that peer influenza vaccination take-up increases individual take-up in a field experiment in a major bank in Ecuador.

substantial health improvements among already targeted and vaccinating individuals aged 65 and older in the Canadian province of Ontario, when free vaccination coverage got extended to all residents in the province. White (2020) confirms these health externality benefits for the United States, in particular when increased vaccination rates result from vaccination mandates for health care workers.

Our work also relates to the literature on spillovers in family health behaviors. Cutler and Glaeser (2010) and Fletcher and Marksteiner (2017) show that exposure to respectively workplace smoking bans and clinical interventions reduces spousal smoking and drinking behavior, while Cawley et al. (2019) find no evidence in favor of genetic nurturing driving the positive association between sibling's body mass indices. Fadlon and Nielsen (2019) show that spouses and adult children increase statins consumption when their spouse/parent experiences a non-fatal, but unexpected heart attack or stroke; and draw attention to the underlying roles of learning about health risks and salience of the health shock. We extend this literature to behaviors preventing infectious diseases. Our work is also related to the broader literature about spillovers in the health domain, but avoids the potential concern of artificial social groups that arises in studies using random group assignment to overcome end ogenous social group formation (Duncan et al., 2005; Kremer and Levy, 2008; Carrell et al., 2011; Yakusheva et al., 2011; Golberstein et al., 2016; Bruhin et al., 2020).

The remainder of the paper proceeds as follows. Section 2 provides more background of the Dutch influenza vaccination program, and Section 3 introduces the data. We provide more background of the age-based RDD design and the exact empirical specification in Section 4. Sections 5 and 6 present our findings; and Section 7 discusses potential mechanisms underlying the asymmetric spillovers between partners and towards adult children. The final section offers some concluding remarks.

2. Dutch free influenza vaccination program

2.1. Dutch influenza vaccinations before 1996

In the Netherlands, increased influenza activity is typically recorded between mid-November and early April, and flu shots are usually administered between October and December. The Dutch Health council and health authorities identify high-risk groups who are targeted for influenza vaccination.

In the eighties, high-risk groups were defined based exclusively on existing chronic disorders, such as diabetes, cardiovascular and pulmonary conditions, HIV/AIDS, renal disease and immune dysfunctions. All healthcare providers received a letter annually to inform them about the influenza vaccine and the definition of the high-risk groups, but the high-risk individuals themselves were not generally informed about the benefits of influenza vaccination. High-risk individuals who were insured by a social sickness fund could get vaccinated free of charge.[3] High-risk individuals who were covered by a private insurer or individuals not belonging to the high-risk group had to pay 38 euros (vaccine price in 1996 expressed in 2019 purchasing power) to obtain the vaccine from the pharmacy, but the actual administration of the vaccine by the general practitioner (GP) was for free as there is no coinsurance or deductible for GP care in the Netherlands. Throughout the eighties, the take-up rate in the high-risk group was rather constant and remained below 30% (Fedson et al., 1995; van Essen et al., 2001).

In 1991, the national health authorities concluded that vaccination coverage among chronically ill and elderly individuals was inadequate. A series of interventions was started to increase vaccination uptake. First, the general public and high-risk patients were informed about the existence and the benefits of influenza vaccination. Second, the position of the GP – who occupies the role of gatekeeper in the Dutch health care system – was strengthened. The GP was encouraged to register and personally invite high-risk patients and organize weekly vaccination walk-ins which do not require a routine appointment and thus reduce the time cost and planning for the vaccine recipient (van Essen et al., 2001).[4] The central role of the GP and the increased publicity rapidly increased influenza vaccination coverage in the high-risk group from about 30% in 1991 to 50% in 1995.

2.2. The Dutch influenza vaccination program since 1996

In 1996, influenza prevention of high-risk individuals effectively evolved into a nationwide preventive care program after a major reform extended the target group to all (healthy) individuals aged 65 and above and making all targeted individuals – including those covered by a private insurer – eligible for free influenza vaccination. Age eligibility was, however, not determined by calendar age. Rather, receipt of an invitation letter in September/October depended on one's program-age defined as whether the individual would be 65 on May 1st of the next calendar year. Hence, all individuals turning 65 between September/October and May 1st received the invitation letter even though their actual calendar age was 64.

The earlier interventions to increase take-up amongst the high-risk group – registration, personal invitation (around late September/early October) and vaccination walk-ins for at-risk patients – were continued. In addition, the reform introduced a GP remuneration per vaccinated target individual, and simplified the provision of influenza

[3] In the eighties, nineties and up to 2005, two thirds of the population - whose earnings (employment or replacement income) fell below an income threshold - were compulsory insured for health care by a social sickness fund. The remaining third could (voluntarily) enroll in private insurance, and fewer than two percent was uninsured.

[4] The fraction of GPs that effectively did these additional tasks increased rapidly. Registration of high-risk patients in a computer program was performed by 54% of the GP in 1994 and 82% in 1997. Personal invitations were sent out by 40% in 1994, 77% in 1997 and 95% in 2000. Vaccination walk-ins were organized by 72% in 1994, 86% in 1997 and 90% in 2000 (Hak et al., 2000; Tacken et al., 2002).

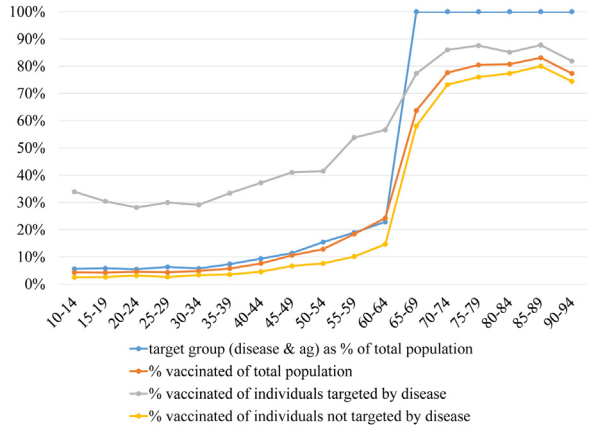

Fig. 1. Vaccination rates and high-risk group fraction by program age for the pooled influenza seasons 1997–1998 up to 2007–2008. Note: Health Interview Survey, own computations.

vaccination to the target group. Prior to the reform, an individual first had to get a prescription from the GP, then go to the pharmacist to buy the vaccine and then revisit the GP for the administration of the vaccine. After the reform, every GP had their own stock of vaccines which were directly administered to eligible individuals. The nationwide influenza program did not directly alter the incentives for untargeted individuals, who still had to pay to obtain the vaccine itself from the pharmacy, but did not face a deductible or coinsurance for the administration of the shot.[5]

In this study we exploit the discontinuity in the free influenza vaccination provision at program-age 65 to study how this preventive care program has affected vaccination behavior of family members of individuals that turned 65 between 1997 and 2008. Note that throughout the study period, flu-related care and treatment (i.e. prescribed drugs, physician visits, hospitalizations) were covered by compulsory social health insurance or private health insurance.

The reform coincided with a sharp increase in vaccination rates of the target population in the influenza season 1996–1997 which further increased until 2005, both in the entire population (up to 19%) and among individuals aged 65 or above (up to 80%) (see Fig. 1). In this time period, the Netherlands had one of the highest vaccination rates of the target population in Europe (Mereckiene et al., 2012, 2014) and higher vaccination rates than in the US (Lu et al., 2005), although the vaccination rate among those not targeted was higher in some other countries (Lu et al., 2008; Blank et al., 2009).[6] Furthermore, Fig. 1 shows that vaccination rates rise up to around program-age 70. This is in line with the findings of Carman and Mosca (2014) showing that individuals above 65 who start vaccinating against influenza within the context of the Dutch influenza program continue to do so in the next year.

In 2008, the national health authorities extended the population targeted by the nationwide influenza prevention care program to include all individuals aged 60 or more, thus effectively lowering the age threshold from 65 to 60. All other interventions remained in place. We use this threshold in a robustness check since we only had data for four influenza seasons since 2008 which give rise to an underpowered RDD design.

3. Data

3.1. Data sources and sample restrictions

Our main data source is the annual cross-sectional health interview survey (HIS) 1997–2008 which constitutes representative cross-sections of the non-institutionalized Dutch population and includes information on vaccination take-up. Each wave includes approximately 10,000 respondents, but only samples one individual per household.[7] We therefore have no information on vaccination take-up or demographics of family members of this individual. We merged the HIS to administrative data of Statistics Netherlands in order to add demographic information (sex and age in months) as well as residential address of the parents and partners of individuals included in the HIS.[8] The formal definition of a partnership in the registry data requires that both partners live at the same address and are married or registered partners.[9]

A typical influenza season in the Netherlands ranges from September to May with an increased activity between November and March. This period does not coincide with the wave period, which follows the civil year.[10] Since the month in which the survey questionnaire was completed is known, we have rearranged the waves into influenza seasons that start in September and end in August in the following year. Influenza season 1996–1997 is therefore only partly observed, and hence the corresponding observations

[5] Note that some employers provide (free or subsidized) flu shots to their employees. In the period 1997–2008, 14% of the individuals who were vaccinated and who did not belong to the high-risk group, indicated that they got vaccinated at the initiative of their employer (own computations, health interview surveys).

[6] Note that the overall vaccination rate of 19% is still well below the rate where marginal benefits of additional vaccination coverage are reduced to zero. Ward (2014) estimates this rate to be 33% in Canada at the time when a similar vaccination program as that in place in the Netherlands was expanded to the full population.

[7] More specifically: 10898 in 1997, 9323 in 1998, 9877 in 1999, 9922 in 2000, 9676 in 2001, 9745 in 2002, 9876 in 2003, 11117 in 2004, 10378 in 2005, 9607 in 2006, 8741 in 2007, 9499 in 2008.

[8] We only observe parents if they were still alive in 1995 (which is when the registry started) and living in the Netherlands. The individuals without observed parents are predominantly older individuals and individuals of foreign origin. If the difference between a child's and a parent's age falls below 15 years or exceeds 40 years, the observations are considered as outliers and excluded from the analysis. The restriction is binding for less than 0.5% of the individuals with parents, and given the sample age requirements laid out below, it effectively restricts the age range of the adult children to 20–51. The registry allows us to link children to their legal parents. Except in a limited number of cases, legal parents coincide with biological parents.

[9] A registered partnership is for partners who do not wish to marry, but the legal and tax implications are identical to those of a marriage.

[10] Interviews in the HIS are conducted over the entire year. The survey months are almost uniformly distributed, with a slight underrepresentation of July and August, which are the main holiday months.

for 1997 are dropped. We further drop all observations corresponding to influenza season 2008–2009 because the age-eligibility threshold for the free vaccination program was lowered to age 60 from 2008–2009 onwards. We pool all remaining influenza seasons (1997–1998 to 2007–2008) and end up with 108,533 observations from which three samples are derived.

Each sample includes vaccination behavior of individuals in the HIS, but merges to the age of different people that are in close vicinity to the age discontinuity induced by the free influenza vaccination program, i.e. the person itself, his or her partner, and his or her parents. Linkage of HIS individual's vaccination take-up to own age is used to estimate the direct policy effects on those turning 65. The sample of HIS individuals whose partner's age is derived from the administrative records allows estimating the spillovers between partners. We consider partners with and without children, but remove individuals whose partner differs less than 12 months in age, and do not consider or link age-*ineligible* (i.e. program age lower than 65) and age-*eligible* (i.e. program-age 65 or above) individuals to the program-age of their younger and older partner, respectively. Third, spillovers from parents to children are estimated on the sample of HIS individuals for whom at least one *parent* in close vicinity to the 65 program-age threshold can be identified from the administrative records. More details on the exact identification strategy for all three samples and the reasons for the sample restrictions are provided in Section 4.

All three samples were restricted to observations within a window of ±2 years around the 65 program-age cut-off of the relevant individual (partner, parent or individual her/himself)[11] and observations without information on the dependent variables and/or controls were removed.[12] This leaves 3183 observations for the estimation of the direct effects, 2068 for the partner spillovers, and 3766 observations for the parent-to-child spillovers.

3.2. Variable description and summary statistics

The most important indicators in the HIS for our analyses are the individual's month and year of birth and vaccination history. Vaccination history includes information on whether the individual ever got an influenza vaccination, and reports month and year of the last flu shot. We created a dummy variable that equals one [zero] if the individual did [not] get vaccinated against influenza during the influenza season. However, since most shots are reported to be taken in the period September to January,[13] individuals interviewed in this period might report not having taken a flu shot, but might still get vaccinated in the near future. For this particular group, the dummy equals one if they got vaccinated in the past influenza season, since this is by far the best predictor for a renewed vaccine take-up (Carman and Mosca, 2014).[14]

The HIS also provides us with additional control variables: sex, educational attainment (primary, lower secondary, upper secondary, post-secondary), household composition (single, couple, household with children, other), number of household members, influenza season (dummy for each season), population density of place of residence (500, 500–2500, 2500+ inhabitants per squared kilometer), existing medical conditions (in order to identify individuals who belong to the high-risk group based on chronic disorders), and presence of a chronic illness.[15]

Table 1 presents summary statistics for each of the three different samples.[16] The column 'direct effect' shows that around 40% of the population between 63 and 66 years old got vaccinated during the current influenza season. About half suffer from a chronic illness and one quarter qualifies for free influenza vaccination based on existing medical disorders. These shares are similar for the partners (column 'partner-to-partner spillover').[17] The column 'parent-to-child spillover' shows the average characteristics of the adult children whose older parents' age falls within the 63–66 program-age bandwidth. When both parents can be identified from the administrative records, we use the age of the older parent to guarantee that no other parent already qualifies for free influenza vaccination based on age-eligibility; when only one parent can be identified,

[11] The criterion by Imbens and Kalyanaraman (2012) indicates an optimal bandwidth of ±2 years for the spillovers between partners and the direct effects; and an optimal bandwidth of ±1.4 years for the spillovers to children. We use the ±2 year bandwidth for all analyses, but show robustness to another bandwidth in section 5.2.

[12] We removed 1.0, 0.8 and 0.3 percent of the observations due to missing in the samples for respectively the direct effects, the partner spillovers and the child spillovers. RDD models with a binary dependent variable indicating whether the observation was removed due to missing dependent and/or control variables could not reject the null hypothesis that those missing are random. The respective RDD estimates are 0.032 (*p-value*: 0.237; $n = 4009$) for the direct effects, 0.005 (*p-value*: 0.419; $n = 2085$) for the partner spillovers, and 0.008 (*p-value*: 0.760; $n = 4688$) for the child spillovers.

[13] More specifically, 95% of all reported influenza shots are reported to be taken in this period, and 86% in October and November, which is the recommended vaccination period. 4.5% in September, 2% in December and 2% in January.

[14] Not unexpectedly, this procedure has the biggest effect on individuals surveyed in the month September, and a much smaller effect in October-January. While it could lead to non-random measurement error in our dependent variable, we scrutinize this in Section 5.2 and find that there is little reason for concern.

[15] Presence of a chronic illness is a binary indicator for reporting a long-term illness, infirmity or handicap. Existing medical conditions is a binary indicator taking one when an individual reports at least one of the following conditions: (a) asthma; (b) heart disease; (c) liver disease; (d) kidney disease; (e) diabetes; (f) rheumatism; and (g) cancer. The information on liver disease is not available after 2000.

[16] Table A.1 provides additional summary statistics for a ±3 year bandwidth.

[17] The partner-to-partner sample contains a small, but non-zero share of single persons. This happens because the partner-to-partner sample is created based on the household composition registered in administrative records on 1 October at the start of the influenza season. This date was deliberately chosen, because it corresponds closely to the date of receipt of the invitation letter to the vaccination program. As HIS interviews are evenly spread throughout the calendar year, this date, however, differs for most individuals from the month of interview. Hence marital status as reported in the HIS might differ from the administrative records as individuals might e.g. divorce between October and the month of HIS interview. Alternatively, as marital status is self-reported in the HIS, there could be a difference in interpretation, e.g. individuals could be legally married at the time of the HIS interview, but in the process of divorce or separation and therefore report to live in a single household.

Table 1
Summary statistics.

	Parent-to-child spillover	Partner-to-partner spillover	Direct effect
	63–66	63–66	63–66
Binary variables (0 = no; 1 = yes)			
Vaccination rate	0.06	0.44	0.40
Education level: primary	0.06	0.23	0.23
lower secondary	0.18	0.33	0.31
upper secondary	0.44	0.25	0.27
post-secondary	0.32	0.18	0.19
Household type: single person	0.12	0.01	0.20
couple	0.20	0.89	0.72
household with children	0.67	0.10	0.07
other	0.01	0.00	0.01
Population density: < 500 inhabitants/km^2	0.14	0.15	0.15
$500 \leq$ inhabitants/km^2 ≤ 2500	0.68	0.72	0.70
$2500 <$ inhabitants/km^2	0.18	0.13	0.16
Male	0.49	0.46	0.51
High-risk group	0.07	0.26	0.25
Chronic illness	0.26	0.45	0.47
Older parent male	0.64		
Continuous variables			
Own program age	35.83 (4.39)	64.07 (6.00)	64.93 (1.15)
Number of household members	3.14 (1.31)	2.12 (0.49)	1.89 (0.57)
Older parent's program age	64.98 (1.14)		
Other partner's program age		64.93 (1.12)	
Number of observations	3766	2068	3183

Note: All cells show weighted sample means. Standard deviations in parentheses for continuous variables.

we use the age of that parent. 6% of the individuals in that sample got vaccinated during the current influenza season. About half are male, their average program-age is 36, and a quarter suffers from a chronic illness. 7% belongs to the high-risk group, who qualifies for free influenza vaccination based on existing disorders.

4. Identification strategy

4.1. Age-based RDD design

Our empirical approach exploits the fact that the nationwide influenza prevention program reduces the barriers to vaccination take-up exactly at program-age 65. We exploit this age discontinuity to estimate how vaccination take-up was impacted among those directly affected by the program upon turning 65, and use the same quasi-exogenous variation to identify spillover effects on vaccination take-up of partners (within household) and adult children (across households, except for those adult children living with their parents). The identifying assumption of the age-based RDD imposes that the discontinuity at age 65 does not correlate with any discontinuity in the observable or unobservable determinants of vaccination uptake. We show this assumption is credible for the direct and both spillover effects, and report a set of internal validity checks in Section 5.2. Here, we highlight the most important conceptual arguments in favor of our age-based RDD identification strategy.

Since age cannot be manipulated, age-triggered RDD's may be invalid when (1) the age-cut-off also determines eligibility for other programs which could affect the outcome of interest, (2) the policy impact on the outcome is not immediate, and (3) individuals anticipate the program-age cut-off (Lee and Lemieux, 2010). Eligibility for other programs should not matter for the direct impact of the Dutch influenza vaccination policy, since program age for the influenza program is different from calendar age which matters for eligibility for other social programs in the Netherlands. The program eligible age of individuals is computed on May 1st, at the end of an influenza season, and determines whether an individual will receive an invitation in September/October of the preceding year. Therefore, newly invited individuals at the margin around the 65 program-age threshold on May 1st are in fact 64 years old when they receive their first invitation in September/October. This feature of the vaccination program removes much of the concern that eligibility for other benefits that start on the day one turns 65 in the Netherlands (e.g. pension claims, benefits for the elderly) interferes with eligibility for influenza vaccination.[18] Eligibility for other programs neither affects the estimation of spillover effects as the sample restrictions ensure that partners do not cross the program-age cut-off and adult children are between 20 and 51 years old (see also Section 3.1 and footnote 14). The other two concerns – immediate effect and no anticipation – should be small for the direct and spillover effects since influenza vaccination is protective after about 2 weeks, it needs to be taken yearly as

[18] In the Netherlands, the health insurance system was dramatically reformed in 2006. However, neither before 2006 nor after 2006, there has been a discontinuity in the coverage of medical care insurance (in general and for flu shots) at the age of 65 (Mossialos and Thomson, 2002; Roos and Schut, 2008). Only for Dutch retirees actually living outside the Netherlands (in another EU country) the reform has led to a reduction in generosity, which was only later acknowledged by the European Court of Justice (case C-345/09, October 14th, 2010). Note that Dutch retirees residing outside of the Netherlands are not part of our sample.

antibodies decline within the year, and the influenza virus mutates every year.

The age discontinuity further addresses several potential biases, such as ensuring that herd immunity faced by individuals and their family members should be similar just right and left of the program-age cutoff. The same should be true for shared family genes and habits between family members, and the age discontinuity also safeguards against potential bias resulting from assortative mating and endogenous social group formation that might apply when studying spillovers to partners and children (Becker, 1973, 1974; Manski, 1993; Moffitt, 2001).[19]

4.2. Empirical specification

Let us first consider the effect of the nationwide influenza preventive care program among all individuals (with and without partners and/or children) that reach the 65 program-age eligibility cutoff. In our RDD setting, their vaccination behavior can be linearly modelled as

$$V_i = \alpha + g(age_i - 65) + h(age_i - 65)D_i + \lambda D_i + X\prime_i\beta + \varepsilon_i \quad (1)$$

where V_i is a binary variable indicating whether individual i gets vaccinated, α is a constant, age_i is her/his age on May 1st in years and months, D_i is a treatment dummy that equals one if $age_i \geq 65$, $g(.)$ and $h(.)$ are unknown functional forms, ε_i are unobservable variables, and there are no time subscripts as each individual is only observed once in our set of pooled cross-sections. We also include a vector of observable control variables X_i (with associated parameter vector β) to increase precision, and which includes sex, educational attainment, household size and composition, influenza season, population density of place of residence, member of high-risk group based on chronic disorders, and presence of a chronic illness (see also Section 3.2). The change in vaccination take-up at the threshold – the divergence in vaccination rates between individuals just left and right of the program-age cutoff – is captured by λ which measures the combined effect of the program incentives.[20] High risk individuals under 65 with existing chronic disorders also qualify for the vaccination program, and might not or to a lesser extent adjust their vaccination take-up when turning 65. Such treatment heterogeneity is ruled out in the description of all empirical specifications in this section but will be revisited in Section 5.

Parameter λ in Eq. (1) informs on the direct impact of the program-age cutoff at 65, but also helps gauging the relative magnitude of spillovers onto partners and children. The spillovers on the partners are modelled as in Eq. (1), but here we include the program-age of one's partner (age_i^{Pr}) instead of own age. We consider individuals with and without children:

$$V_i = \alpha^{Pr} + g^{Pr}\left(age_i^{Pr} - 65\right) + h^{Pr}\left(age_i^{Pr} - 65\right)D_i^{Pr} + \lambda^{Pr}D_i^{Pr} + X\prime_i\beta^{Pr} + G_i^{Pr}\gamma^{Pr} + k(age_i) + \varepsilon_i^{Pr} \quad (2)$$

where the superscript Pr refers to 'partner', D_i^{Pr} is a treatment dummy that equals one if one's partner turns 65 on May 1st, G_i^{Pr} (and associated parameter γ^{Pr}) refers to one's partner's sex and $k(.)$ is an unknown functional form. Inclusion of own age overcomes a potential bias in the spillover effect λ_{pr} in Eq. (2) resulting from correlation between partner's ages. Own age – i.e. $k(age_i)$ – is included as a set of 3-monthly dummies, except in the lower and upper tail of the age distribution where we used wider age dummies. As mentioned in Section 3.1, we delete all partners whose program-ages differ less than 12 months to avoid both partners receiving the first invitation for influenza vaccination in the same influenza season which makes it impossible to distinguish the spillover effect from the direct effect. We also delete age-eligible individuals whose older partner also qualifies for a free vaccination, and age-ineligible individuals whose younger partner also falls below the age cutoff of 65. This ensures that the age-eligibility status of individuals in the HIS is independent of the age-eligibility of their partner, and that we can estimate the spillover in isolation from the direct effects.[21]

The spillovers on the adult children are obtained by linking individuals to their older (or only) parent's age on May 1st (age_i^{Pt}):

$$V_i = \alpha^{Pt} + g^{Pt}\left(age_i^{Pt} - 65\right) + h^{Pt}\left(age_i^{Pt} - 65\right)D_i^{Pt} + \lambda^{Pt}D_i^{Pt} + X\prime_i\beta^{Pt} + G_i^{Pt}\gamma^{Pt} + l(age_i) + \varepsilon_i^{Pt} \quad (3)$$

where superscript Pt refers to 'parent', D_i^{Pt} is a treatment dummy that equals one if $age_i^{Pt} \geq 65$, G_i^{Pt} refers to the older (or only) parent's sex and $l(.)$ is an unknown functional form. Own age – i.e. $l(age_i)$ – is included as a set of 3-monthly dummies, except in the lower and upper tail of the age distribution where we use wider age dummies.

Eqs. (1)–(3) allow intention-to-treat estimation of the policy's direct impact and spillovers: λ reflects the combined direct effects of all program incentives, but λ^{Pr} and λ^{Pt} do not distinguish between competing mechanisms driving the spillovers as only program-age and not the actual vaccination take-up of those turning 65 is known. For

[19] The alternative identification strategy of relying on random social group assignment (Duncan et al., 2005; Kremer and Levy, 2008; Yakusheva et al., 2011; Carrell et al., 2011) often raises concerns about the potentially artificial nature of the social group.

[20] In a sensitivity analysis that is available from the authors upon request, we find a similar effect size estimate using a logit model.

[21] We address the fact that each partner faces two age discontinuities – i.e. the individual reaching program-age 65 (direct effect) and the partner reaching program-age 65 (spillover effect) – by combining four features: (1) separate RDD models in partner's age for the spillover and direct effects, i.e. Eqs. (1) and (2); (2) removing partners differing less than twelve months; (3) ensuring independent age-eligibility status of both partners; and (4) using a flexible specification for own age to account for potential correlation between own and partner's age. An alternative identification strategy, which we term double RDD, relates an individual's vaccination take-up to trends in own and partner program-age at either side of the relevant age thresholds, i.e. basically replacing $k(age_i)$ in Eq. (2) by $g(age_i - 65) + h(age_i - 65)D_i + \lambda D_i$. This allows estimating the direct and spillover effects from one specification but at the cost of imposing a trade-off between (1) sufficient age difference between both partners to identify spillovers separately from direct effects; and (2) focusing on a sufficiently narrow window around the age discontinuities in own and partner-program age to satisfy RDD identifying assumptions.

example, we cannot distinguish between parental program eligibility driving child spillovers only via that parent's vaccination take-up versus child spillovers being driven by the vaccination uptake of both parents, i.e. both the vaccination take-up of the individual turning 65 and the vaccination take-up of the other partner due to the spillover effect. It is nevertheless feasible to get a deeper understanding of the nature of parent-to-child spillovers by exploiting the fact that a child has two parents who each will eventually cross the program age threshold, except when they die before. We therefore perform subgroup sensitivity analyses among those children for whom we can identify both parents: children's vaccination take-up can be linked to the older parent, while linking to the younger parent implies that the adult child has already been exposed to the older parent turning 65. The difference between both estimates provides indirect evidence on the underlying mechanisms driving the spillovers. Parent-to-child spillovers might also differ when parents respond differently to the vaccination policy depending on whether they are the first or second in the household to qualify for free vaccination, and we explore this by re-estimating our models separately for the subpopulations of older and younger partners, but also by allowing for treatment heterogeneity in the spillovers between partners (i.e. from older to younger partner and vice versa) (see Section 6).[22]

Finally, power calculations with power 0.8 as in Schochet (2009) confirm that the pooled samples for Eq. (1), (2) and (3) have sufficient power to identify age discontinuities in vaccination take-up of respectively 8.1, 10.2 and 3.7 percentage points (see Section 3.1 for exact sample sizes). The analyses based on the older/younger subdivision have inevitably less power, in particular the subsamples used for the spillovers between partners and the direct effects.[23] We therefore first present our main analyses derived from the pooled samples and check robustness of these estimates in Section 5. Estimates allowing for the older-younger hetereogeneity are discussed in Section 6.

5. Direct policy effects and spillovers to partners and children

5.1. Main analyses

Let's first consider the direct policy effect at age 65. Panel A in Fig. 2 shows an RDD graph of the influenza vaccination take-up of individuals around the 65 program age cut-off grouped in 3-month age bins.[24] Vaccination rates jump from just above 30% to around 45% at the cut-off, and display an increasing slope with age among age-eligible individuals most likely due to a further influx into the immunization program and a negligible drop-out rate (Carman and Mosca, 2014). The corresponding RDD estimate of 9.8 percentage points (p-value: 0.002; n =3183) in the final column of Table 2 – which assumes linear, but different trends on either side of the cut-off – is precisely estimated, and corresponds to a 30% relative effect size. This estimate is obtained after controlling for eligibility for free vaccination based on existing medical conditions, but assumes homogeneity with respect to this eligibility criterium. This might be an unrealistic assumption since those with existing chronic disorders already benefit from the vaccination program and therefore might not (or differently) update their vaccination behavior. Estimates obtained from subsamples of individuals with and without existing medical conditions, confirm that the homogeneous treatment effect in Table 2 is mostly driven by low-risk individuals, i.e. their vaccination uptake increases with 10.5 percentage points (p-value: 0.005; n =2376) while the RDD estimate is substantially smaller and insignificant for those belonging to the high-risk group (0.058, p-value: 0.370; n =807).[25] We further find no evidence for heterogeneity of the direct program impact between individuals with and without adult children. Restricting the sample to parents only yields a point estimate of 9.6 percentage points (p-value: 0.005; n =2831), while further restricting the sample to older parents (single or older parent in a couple) – which is the subsample most likely resembling the parents involved in the parent-to-child spillovers – leads to a similar result of 8.5 percentage points (p-value: 0.036; n =1843).

Panel B in Fig. 2 presents the RDD graph for the vaccination behavior of individuals according to the program-age of their partner – grouped in 3-month program-age bins. These partner-to-partner spillovers include the spillovers that arise both when the older partner turns 65 as well as when the younger turns 65. The figure shows no support for spillovers in vaccination rates among partners, which corresponds to the relatively small and imprecisely estimated RDD estimate of 0.037 (p-value: 0.377; n =2068) reported in column 'partner-to-partner spillover' in Table 2. Leaving out individuals with existing medical conditions does not change the result, although this might also reflect the reasonably small number of individuals with existing medical conditions (n = 546); and analyzing treatment heterogene-

[22] Note that we do not use a two sample IV estimator (Angrist and Krueger, 1992) to scale the spillover effect from parent to child (and partner to partner) by the direct effect at age 65 as the two-sample IV estimator is inappropriate when the vaccination policy affects children's (partner's) vaccination decisions via *both* its direct impact on the vaccination behaviour of the parent (partner) and the indirect spillover impact on his/her partner (child) which is an untestable possibility in our setting. In addition, there is the concern – due to the time-invariant nature of the parent-child relationship, and the possibility that parents separate and form new partnerships – that the older/younger parent of the child might no longer be the older/younger partner.

[23] We can detect age discontinuities of 13.2–15.8 percentage points for the direct effects, 13.7–14.7 percentage points for the partner-to-partner subsamples, and 4.7–4.9 percentage points for the parent-to-child subsamples with power 0.8.

[24] The procedure described in (Lee and Lemieux, 2010) indicated a 3-month bin size for the target and partner spillover sample, but a 6-month bin size for the children spillover sample. To ease visual comparison of RDD graphs across samples, we use 3-month bin sizes for all panels in Fig. 2. Appendix Figure A.1 shows the corresponding RDD graphs for a ±3 year program-age bandwidth for panels A–C.

[25] This is further corrobarated by age discontinuities in the reason for vaccination among individuals that got vaccinated. Low-risk individuals were more likely to be invited by their GP, and less likely to get vaccinated on own request; while there is no such discernible pattern among high-risk individuals. Results available upon request.

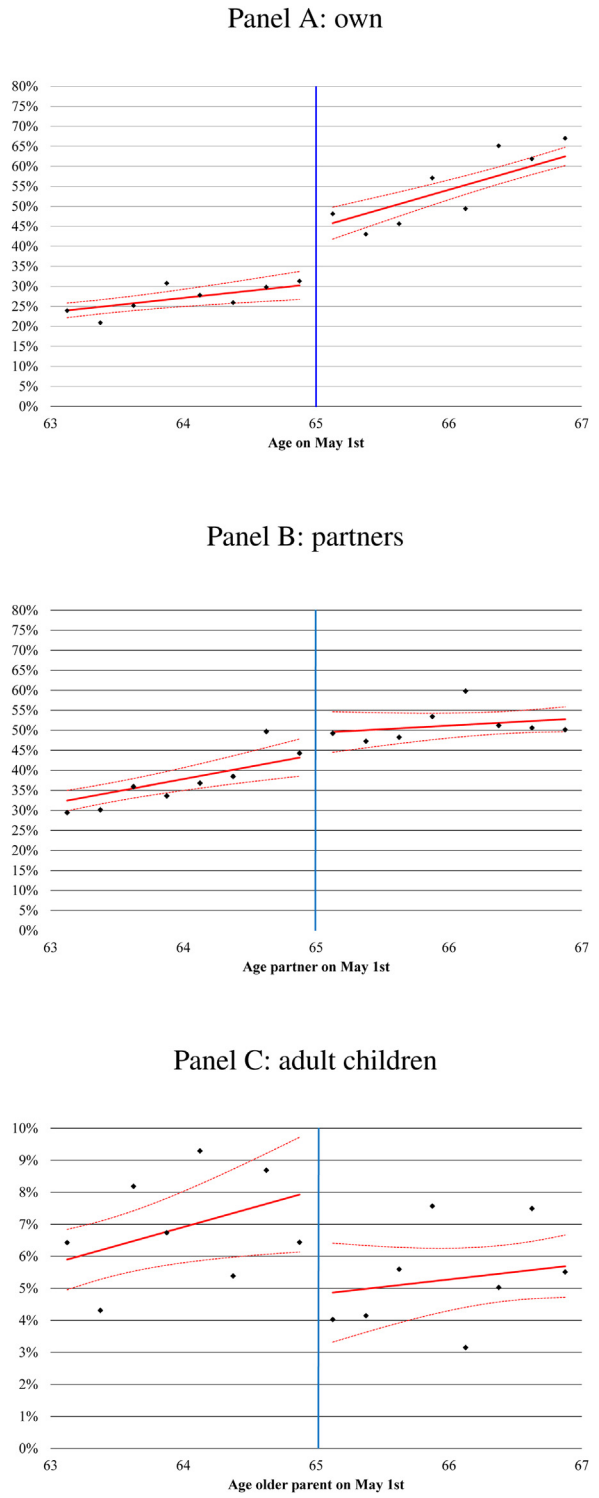

Fig. 2. Influenza vaccination rate according to own/other partner's/older parent's program-age (bandwidth = 63–66). Note: 3183 (panel A), 2068 (panel B), 3766 (panel C) observations from the pooled data of the influenza seasons 1997–1998 to 2007–2008 are used. The running variable is own (panel A), the other partner's (panel B), the older parent's program-age (panel C). Diamonds represent the weighted average influenza vaccination rate of the individuals (panel A), partners (panel B) or adult children (panel C), grouped in 3-monthly bins based on the running variable. The red line shows a linear trend based on the observations. The dotted lines represent the 95% confidence intervals. The vertical line represents the age threshold.

Table 2
Pooled RDD estimates of the (spillover) effect of the vaccination policy at the age threshold.

	Parent-to-child spillover	Partner-to-partner spillover	Direct effect
	63–66	63–66	63–66
Treatment effect	−0.025* (0.015)	0.037 (0.042)	0.098*** (0.032)
Number of observations	3766	2068	3183

Notes: See text and Section 3 and 4 for details on the RDD set-up. The estimates are based on pooled data of the influenza seasons 1997–1998 to 2007–2008 that is restricted to a ±2 year window around the age threshold. Linear trends in program age – that can differ at each side of the cutoff – are used. Control variables include dummies for sex, member of risk group based on existing disorders, population density, chronic illness, education level, number of household members, household type, and influenza season. The parent-to-child spillovers additionally control for the older parent's sex and child age. The partner-to-partner spillovers additionally control for partner's age and the other partner's sex. OLS regression estimates are reported that use sampling weights. Clustered standard errors at the wave-municipality level to mimic the sampling design are between brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ity with respect to one's partner program eligibility based on medical conditions is not feasible as the administrative data only records program-age, sex and residential address of partners (see also section 3).

Panel C in Fig. 2 presents an RDD graph of the vaccination behavior of adult children according to the program-age of their older parent – grouped in 3-month program-age bins. We find a decline in vaccination uptake from around 8 percent to 5 percent, which is a substantial drop in relative terms. This negative parent-to-child spillover is confirmed by the RDD regression estimate of −2.5 percentage points (p-value: 0.094; n =3766) in column 'parent-to-child spillover' of Table 2. Our sample has insufficient power to check treatment heterogeneity across program eligibility based on medical conditions – only 7 percent ($n = 281$) of the adult children reports existing medical conditions –, but it is reassuring that the spillover effect obtained from the subsample of adult children without existing medical conditions was only marginally larger (−0.021, p-value: 0.095, n = 3485), and not smaller as one would expect when those with existing medical conditions do not not alter their vaccination take-up when their older parent crosses the age threshold. Treatment heterogeneity with respect to existing medical conditions of parents is not feasible as program eligibility for this reason is not recorded in the administrative data.

Overall, our estimates provide no evidence in favor of spillovers between partners, but do indicate non-negligible parent-to-child spillovers in comparison to the direct impact of the vaccination policy at age 65 which is similar for individuals with and without children. We further find that program eligibility based on medical conditions weakens the direct policy impact, while spillovers do not strongly depend on these partners and adult children already qualifying for the program based on medical conditions.

5.2. Robustness and internal validity

A series of checks is done to confirm the internal validity of our RDD estimates. A first set of tests concerns sensitivity to the assumptions imposed to capture trends in vaccination take-up, i.e. the choice of window size and the parametric form used to model trends on either side of the cutoff. Appendix Table A.2 shows that the effects are stable or slightly larger for a +/- 3 year window, while quadratic trends mainly reduce the precision of the estimated effects. We further confirm that our baseline findings using linear trends are robust to the exclusion of control variables, and to the inclusion of province fixed effects for each influenza season which capture unobserved regional differences and trends in, for example, GP medical practices, intensity and spread of previous influenza seasons, incidence of other infectious diseases, information campaigns, etc. Next, clustering at the level of the running variable – own/partner/parental program-age in months – as compared to clustering at the municipality-wave level increases the precision of the estimates. Finally, we restrict the analysis to the survey months February to August, i.e. those survey months for which we know exactly whether someone got vaccinated or not (see footnote 14). For the direct effect at age 65, the imputation procedure for vaccination take-up for the months September to January may lead to an underestimation of the true effect at age 65 since the vaccination dummy for the newly invited individuals (i.e. those that are 65 on May 1st) in those months is partly based on their vaccination behavior in the previous year when they did not yet qualify for free influenza vaccination. The point estimates in Appendix Table A.2 increase after restricting the sample to the months February to August, but do not differ significantly from those based on a sample when observations for September to January are included. For the spillovers, restricting to February-August would only make a difference if the month in which the children or partners are interviewed matters for vaccination take-up. Our results in Appendix Table A.2 do not provide evidence for this.

The estimates in Table 2 represent intentions-to-treat (ITT). These ITT's are internally valid when treatment is random, i.e. when the only discontinuous change within a small window around the 65 program-age threshold is the change in eligibility for the vaccination policy. As discussed in Section 4.1, there is little a priori reason to question this assumption: age cannot be manipulated and anticipation makes little sense as influenza vaccination only protects for one year. Moreover, program age-eligibility does not coincide with calendar age, so that calendar-age triggered eligibility for other programs (such as pension benefits) cannot explain the found direct and spillover effects; and for the spillovers onto children there is the additional safeguard that children cannot choose their parents. We did five additional tests that confirm the robustness of our findings. First, we find that the control covariates are reasonably balanced around the 65 program-age threshold. The only exception is household type (see Appendix

Table A.3) although it is reassuring that inclusion or exclusion hardly affects the estimates (see Appendix Table A.2). Second, the *p*-values of McCrary's test, respectively 0.83, 0.14, and 0.24 for the direct effects, partner-to-partner and parent-to-child subsamples, indicate smooth program-age densities and thereby remove any concerns about heaping bias, or age-related attrition (McCrary, 2008). Third, as part of a placebo check, we find no discontinuous jumps in vaccination take-up when the program-age cutoff is set at value 62, 68 or 60 where no jump would be expected (see Appendix Table A.4).[26] While these three additional tests support the internal validity of our results, one might still worry that, due to their ITT-nature, the spillover effects reflect the total effect of parents/partners becoming eligible for the vaccination program, and not just the impacts of partner's/parent's increased likelihood to get vaccinated (see also footnote 22). When parents or partners update their vaccination behavior *but* also change other behaviors – such as preventive behavior – as a result of becoming eligible for the vaccination program, our spillover estimates will reflect both channels. We test for this possibility by running placebo RDD models on behaviors of individuals at age 65 in the domains of prevention, healthy behavior, altruism and adherence to alternative medicine as a change in any of these might also affect the vaccination behavior of their adult children or partners. Our estimates in Appendix Table A.5 provide no evidence for this.[27] . The fifth and final set of falsification/robustness checks confirm, although inevitably with far less power due to much smaller sample sizes, that (1) crossing the program-age threshold of 65 in calendar years 1992–1996, when program-age played no role in qualifying for free influenza vaccination (see Section 2), did not affect own vaccination take-up (compare rows A and C in Appendix Table A.6); and that (2) crossing the program-age cutoff of 60 (the relevant program-age cutoff in place since 2008) led to somewhat larger (in absolute size) but otherwise similar direct and spillover effects in influenza seasons 2008–2009 to 2011–2012 (compare rows A and B in Appendix Table A.6).[28]

6. Spillovers linked to older versus younger parent/partner

In this section we study how vaccination take-up among partners and children differs depending on whether the older or younger partner/parent crosses the 65 program-age cutoff. We first consider within-household spillovers onto partners. Younger partners differ from their older partner in some systematic ways, i.e. they tend to be female and are lower educated (see rows 'male' and 'education level' in column 'Spillover from older/younger to other partner' in Table 3). They are also far less likely to get vaccinated (see row 'vaccination rate').[29] Column 'Spillover from... to other partner' in Table 4 (and panel C and D in Fig. 3) reveals that the younger partner becomes 10.2 percentage points (*p*-value: 0.064; $n = 1108$) more likely to get vaccinated when the older partner crosses the 65 program-age cutoff, while the older partner's vaccination behavior does not change when the younger partner turns 65.[30] The lack of a spillover from young to old is perhaps not that surprising since the older are already age eligible for the influenza vaccination program, which is in line with Carman and Mosca (2014) who show that individuals continue vaccinating as soon has they have ever responded to the invitation from the Dutch vaccination program. More interestingly, the spillover effect from older to younger in panel II is of similar magnitude as the direct effect at age 65 (see column 'Direct effect among... parent' in Table 4), despite the fact that the younger partners do not face reduced barriers to influenza vaccination.

The vaccination policy might also differentially affect adult children depending on whether the older or younger parent crosses the age-threshold. In contrast to the partner-to-partner spillovers, we need to impose additional sample restrictions on the pooled sample in Table 1 before we can estimate separate child spillovers deriving from the older or younger parent turning 65. We additionally (a) remove children for whom only one parent could be identified as we don't know whether this is the older or younger par-

[26] The artificial program-age cutoffs of 62 and 68 avoid that the true program cutoff age of 65 is within the window, and 60 is the program-age cutoff that applied from influenza season 2008–2009 onwards.

[27] We consider mammography screening, body mass index, blood donation, consumption of alternative medicine, and visiting an acupuncturist, naturopathic doctor or psychic healer. We found no discontinuous jumps at program-age 65 for any of these indicators. This was also the case for an indicator of exercising, which has previously been shown to increase upon early and normal retirement in Germany and the United States (Eibich, 2015; Kampfen and Maurer, 2016)

[28] The HIS before and after 1996 are not perfectly comparable. First, we limit to the years 1992–1996 since HIS does not record influenza vaccination take-up prior to 1992. Second, we only report direct effects at age 65 (and not the spillovers) since the administrative records are not linked to the HIS prior to 1996 which is required to identify parental and partners program-age. Third, the HIS prior to 1996 does not record the interview month which is necessary to construct influenza seasons. Instead we work with calendar years and replace the controls for influenza season with controls for calendar year. Absence of information on interview month also implicates the dependent variable: the vaccination binary dependent variable for year t takes one when the individual received an influenza vaccination in year t or in the months September to December of year $t-1$ (see Section 3.2). Fourth, we use robust standard errors and do not account for clustering at the level of the municipality since the HIS does not record municipality of residence before 1996. Fifth, replication of the baseline direct effect estimates in row A using calendar years, calendar year dummies, alternative dependent variable definition and no clustering at the municipality level (see row D of Appendix Table A.6) confirms that the difference between the baseline estimates (row A) and the falsification test using data of 1992–1996 (row C) cannot be attributed to differences in methodology.

[29] The higher number of available observations to estimate spillovers from older-to-younger versus younger-to-older partners merely reflects a survival effect: both running variables equal on average approximately 65, but own program age is lower (higher) for spillovers towards the younger (older) partner (see row 'own program age'). This survival effect derives from (a) older individuals having higher mortality rates, (b) older individuals more often living in single person households, and (c) our sample restrictions which impose that both partners differ at least one year in age and that the partner who is affected by the influenza program via a spillover has the same eligibility status around the cutoff (younger than 65 for spillovers on the younger partner, and older than 65 for the spillovers on the older partner).

[30] The procedure described in (Lee and Lemieux, 2010) indicated a 3-month bin size panels A and B of Fig. 3, but a 6-month bin size for panels C-F. To ease visual comparison of RDD graphs across samples, we use 6-month bin sizes for all panels in Fig. 3.

Table 3
Summary statistics: older-younger subsamples within 63–66 program-age bandwidth.

	Adult children linked to... parent		Spillover from... to other partner		Direct effect among... parent	
	Older	Younger	Older	Younger	Older	Younger
Binary variables (0 = no; 1 = yes)						
Vaccination rate	0.06	0.06	0.24	0.67	0.37	0.43
Education level: primary	0.05	0.05	0.24	0.21	0.19	0.26
lower secondary	0.17	0.18	0.39	0.26	0.26	0.40
upper secondary	0.45	0.45	0.22	0.30	0.31	0.23
post-secondary	0.32	0.32	0.15	0.23	0.24	0.11
Household type: single person	0.13	0.11	0.01	0.01	0.01	0.01
couple	0.21	0.16	0.85	0.93	0.88	0.93
household with children	0.66	0.73	0.13	0.06	0.11	0.05
other	0.01	0.00	0.00	0.00	0.00	0.00
Population density: $<$ 500 inhabitants/km^2	0.14	0.15	0.15	0.15	0.16	0.18
500$\leq$ inhabitants/km^2 $\leq$ 2500	0.68	0.70	0.73	0.71	0.71	0.71
2500 $<$ inhabitants/km^2	0.18	0.16	0.12	0.15	0.13	0.11
Male	0.49	0.51	0.14	0.84	0.85	0.14
High-risk group	0.07	0.07	0.22	0.31	0.26	0.24
Chronic illness	0.26	0.26	0.45	0.46	0.44	0.47
Parent male	0.84	0.16				
Continuous variables						
Own program age	35.00 (4.10)	38.13 (3.96)	60.03 (4.86)	68.84 (2.94)	64.78 (1.11)	65.06 (1.12)
Number of household members	3.09 (1.30)	3.37 (1.36)	2.17 (0.58)	2.07 (0.34)	2.14 (0.49)	2.05 (0.30)
Number of observations	2310	2045	1108	960	1185	828

Note: All cells show weighted sample means. Standard deviations in parentheses for continuous variables. The summary statistics in column 'Older' are obtained after linking to the older parent/partner/own program-age, making sure that the other parent/partner can be identified and asserting that the younger parent/partner's program-age is smaller than 65. The summary statistics in column 'Younger' are obtained after linking to the younger parent/partner/own program-age, making sure that the other parent/partner can be identified and asserting that the older parent/partner's program-age is 65 or older.

ent; (b) remove children whose parents differ less than 12 months in age to avoid contamination by both parents receiving the first invitation in the same influenza season; and (c) ensure that the eligibility status of the 'untreated' parents is the same on either side of the program-age cutoff (cf. Section 3.1). Next, we subdivide the remaining sample in two subsamples where we either link to their younger or older parent. Summary statistics in Table 3 reveal little difference between children linked to their older or younger parent, except that adult children linked to their younger parent are 3 years older and have more children themselves.[31] The resulting RDD estimates are presented in column 'Adult children linked to... parent' in Table 4 and corresponding RDD graphs are presented in panel E and F of Fig. 3.[32] The results in Panel II show that adult children are 4.5 percentage points (p-value: 0.015; n =2310) less likely to get vaccinated when their older parent turns 65. This is almost double the effect size obtained in the baseline analysis in Table 2. At the same time, we find no evidence that adult children update their vaccination behavior when the younger parent turns 65 (Panel III). This suggests that spillovers of the vaccination program on the adult children crucially differ depending on whether their parent is the first or second in the family to receive an invitation for a free influenza vaccination. Our data do not allow testing whether this difference is related to older parents typically being male, but we do not reject the null hypothesis of similarly sized spillovers towards daughters and sons.[33] The different spillover between first and sec-

[31] The same reasoning as in footnote 29 explains why more adults can be linked to their older parent than to their younger parent. Note further that the same child might appear in both subsamples when both parents are within the 63–66 age range which explains why the combined number of observations in columns 'Adult children linked to Older/Younger parent of Table 3 surmounts that in column 'Parent-to-child spillover' of Table 1, even though the latter includes all baseline observations. This is different in the partner-to-partner subsamples: an indivdiual only features in one of both subsamples because the HIS survey only interviews one household member.

[32] All RDD estimates and graphs in Table 4 and Fig. 3 are replicated with a +/- 3 year bandwidth in Appendix Table A.8 and Appendix Figure A.2. We also present summary statistics with a +/- 3 year bandwith in Appendix Table A.7. Appendix Table A.8 also presents estimates obtained after excluding all parents differing less than 24 months in age which confirms robustness to a mechanical composition effect that occurs when only excluding parents differing less than 1 year. In the latter case, the age difference is 12 months or more to the left of the cutoff and for the first year to the right of the cutoff – when linking to the older parent –, but the minimum age difference linearly increases from 12 to 24 months for the second year to the right of the cutoff. A similar reasoning applies when linking to the program-age of the younger parent.

[33] Absence (presence) of spillovers from younger (older) parents to adult children might depend on children's sex (for example, Carpenter and Lawler (2019) find that girls react stronger to school mandates than boys in US). This possibility is analyzed with t-tests of the difference in treatment effects as derived from RDD models using sex-specific and younger/older parent-specific subsamples. We find that spillovers from older parents to both female and male adult children are negative and the estimate for female children is only 0.010 percentage points larger than that for male children (p-value = 0.788). Spillovers from younger parents to female and male adult children are not significantly different from zero (p-values of 0.315 and 0.500; and the spillover to females is 0.048 percentage points larger than that to males (p-value = 0.234)). Treatment heterogeneity of children's spillovers by parental sex, and treatment heterogeneity of the direct effects and partner spillovers by sex of the partners/parents was not analyzed because the relevant subsamples are too small as men tend to be the older partner in around 85 percent of couples (see row "male" in columns "spillovers from. . ." and "direct effect. . ." in Table 3).

ond parent reaching the 65 program-age cutoff could in principle also result from a stronger direct response to the vaccination policy among the older versus younger parent. We find no evidence in the HIS sample that older parents respond more vigorously when becoming eligible for free influenza vaccination compared to the younger parents (see column 'Direct effect among... parent' and panel A and B in Fig. 3).[34] In Section 7 we further discuss potential mechanisms underlying these findings.

7. Further evidence and discussion

In this section we further explore the potential mechanisms underlying the asymmetric spillovers between partners and from parents to children. We already established in sections 5 and 6 that neither anticipating one's 65th birthday, changes in parental prevention, health behavior, altruism, adherence to alternative medicine or gender heterogeneities are likely mechanisms underlying the spillovers. The data also does not support any confounding by regional and/or time variation in GP medical practices, spread of influenza and other infectious diseases, or information campaigns. In this section, we mostly follow Kremer and Miguel (2007) and Fadlon and Nielsen (2019), and discuss the relevance of the following mechanisms in light of the results presented in sections 5 and 6: (a) epidemiological externalities (and the related concept of cross-protection); (b) information transmission, covenience, updating of target group membership and risk perceptions, (c) perceived shortages of vaccines; (d) salience; and (e) we also reflect on the health impact of the spillover effects. Additional empirical evidence is presented in this section, but should be interpreted cautiously since (a) statistical power weakens when checking treatment heterogeneity on increasingly smaller subpopulations (see Section 4.2), (b) we cannot follow influenza vaccination behavior of the same individual over time, and (c) unfortunately the data provide no information on whether the linked partners/parents are getting vaccinated before and after reaching the 65 program-age threshold which would be required to go beyond ITT estimates (more discussion in Section 4.2 and footnote 22).

Negative epidemiological externalities arise when individuals realize their probability to be infected depends negatively on the vaccination behavior of other individuals they frequently interact with. An individual might therefore stop getting vaccinated when their parent or partner takes up vaccination as it reduces their marginal benefit of getting vaccinated (Kremer and Miguel, 2007). We cannot test this hypothesis directly because we cannot follow influenza vaccination behavior over time, we do not observe partner's/parent's vaccination behavior, and our data do not inform on the frequency of interactions between different family members. However, we can do an indirect test, as one would expect more frequent contact, and thus a bigger role for negative epidemiological externalities, between family members that belong to the same household (within-household spillovers) as compared to those that are part of different households (between-household spillovers). This would imply that the spillover-estimates in Tables 2 and 4 should be smaller for the partner-to-partner than the parent-to-child spillovers. This is exactly opposite to what we find suggesting that, if anything, negative epidemiological externalities matter for the parent-to-child spillovers, in particular when the older parent turns 65, but not for the spillovers between partners. We scrutinize this possibility by proxying frequency of contact between parents and children by the geographical distance between their residential addresses which should be a reasonable proxy since the Netherlands is a small and densely populated country,[35] and test whether parent-to-child spillovers are more negative for parents and children living in close vicinity compared to those living further apart. However, when estimating a model allowing for an interaction between the discontinuity and parent-children proximity[36] we could not reject homogeneity of the treatment effect (p-value = 0.374). Furthermore, the absence of a spillover effect from the younger parent to the child is further evidence at odds with negative externalities explaining the parent-to-child spillovers. An alternative, but related explanation for negative child spillovers might be positive epidemiological externalities. Such cross-protection effects consist of children altruistically getting vaccinated to protect their parents from getting the flu. When parents turn 65 and get vaccinated themselves, children might no longer see the need for getting vaccinated. This could in principle play a role in the Dutch context since children providing informal care for their parents are recommended to obtain a flu shot, although lack of data on informal caregiving prohibits test-

[34] Compared to the baseline analyses in row 'baseline' of Table 4, we additionally removed (a) childless people belonging to the 63–66 age bandwidth; (b) single parents, (c) parents differing less than one year in age, and (d) ensured that the eligibility status of 'untreated' parents is the same on either side of the cutoff. Since parents can only appear in the younger or older subsample, these sample restrictions leads to a lower combined number of observations in panels II and III compared to the number of observations in the basline analysis in panel I of Table 4. Summary statistics in column 'direct effect among... parent' of Table 3 indicate that older parents are more likely male and higher educated compared to younger parents. In addition, the reasoning in footnote 29 partially explains the larger sample size for 'direct effect among older (versus younger) parent', but there is an additional mechanism that increases the difference in sample size between older and younger parents: all are on average 65 years old (all are within the 63–66 age interval), but their partners are respectively on average 60 and 69 years old. This affects household composition because older individuals are more likely to be the older partner. Hence at age 69 individuals are more likely to be the older partner than at age 65 ($n = 960$ vs $n = 828$ in row 'Number of observations' of Table 3), and the reverse argument holds for ages 60 and 65 ($n = 1108$ vs $n = 1185$ in the same row of Table 3).

[35] For each individual in the HIS and their parents in the administrative register, we have an encrypted address location which allows calculating the distance between children's and parent's municipalities of residence. As vaccination invitation letters are sent out by the end of September, we take the addresses on October 1st for each influenza season.

[36] In our sample, 3.63 percent of the adult children co-habit with their parents, 45.29 percent live separately but still within the same municipality, 23.25 percent live in a different municipality but still within 20 km, 13.64 percent within 20-60 km, and the remaining 14.19 percent live more than 60 km apart.

Table 4
RDD estimates of the (spillover) effects of the vaccination policy – older/younger partner/parent.

	Adult children linked to... parent	Spillover from... to other partner	Direct effect among... parent
	63–66	63–66	63–66
I. Baseline	−0.025* (0.015) [3766]	0.037 (0.042) [2068]	0.098** (0.032) [3183]
II. Older (conditional on younger parent/partner being under age 65)	−0.045** (0.019) [2310]	0.102* (0.055) [1108]	0.099* (0.052) [1185]
III. Younger (conditional on older parent/partner being 65 or above)	0.002 (0.021) [2045]	−0.024 (0.065) [960]	0.118* (0.066) [828]

Notes: See text and Section 3 and 4 for details on the RDD set-up. The estimates are based on pooled data of the influenza seasons 1997–1998 to 2007–2008 that is restricted to a ±2 year window around the age threshold. Linear trends in program age – that can differ at each side of the cutoff – are used. Control variables include dummies for sex, member of risk group based on existing disorders, population density, chronic illness, education level, number of household members, household type, and influenza season. The parent-to-child spillovers additionally control for the older parent's sex and child age. The partner-to-partner spillovers additionally control for partner's age and the other partner's sex. The estimates in Panel II reflect the age discontinuities in vaccination take-up obtained after linking to the older parent/partner/own program-age, making sure that the other parent/partner can be identified and asserting that the younger parent/partner's program-age is smaller than 65. The estimates in Panel III reflect the age discontinuities in vaccination take-up obtained after linking to the younger parent/partner/own program-age, making sure that the other parent/partner can be identified and asserting that the older parent/partner's program-age is 65 or older. OLS regression estimates are reported that use sampling weights. Clustered standard errors at the wave-municipality level to mimic the sampling design are between brackets. Number of observations between square brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ing this hypothesis directly. Nevertheless, as women are much more likely to provide informal care for their parents compared to men (Oudijk et al., 2010), the absence of differential spillovers towards male or female children reported in Section 6 and footnote 33 is not in line with the hypothesis of positive cross-protection effects. All in all, we therefore tentatively conclude that epidemiological externalities and cross-protection are an unlikely explanation for the partner-to-partner and parent-to-child spillovers.

Well-informed individuals are known to have higher vaccination take-up rates (Bronchetti et al., 2015; Nuscheler and Roeder, 2016). When becoming age-eligible for the free influenza vaccination program, individuals learn about costs and benefits of influenza vaccination and might transfer this information to their partners, including information on tangible costs, perceptions about the effectiveness and the safety of the vaccine, fear about side effects, needles and pain, and knowledge on time costs (van Essen et al., 1997; Chapman and Coups, 1999; Wu, 2003; Zijtregtop et al., 2010; Bish et al., 2011; Nagata et al., 2013). This explanation is line with the positive spillover from the older to younger partner; while the absence of spillovers from the younger to older partner is in line with the older already having received invitations (and the associated information) before.[37] These patterns are also consistent with both partners visiting the GP together due to the younger partner experiencing a reduced barrier to access the GP. Such a convenience argument is sensible in countries with substantial barriers to access health care providers. For example, Carpenter and Lawler (2019), Schaller et al. (2019) and Hoffmann et al. (2019) have shown that vaccination take-up and/or spillovers towards non-mandated vaccines mostly arise through overcoming the barriers and opportunity costs of accessing health care providers. Although our data rule out formal testing of this convenience argument, it should be less relevant in the Netherlands: all GPs organize special walk-ins throughout the week, there are no co-payments or out-of-pocket expenditures for GP visits; and non-eligible individuals (the younger partners) first need to visit the GP for a prescription, next have to buy the vaccine from the pharmacy and finally revisit the GP to get vaccinated; while the eligible individuals (the older partners) can get vaccinated during their first GP visit (see also Section 2). Getting vaccinated together therefore requires a substantial amount of coordination between both partners; but at the same time does neither substantially reduce the opportunity cost of getting vaccinated, nor does it increase access to GPs.[38]

A special type of information transmission concerns updating of the target group which could explain the negative parent-to-children spillovers, in particular since Dutch citizens are among the European citizens that put most trust in their government and attach importance to governmental information on health issues (Bults et al., 2010; van der Weerd et al., 2011). When individuals receive their first invitation for influenza vaccination, their offspring might learn that they do not belong to the target group of the influenza vaccination policy, while target group membership is already internalized by the time the younger parent turns 65.[39] This is in line with previous descriptive Dutch

[37] This is corroborated by the patterns in the reported reasons for influenza vaccination: the share of vaccinations on own request among younger partners showed the largest increase (8 percentage points) when their older partner became age-eligible (compared to vaccination because of program-related GP invitation, program-unrelated GP recommendation, specialist/employer/community health care centre recommendation, other reason); while no meaningful change in any reason was found among the spillovers towards older partners.

[38] The convenience argument could be more reasonable when both partners visit a GP with an in-practice pharmacy, because the additional steps of getting a prescription and buying the vaccine from the pharmacy can be avoided. Since only 141 younger partners (out of 1108) visit a GP with an in-practice pharmacy we lack the statistical power to test whether the partner-to-partner RDD estimates are stronger for this group. In addition, since only age and sex of the older partner are linked from the administrative registers, we do not know whether both partners visit the same GP.

[39] This interpretation might seem at odds with the earlier reported homogeneity in the negative parent-to-children spillover across physical proximity between parents and children in the context of epidemiological externalities. One should however realize that actual and frequent physi-

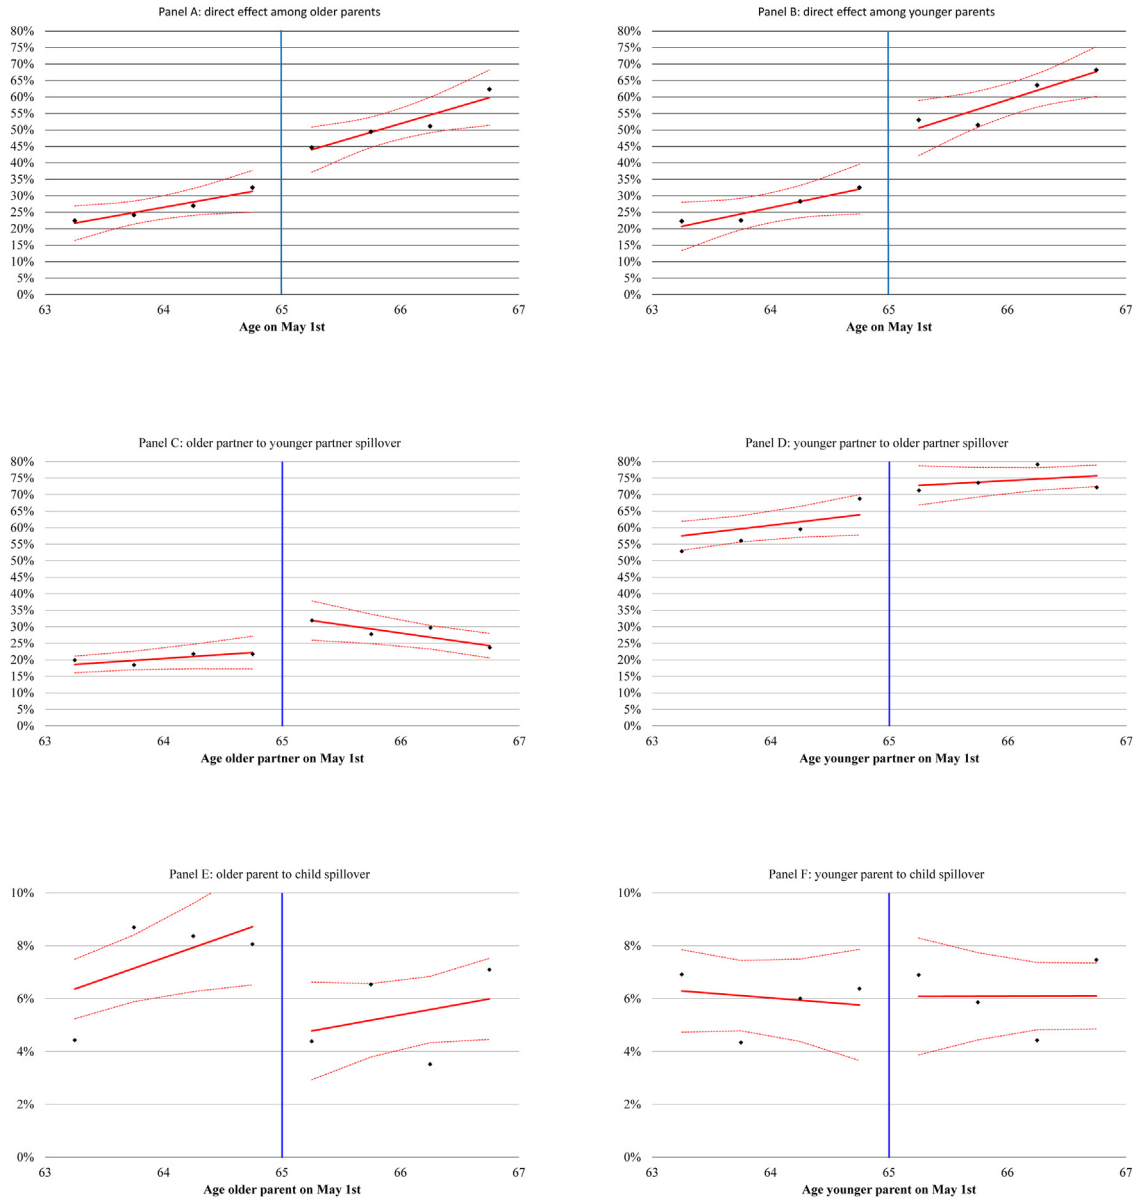

Fig. 3. Own/partner's/children's influenza vaccination rate around the 65 program-age threshold: young vs old (bandwidth = 63–66). Note: 1185 (panel A), 828 (panel B), 1108 (panel C), 960 (panel D), 2310 (panel E), 2045 (panel F) observations of the pooled data of the influenza seasons 1997–1998 to 2007–2008 are used. The running variable is program-age. The RDD graph in panel A (B) is obtained after linking to the individual's own program-age and making sure that the younger (older) parent is younger (older) than 65. The RDD graph in panel C (D) is obtained after linking to the program-age of the older (younger) partner and where the younger (older) partner is younger (older) than 65. The RDD graph in panel E (F) is obtained after linking to the older (younger) parent's program-age and using children for whom both parents can be identified and where the younger (older) parent is younger (older) than 65. Diamonds represent the weighted average influenza vaccination rates of parents, partners and adult children grouped in 6-monthly bins based on program-age. The red line shows a linear trend based on the observations. The dotted lines represent the 95% confidence intervals. The vertical line represents the age threshold.

evidence showing that not being part of the target group is mentioned as (second) most important reason for not getting vaccinated against seasonal or pandemic influenza, in particular for individuals that do not qualify for the program based on existing medical conditions (Kroneman and Verheij, 2003; Bults et al., 2010; van der Weerd et al., 2011). The importance of target group membership updating is further corroborated by the exact wording used in the invitation letters sent by GPs. We obtained example invitation letters drawn up by the Dutch college of general

cal contacts are required for epidemiological externalities while updating of the target group, and more generally information transmission, only requires non-physical contact. Unfortunately, we have no accurate measurements on the frequency of physical and/or non-physical contact (e.g. phone, email, etc.).

practitioners that serve as example for the GPs to personally invite the population at risk.[40] These letters underline that getting vaccinated offers protection against influenza, that vaccination is free of charge and provide the following definition of the target group: "Individuals aged 65 or more,[41] and individuals of all ages with heart, lung or renal disease, diabetes or with immune dysfunctions, are at increased risk to fall seriously ill due to influenza. You belong to one of these groups. That is why you are eligible for influenza vaccination." The reference to the 65 age threshold can be confusing for individuals who receive the letter for the first time. The invitations are received around September/October, and since program-age is determined as of May 1st in the following year, the real age of individuals whose month of birth is between September/October and April is only 64. For this group, the receipt of the letter can come across as a mistake made in the invitation procedure, which can contribute to discussing the letter more intensely with family. The mechanism of updating target group membership seems less important for the spillovers between partners as the vaccination rate of younger partners increases when their older partner turns 65. This increase is mostly driven by couples that differ only a few years in age. Younger partner's vaccination rates do not decrease in couples with a larger age difference, i.e. the pattern one would expect when updating of target group membership occurs mostly among partners that are sufficiently younger than the 65-year program age-cutoff.[42]

Perceived (or actual) vaccine shortages could further strengthen the reduction in vaccination uptake due to updating of target group membership. Under real (or perceived) pressure on vaccination capacity, adult children might forego getting vaccinated in order to give priority to individuals that are more at risk. While it is reassuring that the estimates in Table A.2 confirm robustness of the pooled parent-to-children spillovers to inclusion of province fixed effects for each influenza season (that account for region-specific trends in perceived and actual shortages), our estimates could still be driven by perceived vaccine shortage coinciding exactly with one's older parent turning 65. We have no information on individual's perceived shortages, but there were no actual influenza vaccine shortages between 1997–1998 and 2007-2008 in the Netherlands, except for a mild shortage in 2005 which was solved by December 2005 (Tacken et al., 2006) and which attracted some attention in newspapers (Scholtens, 2005). Furthermore, the bird flu spread around the Eurasian continent in 2003–2006, and in 2003 it reached the Netherlands. So, if anything, we would expect, in decreasing order, more confounding due to perceived vaccine shortages in 2005, 2003, 2004 and 2006. However, when re-estimating the older parent-to-children spillovers with a model allowing for an interaction between the discontinuity and an indicator for influenza seasons 2003–2004 and 2005–2006, we could not reject homogeneity of the treatment effect (p-value: 0.831). The same finding was obtained when interacting with an indicator for 2003–2004 to 2006–2007 (p-value: 0.119). We thus conclude that perceived (or actual) shortages are an unlikely explanation for the negative parent-to-child spillovers.

We thus hitherto conclude that the opposite signs of the spillover effects between partners and towards children are most likely explained by an information transmission and updating of target group membership. The data patterns are however also in line with underestimation of the flu mortality risk at older ages – information transmission should then lead to an upward update of the flu mortality risk; with a more modest (or no) shift the lower the age of the younger partner –; and overestimation at younger ages - where updating of group membership coincides with a negative update of the flu mortality risk (Carman and Kooreman, 2014). While the exact distinction between risk updating on the one hand, and information transmission (partners) and updating of the target group (children) on the other hand is impossible to establish with our data, a perhaps equally important distinction is that between spillovers being driven by learning (learning about costs and benefits of vaccination, updating target group or updating flu risk) and/or the salience of the paternal/partner's receipt of the invitation for free influenza vaccination. Recent work on non-communicable health shocks shows that within family spillovers in preventive care uptake and harmful habits partly reflect learning but mostly salience of the health shock (Fadlon and Nielsen, 2019), but there is no comparable literature on within family vaccination spillovers for communicable diseases. In order to learn more about the potential role of salience, we compare spillovers towards non-eligible individuals with spillovers towards individuals qualifying for free vaccination based on existing chronic disorders, but not based on their age. A change in vaccination uptake in the latter group that coincides with their older partner turning 65 is in line with a larger role of salience versus learning compared to a similar change in vaccination uptake in the former group.[43] Our estimates in Section 5.1 did not indicate that spillovers towards partners or adult children strongly depend on qualifying for the program based on medical conditions; thus suggesting that salience plays a role for the spillovers. How important salience is versus learning is difficult to establish, but the negative sign of the spillovers towards children and the more positive spillovers towards partners

[40] Note that the GP is not obligated to follow the example invitation letter, and is free to formulate an alternative invitation letter or not to send out an invitation at all. As we discussed in footnote 4, most GPs do send out a personal invitation.

[41] We could only obtain access to example invitation letters for 2009, 2011 and 2012 (not for the years 1997–2008) and therefore these letters actually mentioned "60 or more".

[42] The RDD estimate of the older-to-younger partner spillovers among couples that differ less than 4 years in age amounts to 0.121 (p-value: 0.093, $n = 575$), while that among couples differing 4 or more years in age amounts to 0.041 (p-value: 0.673; $n = 533$). Treatment heterogeneity along groups based on smaller (3 or 2 years) or larger age differences (5 or more) was not pursued because such groupings reduced sub-sample sizes too much for credible RDD estimation.

[43] While those with existing chronic disorders qualify for free influenza vaccination, and the group without not; the cost of getting vaccinated is similar on both sides of the RDD spillover threshold for both groups as the younger partners cannot qualify for the program based on age when their older partner turns 65.

with a more similar age, as compared to partners with a much younger age, underlines that learning also matters (see also footnote 42).

Unraveling the mechanisms is important to better understand the underlying behavioral responses, but does not inform on the health impact of the spillover effects. In a companion paper, we show that the vaccination policy had a substantial direct health and health care impact at age 65 (Van Ourti and Bouckaert, 2020). During each month with high influenza virus circulation, we found, on average, 1.5 per 100,000 fewer influenza and pneumonia deaths, a 15 percentage point lower probability to use prescribed medicines and 0.13 fewer GP visits at program-age 65; while no discontinuities were observed during months with low virus circulation.[44] These estimates show that the policy had a non-negligible direct effect – both directly but also via the spillovers – on health care use at program-age 65; while the mortality effect, even though statistically significant, was small.[45] While corresponding estimates of how the policy 'spilled over' to the health care use and mortality experiences of the partners and children of individuals that turned 65 are not presented in this paper due to insufficient statistical power,[46] the findings in our companion paper do suggest that the policy, via increased vaccination take-up among younger partners, further reduced mortality and GP visits and prescribed medicines use. The net impact is however unclear because adult children of parents at the 65 program-age discontinuity reduced vaccination uptake leading potentially to more health care use. How much this counterbalances the reduced health care use among parents (and partners) is unclear, but it could be substantial as the influenza virus' effectiveness at preventing influenza infections, i.e. the antibody response to influenza vaccination, is inversely related to age (Lang et al., 2010; Grohskopf et al., 2013). This counterbalancing argument is probably less important for the mortality impact because influenza-related mortality hardly occurs at younger ages (Thompson et al., 2004), and because our identification strategy accounts for increased parental infection probability when children are getting vaccinated less.

[44] These estimates additionally confirm that the estimates reported in Tables 2 and 4 are not just the result of social desirability bias.

[45] Similarly small mortality impacts at age 65 have recently been reported by Anderson et al. (2020) for the UK.

[46] We used the same data to estimate the health care use impact in our companion paper, but the identification strategy was extended with a comparison between RDD estimates in months with low versus high virus circulation. This reduces the sample size considerably, in particular for months with high virus circulation because there are only between 0 and 3 such months per influenza season. This leaves sufficient statistical power for identification of the direct policy impact at age 65, but the additional reduction in sample size required to link to older versus younger parent or partner reduces sample sizes too much. This holds also true for the mortality effects, even though they were estimated from register data for the entire Dutch population, as the combination of reduction in sample size from linking to older or younger partner/parent with influenza-related deaths being rare, in particular among children, leaves too little statistical power.

8. Conclusion

In this paper, we investigate spouses' and adult children's influenza vaccination response when partners or parents become age-eligible for free influenza vaccination in the Netherlands. This focus on within-family spillovers is justified as the household is the most important transmitter of the influenza virus, more so than childcare, school or workplace environments. Our results thus contribute to the design of preventive care policies by pointing to the importance of spillovers, neglect of which might lead to under- or overestimation of overall program (cost-)effectiveness.

Using a rich dataset that combines survey data on individual vaccination uptake with administrative records on demographics of partners and parents, and exploiting the age-eligibility rules of the Dutch influenza vaccination program, we show that vaccination spillovers in the family context are almost as important as the direct impact on vaccination uptake. We find a positive vaccination response when an individual's older partner turns 65, while the adult children, living mostly in other households, reduce their vaccination take-up by almost half when the older parent becomes age-eligible.

These findings have several important implications. We are the first to present evidence of important and asymmetric within-family influenza vaccination uptake spillovers in a high income country. Our findings suggest that family networks contribute to information transmission of costs and benefits of influenza vaccination, but the salience effect of the invitation for a free flu shot also runs through families. Furthermore, our results underline the importance of public health campaigns paying attention to the effects of public perceptions and attitudes on (voluntary) preventive care participation. In the case of influenza vaccination, participation outside the target group is not harmful for the personal health of the program-ineligible individual, and generates indirect protection to individuals at risk. If ineligible individuals choose to finance their own vaccination, it should not be discouraged by policymakers, but directly or indirectly encouraged, e.g. by adjusting the information message that is distributed. One possible policy adjustment in this respect could be to less explicitly focus on the delineation of the target group, or alternatively, more explicitly focus on the benefits of indirect protection (next to direct protection) when addressing the target group and the broader population. The combination of the findings in this paper with those in a companion paper (Van Ourti and Bouckaert, 2020), suggest substantial program-associated health care effects of within-family vaccination spillovers, although the overall program effect in the Netherlands ultimately depends on the persistence of the spillovers and on the positive health care impacts among partners outweighing the negative impact on adult children.

Authors' contribution

Nicolas Bouckaert: Conceptualization, Methodology, Software, Formal Analysis, Writing – Original Draft. Anne Gielen: Conceptualization, Methodology, Software, Formal Analysis, Writing – Editing. Tom Van Ourti: Conceptual-

ization, Methodology, Writing - Original Draft & Review & Editing.

Appendix A. Supplementary Data

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.jhealeco.2020.102386.

References

Anderson, Michael L., Dobkin, Carlos, Gorry, Devon, 2020. The effect of influenza vaccination for the elderly on hospitalization and mortality: an observational study with a regression discontinuity design. Ann. Intern. Med. 172 (7), 445–452.

Angrist, J.D., Krueger, A.B., 1992. The effect of age at school entry on educational attainment: an application of instrumental variables with moments from two samples. J. Am. Stat. Assoc. 87 (418), 328–336.

Becker, Gary S., 1973. A theory of marriage: Part I. J. Polit. Econ. 81 (4), 813–846.

Becker, Gary S., 1974. A theory of marriage: Part II. J. Polit. Econ. 82 (2, Part 2), S11–S26.

Bish, A., Yardley, L., Nicoll, A., Michie, S., 2011. Factors associated with uptake of vaccination against pandemic influenza: a systematic review. Vaccine 29, 6472–6484.

Blank, Patricia R., Schwenkglenks, Matthias, Szucs, Thomas D., 2009. Vaccination coverage rates in eleven European countries during two consecutive influenza seasons. J. Infect. 58 (6), 446–458.

Bronchetti, Erin Todd, Huffman, David B., Magenheim, Ellen, 2015. Attention, intentions, and follow-through in preventive health behavior: field experimental evidence on flu vaccination. J. Econ. Behav. Org. 116, 270–291.

Bruhin, Adrian, Goette, Lorenz, Haenni, Simon, Jiang, Lingqing, 2020. Spillovers of prosocial motivation: evidence from an intervention study on blood donors. J. Health Econ. 70, 102244.

Bults, M., Beaujean, D.J., Zwart, O.D., Kok, G., Empelen, P.V., Steenbergen, J.V., 2010. Mexicaanse griep: risicoperceptie bij de bevolking, eigen maatregelen en vertrouwen in overheidsinformatie. Nederlands Tijdschrift voor Geneeskunde 154, A1686.

Carman, K.G., Mosca, I., 2014. Who takes up free flu shots? Examining the effects of an expansion in coverage. De Economist 162 (1), 1–17.

Carman, K.G., Kooreman, P., 2014. Probability perceptions and preventive health care. J. Risk Uncertain. 49 (1), 43–71.

Carpenter, Christopher S., Lawler, Emily C., 2019. Direct and spillover effects of middle school vaccination requirements. Am. Econ. J. Econ. Policy 11 (1), 95–125.

Carrell, S.E., Hoekstra, M., West, J.E., 2011. Is poor fitness contagious?: Evidence from randomly assigned friends. J. Public Econ. 95 (7), 657–663.

Cawley, John, Han, Euna, Kim, Jiyoon, Norton, Edward C., 2019. Testing for family influences on obesity: the role of genetic nurture. Health Economics 28 (7), 937–952.

CDC, 2010. Prevention and control of influenza with vaccines. Recommendations of the Advisory Committee on Immunization Practices (ACIP). Morbid. Mortal. Wkly. Rep. 59 (RR-8), 1–62.

Chao, D.L., Halloran, M.E., Obenchain, V.J., Longini Jr., I.M., 2010. FluTE, a publicly available stochastic influenza epidemic simulation model. PLOS Comput. Biol. 6 (1), 1–8.

Chapman, G.B., Coups, E.J., 1999. Predictors of influenza vaccine acceptance among healthy adults. Prev. Med. 29, 249–262.

Cutler, D.M., Glaeser, E.L.<Ed>, 2010. Social interactions and Smoking. In: Wise, D.A. (Ed.), Research Findings in the Economics of Aging. The University of Chicago Press, pp. 123–141.

Duncan, G.J., Boisjoly, J., Kremer, M., Levy, D.M., Eccles, J., 2005. Peer effects in drug use and sex among college students. J. Abnorm. Child Psychol. 33 (3), 375–385.

Eibich, P., 2015. Understanding the effect of retirement on health: mechanisms and heterogeneity. J. Health Econ. 43, 1–12.

Fadlon, Itzik, Nielsen, Torben Heien, 2019. Family health behaviors. Am. Econ. Rev. 109 (9), 3162–3191.

Fedson, D.S., Hannoun, C., Leese, J., Sprenger, M.J.W., Hampson, A.W., Jørgensen, K.B., Ahlbom, A.-M., Nøkleby, H., Valle, M., Olafsson, O., Salmeron Garcia, F., Gugelman, R., de Andrade, H.R., Snacken, R., Ambrosch, F., Donatelli, I., 1995. Influenza vaccination in 18 developed countries, 1980–1992. Vaccine 13 (7), 623–627.

Ferguson, N.M., Cummings, D.A.T., Fraser, C., Cajka, J.C., Cooley, P.C., Burke, D.S., 2006. Strategies for mitigating an influenza pandemic. Nature 442, 448–452.

Fletcher, J., Marksteiner, R., 2017. Causal spousal health spillover effects and implications for program evaluation. Am. Econ. J. Econ. Policy 9 (4), 144–166.

Golberstein, E., Eisenberg, D., Downs, M.F., 2016. Spillover effects in health service use: evidence from mental health care using first-year college housing assignments. Health Econ. 25 (1), 40–55.

Grohskopf, L.A., Shay, D.K., Shimabukuro, T.T., Sokolow, L.Z., Keitel, W.A., Bresee, J.S., Cox, N.J., 2013. Prevention and control of seasonal influenza with vaccines: recommendations of the advisory committee on immunization practices - United States, 2013–2014. Morbid. Mortal. Wkly. Rep. 62 (45), 1–43.

Hak, E., Hermens, R.P., Hoes, A.W., Verheij, T.J., Kuyvenhoven, M.M., van Essen, G.A., 2000. Effectiveness of a co-ordinated nation-wide programme to improve influenza immunisation rates in the Netherlands. Scand. J. Prim. Health Care 18 (4), 237–241.

Hoffmann, Manuel, Mosquera, Roberto, Chadi, Adrian, 2019. Vaccines at Work, Available at SSRN.

Imbens, G.W., Kalyanaraman, K., 2012. Optimal bandwidth choice for the regression discontinuity estimator. Rev. Econ. Stud. 79 (3), 933–959.

Kampfen, F., Maurer, J., 2016. Time to burn (calories)? The impact of retirement on physical activity among mature Americans. J. Health Econ. 45, 91–102.

Kremer, M., Levy, D., 2008. Peer effects and alcohol use among college students. J. Econ. Perspect. 22 (3), 189–189.

Kremer, M., Miguel, E., 2007. The illusion of sustainability. Q. J. Econ. 122 (3), 1007–1065.

Kroneman, M.W., Verheij, R., 2003. De griepprik in Nederland: motivatie voor deelname en distributiekanalen. NIVEL, Technical Report.

Lang, P.O., Govind, S., Mitchell, W.A., Kenny, N., Lapenna, A., Pitts, D., Aspinall, R., 2010. Influenza vaccine effectiveness in aged individuals: the role played by cell-mediated immunity. Eur. Geriatr. Med. 1 (4), 233–238.

Lee, D.S., Lemieux, T., 2010. Regression discontinuity designs in economics. J. Econ. Lit. 48, 281–355.

Longini, I.M., Nizam, A., Xu, S., Ungchusak, K., Hanshaoworakul, W., Cummings, D.A.T., Halloran, M.E., 2005. Containing pandemic influenza at the source. Science 309 (5737), 1083–1087.

Longini Jr., I.M., Halloran, M.E., Nizam, A., Yang, Y., 2004. Containing pandemic influenza with antiviral agents. Am. J. Epidemiol. 159 (7), 623–633.

Lu, P., Bridges, C.B., Euler, G.L., Singleton, J.A., 2008. Influenza vaccination of recommended adult populations, U.S., 1989–2005. Vaccine 26 (14), 1786–1793.

Lu, P., Singleton, J.A., Rangel, M.C., Wortley, P.M., Bridges, C., 2005. Influenza vaccination trends among adults 65 years or older in the United States, 1989–2002. Arch. Intern. Med. 165 (16), 1849–1856.

MacDonald, Pauline, 2016. Extension of influenza immunization program to children in England – future plans. Hum. Vaccines Immunother. 12 (10), 2707–2708.

Manski, C.F., 1993. Identification of endogenous social effects: the reflection problem. Rev. Econ. Stud. 60 (3), 531–542.

McCrary, J., 2008. Manipulation of the running variable in the regression discontinuity design: a density test. J. Econom. 142 (2), 698–714.

Mereckiene, J., Cotter, S., Nicoll, A., Lopalco, P., Noori, T., Weber, J.T., D'Ancona, F., Levy-Bruhl, D., Dematte, L., Giambi, C., Valentiner-Branth, P., Stankiewicz, I., Appelgren, E., Flanagan, D.O., VENICE project gatekeepers group, 2014. Seasonal influenza immunisation in Europe. Overview of recommendations and vaccination coverage for three seasons: pre-pandemic (2008/09), pandemic (2009/10) and post-pandemic (2010/11). Euro Surveill. 19 (16), 20780.

Mereckiene, J., Cotter, S., Weber, J.T., Nicoll, A., D'Ancona, F., Lopalco, P.L., Johansen, K., Wasley, A.M., Jorgensen, P., Lévy-Bruhl, D., Giambi, C., Stefanoff, P., Dematte, L., O'Flanagan, D., 2012. Influenza A(H1N1)pdm09 vaccination policies and coverage in Europe. Euro Surveill. 17 (4), 103–112.

Moffitt, R.A., 2001. Policy Interventions, Low Level Equilibria, and Social Interactions. In: Durlauf, S.N., Young, H.P. (Eds.), Social Dynamics. MIT Press, pp. 45–82.

Mossialos, E., Thomson, S., 2002. Chapter 6: Voluntary health insurance in the European Union. In: Mossialos, E., Dixon, A., Figueras, J., Kutzin, J. (Eds.), Funding health care in: options for Europe. Open University Press, pp. 128–160.

Nagata, J.M., Hernandez-Ramos, I., Kurup, A.S., Albrecht, D., Vivas-Torrealbas, C., Franco-Paredes, C., 2013. Social determinants of health and seasonal influenza vaccination in adults over 65 years: a

systematic review of qualitative and quantitative data. BMC Public Health 13, 388.
Nuscheler, Robert, Roeder, Kerstin, 2016. To Vaccinate or to procrastinate? That is the prevention question. Health Econ. 25 (12), 1560–1581.
Oudijk, D., Boer, A., Woittiez, d., Timmermans, I., Klerk, J.M.D., 2010. In the Spotlight: Informal Care in the Netherlands. Technical Report. Netherlands Institute for Social Research, The Hague.
Ourti, Tom Van, Bouckaert, Nicolas, 2020. The Dutch influenza vaccination policy and medication use, outpatient visits, hospitalization and mortality at age 65. Eur. J. Public Health 30 (2), 275–280.
Roos, A., Schut, E., 2008. Evaluatie aanvullende en collectieve verzekeringen 2008. Onderzoek in opdracht van de Nederlandse Patiënten Consumenten Federatie. Erasmus University Rotterdam, Technical Report.
Schaller, Jessamyn, Schulkind, Flisa, Shapiro, Teny, 2019. Disease outbreaks, healthcare utilization, and on-time immunization in the first year of life. J. Health Econ. 67, 102212.
Schochet, P.Z., 2009. Statistical power for regression discontinuity designs in education evaluations. J. Educ. Behav. Stat. 34 (2), 238–266.
Scholtens, B., 2005. Griepvaccin niet meer te krijgen," de Volkskrant, 2005, https://www.volkskrant.nl/mensen/griepvaccin-niet-meer-te-krijgen b52c2946/d (accessed on June 11th, 2020).
Szilagyi, P.G., Bordley, C., Vann, J.C., Chelminski, A., Kraus, R.M., Margolis, P.M., Rodewald, L.E., 2000. Effect of patient reminder/recall interventions on immunization rates: a review. J. Am. Med. Assoc. 284 (14), 1820–1827.
Tacken, M., de Bakker, D., Verheij, R., Mulder, J., van den Hoogen, H., Braspenning, J., 2002. LINH Evaluatie Griepvaccinatiecampagne 2001. Nivel, Technical Report.
Tacken, M., Verheij, R., Mulder, J., van den Hoogen, H., Braspenning, J., 2006. Monitoring griepvaccinatieprogramma 2005. Nivel, Technical Report.
Thompson, W.W., Shay, D.K., Weintraub, E., et al., 2004. Influenza-associated hospitalizations in the United States. J. Am. Med. Assoc. 292 (11), 1333–1340.
van der Weerd, W., Timmermans, D.R., Beaujean, D.J., Oudhoff, J., van Steenbergen, J.E., 2011. Monitoring the level of government trust, risk perception and intention of the general public to adopt protective measures during the influenza A (H1N1) pandemic in the Netherlands. BMC Public Health 11 (1), 575–587.
van Essen, G.A., Kuyvenhoven, M.M., de Melker, R.A., 1997. Compliance with influenza vaccination. Its relation with epidemiologic and sociopsychological factors. Arch. Fam. Med. 6, 157–162.
van Essen, G.A., Kuyvenhoven, M.M., Masurel, N., Diepersloot, R.J.A., van der Graaf, M., Verheij, T.J.M., 2001. Influenza vaccination in the Netherlands: a successful system approach. Int. Congress Ser. 1219, 689–695.
Vrieze, H.A., van Haaren, K.M.A., Drenthen, A.J.M., Schaik, B., Dayan, M., 2016. Griepvaccinatie: de organisatie en uitvoering van de influenzavaccinatie in de huisartsenpraktijk. Nederlands Huisartsengenootschap, Technical Report.
Ward, C.J., 2014. Influenza vaccination campaigns: is an ounce of prevention worth a pound of cure? Am. Econ. J. Appl. Econ. 6 (1), 38–72.
Welliver, R., Monto, A.S., Carewicz, O., Schatteman, E., Hassman, M., Hedrick, J., Jackson, H.C., Huson, L., Ward, P., Oxford, J.S., 2001. Effectiveness of oseltamivir in preventing influenza in household contacts: a randomized controlled trial. J. Am. Med. Assoc. 285 (6), 748–754.
White, C., "Measuring Social and Externality Benefits of Influenza Vaccination," Journal of Human Resources, forthcoming.
WHO, 2019. Influenza (Seasonal) Fact Sheet. Technical Report. World Health Organization, Geneva, accessed March 28, 2019.
Wu, S., 2003. Sickness and preventive medical behavior. J. Health Econ. 22, 675–689.
Yakusheva, O., Kapinos, K., Weiss, M., 2011. Peer effects and the freshman 15: evidence from a natural experiment. Econ. Hum. Biol. 9 (2), 119–132.
Zijtregtop, E.A.M., Wilschut, J., Koelma, N., Van Delden, J.J.M., Stolk, R.P., Van Steenbergen, J., Broer, J., Wolters, B., Postma, M.J., Hak, E., 2010. Which factors are important in adults' uptake of a (pre)pandemic influenza vaccine? Vaccine 28, 207–227.