Journal of Health Economics 94 (2024) 102866

Contents lists available at ScienceDirect

Journal of Health Economics

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

Effects of nurse visit copayment on primary care use: Do low-income households pay the price?

Tapio Haaga [a,b,*,1], Petri Böckerman [c,d], Mika Kortelainen [a,b], Janne Tukiainen [a]

[a] Turku School of Economics, University of Turku, FI-20014, Finland
[b] Finnish Institute for Health and Welfare (THL), P.O. Box 30, FI-00271 Helsinki, Finland
[c] Jyväskylä University School of Business and Economics, University of Jyväskylä, P.O. Box 35, FI-40014, Finland
[d] Labour Institute for Economic Research LABORE, Arkadiankatu 7, FI-00100 Helsinki, Finland

ARTICLE INFO

JEL classification:
I18
I14
I13
H42
I11

Keywords:
Cost-sharing
Copayment
Primary care
Difference-in-differences
Pre-analysis plan
Blind analysis

ABSTRACT

Nurses are increasingly providing primary care, yet the literature on cost-sharing has paid little attention to nurse visits. We employ a staggered difference-in-differences design to examine the effects of adopting a 10-euro copayment for nurse visits on the use of public primary care among Finnish adults. We find that the copayment reduced nurse visits by 9%–10% during a one-year follow-up. There is heterogeneity by income in absolute terms, but not in relative terms. The spillover effects on general practitioner (GP) use are negative but small, with varying statistical significance. We also analyze the subsequent nationwide abolition of the copayment. However, we refrain from drawing causal conclusions from this due to the lack of credibility in the parallel trends assumption. Overall, our analysis suggests that moderate copayments can create a greater barrier to access for low-income individuals. We also provide an example of using a pre-analysis plan for retrospective observational data.

1. Introduction

An aging population puts pressure on primary care system that provides comprehensive and accessible services as early and as close to the patient as possible. One solution to cope with this pressure is cost-sharing policies that aim to curb wasteful health care utilization. Therefore, universal healthcare systems often charge moderate-sized copayments to limit demand and collect revenue. While smaller copayments reduce the financial risks for patients, it is notable that even they can have significant effects on service utilization (Iizuka and Shigeoka, 2022). As needs-based prioritization by primary care professionals is conditional on patients having initiated the contact, cost-sharing may put more screening responsibility on the patients themselves. However, the patients' consumption choices can be far from optimal in health care (Chandra et al., 2023).

A key concern is that copayments constitute a barrier to care that disproportionately affects low-income patients, potentially leading to a reduction in valuable care and contributing to health inequalities. Given the potential for large utilization effects, policymakers should set copayment levels, whether zero or nonzero, to achieve specific objectives (Iizuka and Shigeoka, 2022). Some countries, such as the UK and Germany, do not charge copayments for primary care GP visits. Several others, such as Ireland,

* Corresponding author at: Turku School of Economics, University of Turku, FI-20014, Finland.
E-mail address: tapio.haaga@utu.fi (T. Haaga).

1 We thank Mikko Peltola for support, and editor Owen O'Donnell, two anonymous reviewers, Gustav Kjellsson, Marius Opstrup Morthorst, Heather Royer, Henri Salokangas, Matti Sarvimäki, Markku Siikanen, Lauri Sääksvuori, Jussi Tervola, and several seminar and conference participants for comments and suggestions.

https://doi.org/10.1016/j.jhealeco.2024.102866
Received 22 March 2023; Received in revised form 16 February 2024; Accepted 24 February 2024
Available online 28 February 2024

Sweden, Denmark, and Finland, or regions within them, have recently abolished copayments of specific entry-level health services for specific vulnerable groups, such as low-income individuals, minors, and the elderly, for better accessibility. These services include visits to primary care nurses and general practitioners (GP) and to psychologists. Despite the vital role nurses play in many primary care systems, nurse visits remain understudied in the cost-sharing literature. Nurses are increasingly substituted for physicians to address physician shortages and improve efficiency in primary care (Maier and Aiken, 2016). This includes tasks such as examination, triage, diagnosis, and treatment of patients.

We estimate the effects of introducing a 10-euro copayment for curative primary outpatient care nurse visits on the use of public primary care among Finnish adults, with a focus on heterogeneous effects by income level. Our empirical strategy was specified in a pre-analysis plan (PAP), written based on a placebo exposure (analysis blinding) and before we linked outcome data with the real policy exposure data (Haaga et al., 2022). We employ a difference-in-differences (DD) design and exploit the staggered adoption of the copayment in primary care areas in 2014–2019, as well as comprehensive administrative data. Most Finnish municipalities adopted the copayment at some point in 2014–2019 to collect more revenue, and the exact timing of the adoption in the treated areas is arguably arbitrary. In Finland, nurses triage primary care patients and book appointments if needed. Many patients with acute infectious diseases or chronic conditions are directed to nurse visits rather than GP visits. Nurses can consult GPs, who write the vast majority of prescriptions, or book GP appointments for their patients if needed.

We find that the copayment adoption reduced the number of curative nurse visits by 9%–10% in public primary care during a one-year follow-up. There is statistically significant heterogeneity by income in absolute terms: the estimated decrease in the number of visits is more than two times larger for individuals at the bottom 40% of the income distribution, who also have a higher baseline service use, than at the top 40%. The effect size increases as income decreases. However, heterogeneity by income level is much weaker and statistically insignificant in relative terms (percentage changes). Moreover, we examine whether nurse visit copayments have spillover effects on GP use. As access to GPs is controlled by gatekeeper nurses, we would expect little or no substitution to GP visits. Instead, GP use may decrease, with large reductions interpreted as concerning: if patients do not contact primary care as often as previously, triages are being missed, and a fraction of them could plausibly have led to GP visits. We estimate a 3%–5% reduction in GP visits, but these estimates may partially capture a decreasing trend. In a robustness check that assumes parallel trends only from the last pre-treatment period on, the point estimates on GP visits are no longer statistically different from zero.

The relevant dimension of heterogeneity by income is context-specific. Our findings suggest that copayments can create a greater barrier to care for low-income individuals in terms of visits. This can be important for politicians who prioritize equality and the well-being of the most disadvantaged and consider cost-sharing as a means to influence the distribution of public resources. On the other hand, equal elasticities, defined in relative terms, may be appropriate for both low-income and high-income individuals when authorities make straightforward predictions about the impacts of cost-sharing changes.

We also conduct post-blind exploratory analyses (not specified in the PAP) to examine the extent to which effects on utilization are likely to be consequential for health. This reveals that those with a prescription for diabetes or hypertension, whose baseline nurse use is much higher, responded more strongly in absolute terms to the copayment than those without, but we find no evidence for differences in relative terms. Second, we observe statistically insignificant and imprecisely estimated increases in emergency department visits and unplanned hospitalizations for ambulatory care sensitive conditions. In short, these results do not offer conclusive evidence supporting either the existence or nonexistence of health effects.

The nurse visit copayment was abolished nationwide in July 2021. We also examine the impacts of this reform using a DD design, as stated in the pre-analysis plan. However, we refrain from drawing any causal conclusions from the abolition policy due to the lack of credibility in the parallel trends assumption based on pre-trend patterns. Using the approach of Rambachan and Roth (2023) (post-blind), we demonstrate that even small post-treatment violations of parallel trends, compared to the observed pre-treatment violations of parallel trends, would make the confidence intervals very wide. In essence, the abolition analyses are imprecise and provide evidence supporting neither the existence nor the nonexistence of effects.

Our analysis relates most closely to studies that analyze the impacts of moderate copayments on primary care use in public health insurance systems covering all citizens. Such studies have evaluated the effects of copayments for GP visits in the Nordic countries among children and adolescents (Nilsson and Paul, 2018; Johansson et al., 2019; Landsem and Magnussen, 2018; Olsen and Melberg, 2018) and the elderly (Johansson et al., 2023), and among Irish children (Nolan and Layte, 2017) and the elderly (Ma and Nolan, 2017). Using German data, Farbmacher and Winter (2013) study the impacts of doctor visit copayments, paid once in each calendar quarter, and focus on adolescents — see also the model-based approaches of Farbmacher et al. (2017) and Kunz and Winkelmann (2017) to account for the nonlinear nature of the copayment schemes. Regarding other services, Han et al. (2020) and Iizuka and Shigeoka (2022) examine the impacts of copayments among small children and adolescents in Taiwan and Japan, while Kruse et al. (2022) study the effects of abolishing a copayment for psychologist treatment among Danish adolescents.

Despite the vital role of nurses in many primary care systems, we are among the first to focus on nurse visits and their copayments. We expect that nurse visits are more responsive than GP visits to copayments of equal size: (i) Nurse visits tend to be more accessible, with less gatekeeping and shorter waiting times. This implies lower indirect costs of seeking care and potentially a larger role for the copayment in affecting demand. (ii) The equal-sized copayment represents a larger proportion of the overall costs for nurse visits, i.e., a higher coinsurance rate. Patients may value GP visits more as GPs have more extensive education, the authority to write prescriptions, and they can make referrals to specialists.[2]

[2] Copayment increases for prescription drugs reportedly reduced doctor visits in Germany (Winkelmann, 2004).

We study a large adult population and both the direct effects on nurse visits and also the indirect effects on GP use. Moreover, we employ a staggered DD design with methods that are robust to heterogeneous effects. These features have several important benefits: (i) the availability of several events reduces the risk of external shocks systematically biasing the estimates, (ii) the short lag between the policy decision and its implementation limits anticipation effects, and (iii) we provide estimates for the larger adult population based on a one-year follow-up instead of individual-level effects local to a specific birthday. To our understanding, we are among the first to use a PAP in a nonexperimental study on cost-sharing and to conduct pre-specified heterogeneity analysis by income level. We use the PAP as a self-control tool to enhance the quality and transparency of our research, for example, by avoiding multiple-testing or hindsight bias. Several earlier studies have examined treatment effect heterogeneity by income but by focusing on narrow age groups, mainly children or adolescents (Nilsson and Paul, 2018; Johansson et al., 2019; Han et al., 2020; Kruse et al., 2022), or people in their mid-eighties (Johansson et al., 2023).

The rest of the paper is structured as follows. Section 2 describes the institutional background and Section 3 the data. Section 4 presents our PAP-based workflow and our empirical approach for the staggered adoption. Section 5 reports the results. Section 6 summarizes the analyses for the copayment abolition. Sections 7 and 8 contain discussion and conclusion. The Online Appendix contains additional figures and tables, post-blind supplementary analyses for the copayment adoption, regression analyses for the copayment abolition, a description of the data construction, and documentation of deviations from the PAP.

2. Institutional background

Public primary outpatient care in Finland. Primary outpatient care is provided for adults by three sectors: public primary care, occupational healthcare, and private clinics. The sectors differ with respect to gatekeeping, out-of-pocket costs, and waiting times. Public primary care is available for all individuals. It is disproportionately important for those who do not have access to occupational healthcare, such as pensioners and the unemployed. Occupational healthcare provides preventive work-related services, but a majority of employees are also offered curative services free-of-charge and with fast access. For most of the employed, occupational healthcare is thus an attractive alternative to public primary care which is characterized by gatekeeping, waiting times that may be long for nonurgent care, and copayments (21 euros per GP visit at maximum and approximately 10 euros per curative nurse visit). Private clinics offer services with fast access and no gatekeeping, but the fees are much higher than in public primary care. The number of curative nurse visits per capita in public primary care is negatively correlated with income decile in our data. Importantly for our study, there is still use of public primary care at the top of the income distribution, even in the working age (aged 25 to 64) population.

We focus on the potential impacts of the nurse copayment on the use of public primary care and do not account for potential spillovers to the use of occupational or private healthcare for data availability reasons. There are reasons to believe that these spillovers are small. Rinne and Blomgren (2023) show, using a matched DD design and data from the city of Oulu, that job loss reduces the use of occupational healthcare, with no impacts on the use of public or private healthcare services. We argue that the spillovers on private clinics are small: (i) those who respond to a moderate 10-euro copayment likely do not switch to private services that are many times costlier, and (ii) the network of private clinics is geographically not as accessible as the network of public primary care (Lavaste, 2023).

In our study period, municipalities formed publicly funded primary care areas (health centers) on their own or in cooperation with others. Every citizen has their designated health station determined by where they live. In some primary care areas, all health stations may be available on a visit-by-visit basis. Since 2014, citizens have been able to choose their health station once a year, but active choices have not been common. Municipal services are financed through state transfers, municipal taxes, copayments, and borrowing. The state guides copayment policies by setting which groups or services are exempted (based on the Social and Health Care Client Fees Act) and maximum copayment levels (in the corresponding government decree). Within these constraints, primary care areas set their own policies.

The public system is characterized by limited supply and labor shortages. Cohort sizes in medical schools are fixed, and the public and the private sector compete for doctors. Primary care areas face challenges in hiring nurses at the prevailing wage level determined by sector-based collective bargaining. The challenges are reinforced by the fact that central and local governments have been running fiscal deficits for years, which is expected to continue.

The vital role of nurses in public primary outpatient care in Finland. Patients can contact public health stations via phone (the primary channel) or visit in person. Upon contact, patients undergo triage, a process primarily handled by nurses or public health nurses who then schedule appointments with professionals if necessary. As a result of task-shifting from GPs to nurses, an increasing number of patients are directed toward nurse visits instead of GP visits. Nurses offer a range of services including urgent care for conditions such as flu and stomach flu and monitoring visits for patients with chronic conditions, such as diabetes, hypertension, COPD, asthma, and dementia. Nurses can also provide wound care, administer drug injections and infusions, perform ear irrigation, and direct patients to GPs. Specialist visits in public healthcare require a doctor's referral, and most prescriptions are written by doctors.[3] A recent trend in work arrangements is that the triaging nurses not only guide patients to the appropriate professional but try to resolve the cases immediately. This often involves real-time consultation with a GP based on nurse–GP pairs or larger multidisciplinary teams, without the need for the patient to see the GP directly.

3 Less than 700 nurses had restricted prescription rights in 2023, and they can describe drugs to treat, for example, urinary tract infections or tonsillitis.

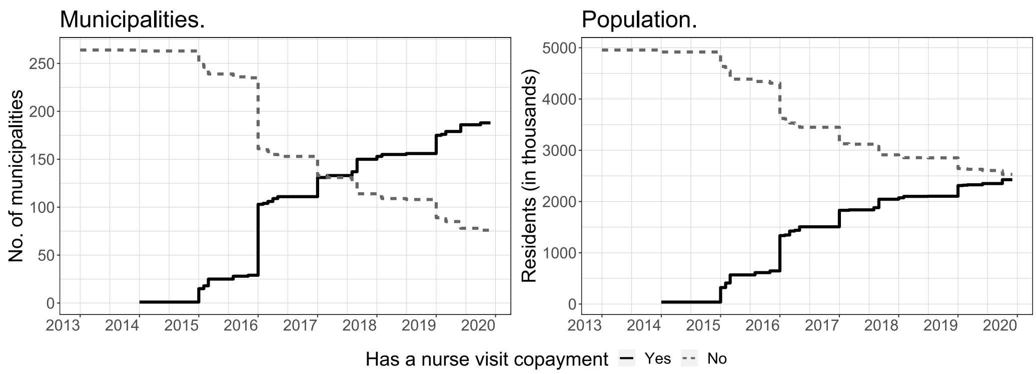

Fig. 1. Staggered adoption of the nurse visit copayment. *Notes:* The figure illustrates the staggered adoption of the nurse visit copayment in terms of the number of exposed municipalities and individuals. We take municipalities in mainland Finland (293 in 2022) and use the 2022 municipal boundaries and population sizes from the end of 2019. The sample contains those municipalities whose policies on copayments for curative nurse visits we observe in our data collection.

In a cross-country comparison that includes Europe, the USA, Canada, Australia, and New Zealand, Finland has been classified as part of a group of countries where there has been "extensive" task shifting from GPs to nurses in primary care (Maier and Aiken, 2016). This group also includes the USA, the UK, Canada, Australia, Ireland, the Netherlands, and New Zealand. Moreover, as of 2019, Finland had the second highest ratio of nurses to doctors among OECD countries (OECD, 2021), with all the countries identified by Maier and Aiken (2016) as having "extensive" task shifting surpassing the OECD average. Regarding curative visits at the municipal level in Finland, utilization of GP visits shows a strong correlation with nurse use, while there still is noticeable variation in the ratio of nurse visits to GP visits (Figure A1, post-blind, i.e., not included in the PAP).

Copayments for nurse visits. Finland adopted restricted prescription rights for nurses in 2010. Related to that, the law on copayments was changed to allow primary care areas to charge a copayment for curative nurse visits, no longer specifying professions (e.g., physicians) whose visits can be subject to copayments. However, the decree continued to mention explicitly only doctor visits. This confusion likely explains why no areas immediately adopted the nurse visit copayment, first introduced in 2014. Many other areas adopted it to collect more revenue once they became aware of the possibility.

The staggered adoption is illustrated in Fig. 1. At the end of 2019, half of the population lived in areas charging the copayment, with the vast majority of municipalities charging it. The set of municipalities that adopted the copayment is geographically diverse (Figure A2). In 2021, 80% of the municipalities with the copayment charged it for three visits annually, and by far the most common per-visit copayment was 11 euros, the population-weighted mean being 12 euros (Figure A3). No significant area-specific changes have been made to the levels after the adoption, except for minor inflation adjustments.

The nurse visit copayment was later abolished nationally in July 2021 when the national government conducted a reform to the act on copayments to reduce barriers to access and health inequalities. The reform included no major changes to GP visit copayments. More than 200 municipalities, on average disproportionately small and rural, and almost three million people were affected by the nurse visit copayment abolition (Figure A4 ; Figure A5 ; Table A1, post-blind).

Several policies protect financially vulnerable patients from healthcare costs. There are annual out-of-pocket caps for public healthcare services and prescription drugs of 692 euros and 592 euros, respectively (in 2022). Households with the lowest incomes and only little wealth can apply for social assistance, a means-tested last-resort benefit, which can also cover out-of-pocket costs for public health care and prescription drugs. The law on copayments requires that, for some public services, financially vulnerable patients can apply for an exemption or a lowered copayment. This right does not apply to nurse visit copayments, but some areas may still exempt individuals based on applications. A few primary care areas provide exemptions to specific low-income groups, such as those with the lowest national pension or those receiving social assistance.

3. Data

We combine several Finnish national administrative registers using person IDs. The data contain contacts in publicly funded primary care and hospitals, social assistance recipients, drug prescriptions since 2018, and socioeconomic characteristics of all individuals who have a permanent residence in Finland at year's end.[4] We observe age and the municipality of residence, which are used to link visits to copayment policies. We construct a variable for equivalized family disposable income for each individual and calculate population sizes for each municipality.

We also use publicly available data on each municipality's primary care area in 2021 (from the Association of Finnish Municipalities). Two publicly available registers listing social and healthcare organizations are linked to primary care contacts

[4] The data on primary care use (Register of Primary Health Care Visits), specialized healthcare use (Care Register for Health Care), and social assistance recipients (Register of Social Assistance) are all administered by the Finnish Institute for Health and Welfare (THL). The socioeconomic data comprise Statistics Finland's FOLK modules "basic", "family", and "income". Prescriptions (Kanta Prescription Center) are administered by the Social Insurance Institution of Finland.

(THL). We create three tables mapping areas to copayment policies. The first reports whether a given municipality had adopted the nurse visit copayment by the end of 2019 and the possible adoption date. These data were collected from municipal documents, websites, and news in local media. The search was based on the publicly available dataset on nurse and GP visits copayments (from THL), which we also use for GP visit copayments in 2013–2018 (the second table). The third table reports the copayments in the summer of 2021, collected from the websites of primary care areas.[5]

The analysis for the staggered adoption of the copayment uses data from 2013–2019, restricting to pre-pandemic (COVID-19) years. We include those individual-by-year observations in which the person is 25 years or older. The aim is to exclude minors, who are exempted from the copayment, and students, who have access to student healthcare. The primary outcomes are the annualized number of curative nurse and GP visits per capita in publicly funded primary care, constructed by multiplying monthly visits per capita by 12. The secondary outcomes include the share of individuals receiving social assistance and the annual sum of received basic social assistance, both defined at the family level.[6] Our pre-specified choice is to estimate the effects separately in all analyses for individuals at the bottom 40% and the top 40% of the equivalized family disposable income distribution, constructed across the entire eligible population.[7]

We discuss in detail how we clean and construct our analysis data in Online Appendix Section A.3. Ultimately, we have an unbalanced panel at the municipality-by-time-period-by-income-decile-by-outcome level. Time period is a month except for the annually measured sum of social assistance. The panel is unbalanced because we exclude some observations due to quality problems, mainly for primary care outcomes. When the national data collection started in 2011, not all areas were able to transfer primary care data from their electronic health record (EHR) systems to the national register. Later changes in the EHR systems may also be visible in the data as a sudden but short drop to a near zero value in aggregate contacts. The details of how we detect and exclude observations with data quality concerns are provided in Online Appendix Section A.3.

The analysis of the nationwide abolition of the nurse visit copayment in July 2021 is based on 12 pre-treatment and 11 post-treatment months (from 7/2020 to 5/2022), requiring a balanced panel. Socioeconomic data are from 2020. The observation window excludes early 2020 when the supply and demand shocks caused by the COVID-19 pandemic were largest. For data construction, we assume that the effects of the earlier adoption of the copayment fully accumulate within one year, excluding municipalities that adopted the copayment less than 12 months before the start of the observation window.

4. Empirical approach for the staggered adoption

Pre-analysis plan (PAP). We use a PAP and analysis blinding with retrospective observational data in the absence of a fraud-proof firewall between planning and analysis (Haaga et al., 2022). Besides being a self-control tool for us, the workflow provides additional transparency to readers. For the analysis of the staggered adoption of the copayment, we had all data available when designing the analyses. Our outcomes and the policies on nurse visit copayments had not been linked previously, which allowed us to blind the causal relation of interest for the PAP. Specifically, we randomly assigned municipalities into placebo policies, using the real adoption dates but randomly assigning municipalities into them, and specified our statistical approach and demonstrated the reporting of the results in the PAP without observing the real results.[8] At the time of writing the PAP, we did not yet have access to microdata on 2021–2022 (post-treatment) outcomes for the copayment abolition analysis. We wrote the statistical programs as if the abolition occurred years earlier on July 1st, 2018. Some aggregated post-treatment outcome data were publicly available at the municipality level; however, these data did not influence our work.

In practice, our PAP included a placebo report based on blinded data, mimicking the structure of the research paper, and the corresponding statistical codes (Haaga et al., 2022). The benefits of using blinded data for PAPs have previously been discussed by Olken (2015) and Nosek et al. (2018). PAPs should not prevent being flexible to reasonable changes *ex post*. Deviations from plans are common, and PAPs can still provide substantial benefit if researchers transparently report all changes and their reasons (Nosek et al., 2018). We document and discuss our changes in Online Appendix Section A.4. This article is a combination of pre-registered confirmatory analyses and post-blind exploratory and supplementary analyses. We separate tables and figures that were not pre-registered by using the label "post-blind" in the text and in the table and figure notes.

One deviation from our PAP is worth highlighting here. Our plan was to examine and report both the staggered adoption and the later abolition of the copayment with equal weight. The final paper focuses on the staggered adoption because its design turned out to be much more credible, and we moved parts of the abolition analyses to the Online Appendix. For the abolition, we could not credibly assume parallel trends, the essential identification assumption, based on pre-treatment outcomes observed *ex post*.

Research design. We use a staggered difference-in-differences (DD) design with an irreversible treatment. For each event, we have both never-treated and later-treated municipalities as controls.[9] The never-treated areas include the six largest cities and differ

[5] We thank Katja Ilmarinen, who had gathered the same information independently, for allowing us to cross-validate our information.

[6] Social assistance is a means-tested last-resort benefit for households. Higher copayments may increase the need for social assistance.

[7] We focus on two groups for parsimony. Using smaller groups than the bottom 40% and the top 40% has two disadvantages: smaller samples and larger variation, and the fact that the share of social assistance recipients is larger at the bottom of the income distribution, potentially attenuating estimates as the benefit can cover copayments.

[8] We had previously examined the impacts of GP visit copayments using similar data from the 2010s by (1) focusing on exempted minors (Haaga et al., 2023), (2) exploiting the abolition of the GP visit copayment in one large municipality in 2013 (Haaga et al., 2024), and (3) conducting exploratory analyses on the impacts of GP visit copayment increases.

[9] To be specific, copayment policies are set at the primary care area level. However, we analyze the data at the municipality level for practical purposes.

from the treated areas. We do not view the decision to adopt the nurse visit copayment as quasi-random. However, the decision on when to adopt conditional on adopting seems much more arbitrary. Our interpretation of municipal decision-making is that there is potential randomness in the timing of when public servants became aware of the possibility to charge the copayment and, consequently, in the treatment timing.[10]

In staggered settings, two alternative identification assumptions are relevant: a model-based parallel trends assumption (PTA) and a stricter design-based assumption of (quasi-)random treatment timing. The causal interpretation of our main analysis relies on the validity of the PTA, i.e., we assume that the outcomes for the treated cohort and for the comparisons (not-yet-treated or never-treated) would have followed parallel trends in the absence of treatment. The PTA may hold in both levels and logs, or in levels but not logs, or *vice versa* (Roth and Sant'Anna, 2023b). As stated in our PAP, we report the main results on primary care use in both levels and logs.[11] Both models seem plausible based on pre-treatment trends in event-study plots reported in Section 5.[12] However, similar pre-trends are not sufficient for the PTA to be valid, and the PTA is inherently untestable. Another caveat is that in DD designs it is not possible to rule out with complete certainty the influence of other exposures concurrent with the intervention of interest. If they exist, we can only identify the net effect. However, the staggered design (several events) reduces the risk of external shocks systematically biasing our results.

Still, two potential threats to identification should be noted. First, the fiscal challenges that led many municipalities to adopt the nurse visit copayment may also have led to other concurrent policy changes affecting health care use. Most importantly, GP visit copayments could have been increased simultaneously. The central government increased the maximum GP visit copayment from 16.10 euros in 2015 to 20.90 euros in 2016. Municipalities responded differently: many made the increase instantly in 1/2016, some made it later, and some not. However, we find that the nurse visit copayment adoption was correlated with only a 1-euro increase in GP visit copayments (small in both absolute terms and percentages) based on the methods of our main analysis (Figure A6, post-blind). We chose to model the setting as a single-treatment design in line with our PAP, not controlling for GP visit copayments.[13] We are not aware of anecdotal evidence for supply reductions motivated by cost savings, such as health station closures, shorter opening times, or staff reductions, concurrent to the copayment adoption. If they existed, we would expect to detect rather similar reductions in both nurse and GP visits in the main analysis, which is not the case.

Second, once the copayment is adopted, preventive and curative visits must be distinguished for charging purposes. If this affects how contacts are recorded, the number of recorded visits may change even if the underlying use does not. However, we find no evidence of preventive-labeled visits crowding out curative-labeled visits: adopting the copayment for curative nurse visits was not correlated with the number of preventive nurse visits based on the methods of our main analysis (Figure A7, post-blind).

We do not have to account for utilization spillovers across areas due to Finland's publicly funded primary care system where nonurgent care is provided by a designated health station determined by the location of residence. We do not expect noticeable anticipation effects. The implementation time from the political decision is often a month or less, and few citizens likely pay much attention to the minutes of the municipal committees.

Econometric methods. We use the stacked regression estimator (Gormley and Matsa, 2011; Cengiz et al., 2019) as our baseline. For robustness checks, we use the Callaway and Sant'Anna (2021) (CS) estimator. Both estimators are robust to biases in conventional two-way fixed effects (TWFE) regression models caused by staggered treatments and treatment effect heterogeneity (Baker et al., 2022) and thus were attractive *ex ante* when designing the PAP.

In practice, stacking ensures that earlier-treated units are not used as controls for later-treated units. It transforms the staggered setting into event-specific datasets that are ultimately stacked (or pooled) together before conventional TWFE regressions are fitted. We first create a separate dataset for each event, including the treatment cohort and all clean controls that are unexposed (not-yet-treated) in the window of 24 months before and 12 months after the copayment adoption. We only use data from the 36-month window and include events with at least 12 post-treatment months. Depending on the outcome, we exclude several municipality-year observations due to data quality concerns (see Online Appendix Section A.3). We require balanced panels in event (or relative) time as the baseline. These event-specific datasets are stacked for estimation. Our static TWFE specification includes event-specific unit and time fixed effects:

$$y_{mte} = \alpha_{me} + \gamma_{te} + \delta^{DD} D_{mt} + \varepsilon_{mte}. \tag{1}$$

Here, subscripts m, t, and e denote municipality, month, and event-specific datasets, respectively, and α_{me} and γ_{te} represent event-specific municipality and calendar month fixed effects. D_{mt} is a dummy for post-treatment periods in the treated municipalities. We weight by population due to heterogeneity in municipality size. Standard errors are clustered by municipality.[14] We estimate the model separately for individuals at the bottom 40% and the top 40% of the distribution of equivalized family disposable income.

We complementarily use two other stacked specifications modified from Model (1). Model (1.1), a dynamic event-study version of Model (1), is used for event-study plots:

$$y_{mte} = \alpha_{me} + \gamma_{te} + \sum_{l=-24, l\neq -1}^{11} \mu_l D^l_{mte} + \varepsilon_{mte}. \tag{1.1}$$

[10] The earlier-treated and later-treated municipalities, defined by median event date, are rather similar in some key features (Table A1, post-blind).

[11] The results in levels are shown first while most of the results in logs are documented in the Online Appendix.

[12] Complementarily, we will estimate models that allow for smooth (linear) pre-trend differences and models that do not use the never-treated municipalities.

[13] With this choice, we likely trade off some bias for lower variance and more external validity compared to using the estimator proposed by de Chaisemartin and D'Haultfœuille (2023).

[14] Our stacked data contain multiple copies of the same observation, as a municipality can belong to several event-specific datasets. Thus, clustering at the municipality-by-event level is not appropriate.

The coefficients of interest, μ_l, represent all leads and lags ($t = -1$ omitted as a reference). D^l_{mte} is a dummy for the treated areas for observations l months from the copayment adoption in the event-specific dataset e. Both relative time l and the treated are defined by e.

Next, Model (1.2) allows for a linear pre-trend difference between the treated and the comparisons in contrast to Model (1). Its PTA concerns deviations from a linear pre-trend difference: Model (1.2) is preferable to Model (1) if there exists a linear pre-trend difference that would have continued in the absence of treatment.

$$y_{mte} = \alpha_{me} + \gamma_{te} + \theta d_{me} t_e + \sum_{l=0}^{11} \mu_l D^l_{mte} + \varepsilon_{mte}. \tag{1.2}$$

The coefficients of interest, μ_l, represent all lags. D^l_{mte} is a dummy for the treated areas for observations l months from the copayment adoption in the event-specific dataset e, as in Model (1.1). Regarding the linear pre-trends, d_{me} is a dummy for the municipality being treated in the event-specific dataset e, and t_e denotes time relative to that event. We report the mean of μ_l over l as our point estimate.

The stacking estimator is efficient as it uses OLS to derive weights on the event-specific DD estimates, trading off bias for efficiency. However, the use of variance weighting may lead to inconsistency for the *sample-average* ATT (average treatment effect on the treated) (Baker et al., 2022). In fact, Gardner (2021) shows that the estimator identifies an average of event-specific ATTs, weighted by event-specific treatment variance and sample size. For the PAP, we valued the simplicity and implementability of stacking and its ability to accommodate triple difference models for testing treatment effect heterogeneity.

As an alternative to stacking, we use the CS estimator (Callaway and Sant'Anna, 2021). The aim is to identify a group-time average treatment effect, allowing for treatment effect heterogeneity over cohorts and time. The group-time ATTs can be aggregated to construct measures of overall treatment effects. We provide both event-study-type estimates and a static estimate that is the average of all group-time ATTs, weighted by group size.[15] The authors propose several two-step plug-in estimators for group-time ATTs: first estimate nuisance functions and then plug their fitted values into the sample analogue of the group-time ATT. When the never-treated units are used for comparison, the PTA is assumed only from the last pre-treatment period on.[16] We use outcome regression, weight by population, and cluster standard errors by municipality. Events with at least 12 follow-up months are included. The dataset is balanced in calendar time, excluding municipalities with data quality concerns in the study window. Our baseline is to exclude the years 2013 and 2019 when analyzing primary care use to increase the number of sample municipalities.[17] Regarding social assistance use, we use all data from 2013–2019. The data are aggregated to the municipality-by-time-period level for estimation.

Assessing treatment effect heterogeneity. We focus on whether the copayment had heterogeneous utilization effects by income. The analysis compares the bottom 40% of the income distribution (equivalized family disposable income) to the top 40% throughout the study, as defined in our PAP. First, we estimate the effects separately for individuals at the bottom 40% and the top 40% of the income distribution by using the DD framework described above to (a) visually illustrate the design and results and (b) to estimate the magnitude of the effects separately for income groups.

Second, we use a triple difference (DDD) model with the stacked data to formally test whether there is heterogeneity in effects by income. In the DDD model, we compare the evolution of outcomes at the bottom 40% of the income distribution to that at the top 40% in both treatment and comparison areas. The PTA is now assumed in ratios (Olden and Møen, 2022), concerning the relative outcomes of the income groups. We use the following specification:

$$\begin{aligned} y_{mgte} = {} & \alpha + \beta_{1e} Treat_{me} + \beta_{2e} Bottom40_{ge} + \beta_{3e} Post_{te} + \beta_{4e} Treat_{me} \times Bottom40_{ge} \\ & + \beta_{5e} Treat_{me} \times Post_{te} + \beta_{6e} Bottom40_{ge} \times Post_{te} \\ & + \gamma Treat_{me} \times Bottom40_{ge} \times Post_{te} + \varepsilon_{mgte}. \end{aligned} \tag{2}$$

Here, subscripts m, g, t, and e denote municipality, income group, time (month), and event-specific datasets, respectively. $Treat_{me}$ and $Post_{te}$ are two sets of indicators. $Treat_{me}$ indicates whether municipality m is used in the treatment group in the dataset e. Consequently, the never-treated municipalities always have $Treat_{me} = 0$, but the later-treated can get either 1 or 0 for $Treat_{me}$ depending on whether they are used in the treatment group or in the comparison group for the earlier treated in the dataset e. $Post_{te}$ indicates the calendar months following the treatment month in the dataset e. $Bottom40_{ge}$ is a dummy for the bottom 40% of the income distribution (0 for those in the top 40%), and γ is the coefficient of interest. Other coefficients (from β_{1e} to β_{5e}) are event-specific. We again weight by population size and cluster standard errors by municipality.

[15] Our stacking analyses use a 12-month follow-up. For the CS estimator, follow-up varies by treatment group. This also implies putting more weight on the earlier-treated cohorts.

[16] Thus, the assumption does not restrict pre-treatment trends. However, the PTA is different and restricts pre-trends when the not-yet-treated are used for comparison (Callaway and Sant'Anna, 2021).

[17] The exclusion of 2013 trades off one event and 12 months of data for a larger number of municipalities. In many cases, the primary care data quality concerns reported in Online Appendix Section A.3 occurred early in the panel. The exclusion of 2019 leads us to keep one large never-treated municipality that changed its EHR system in Spring 2019.

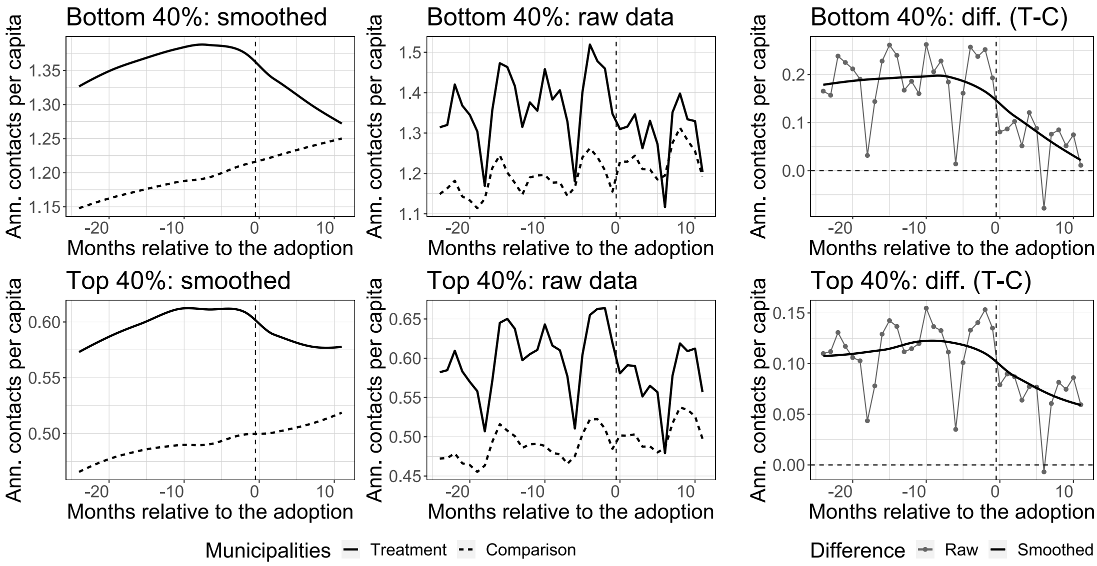

Fig. 2. Adoption: Evolution in nurse visits. *Notes:* The figure shows monthly evolution in the annualized number of curative nurse visits per capita separately at the bottom 40% and the top 40% of the distribution of equivalized family disposable income. Treatment municipalities adopted the nurse visits copayment at time 0 in relative time. We divide the number of visits by individuals in income group G in municipality M in month T by the number of individuals in income group G living in municipality M, multiply the monthly rates by 12 to get annualized utilization, and average over municipalities by weighting by population. The dataset is stacked, and event-specific datasets balanced. Sample sizes: 17 events and 245 municipalities of which 152 are ever treated. The left column contains smoothed conditional means, fitted with local linear regression. The raw data is illustrated in the middle column, while the difference between treatment and comparison areas is depicted in the right column. The observed reductions in nurse use occurring every twelve months are likely explained by summer holidays and reduced supply in July, which appears to disproportionately affect the treated municipalities that are smaller on average. Fig. 1 shows that January was a common adoption month.

5. Results: Staggered adoption

5.1. Primary outcomes

Trend and event-study plots. Fig. 2 plots the trends in curative nurse visits in public primary care for individuals at the bottom 40% and the top 40% of the income distribution in treatment and comparison municipalities based on the stacked dataset. The utilization rates of income groups vary significantly, with the bottom 40% exhibiting more than double the utilization compared to the top 40%. We find that nurse visits decreased in the treated municipalities after the adoption of the copayment compared to the comparison municipalities. The decrease was 0.10–0.15 annualized visits at the bottom 40% of the income distribution and approximately 0.05 visits at the top 40%. Nurse visits were increasing in both policy groups before the copayment adoption, and the PTA is arguably plausible. After the adoption, the growth continued in the comparison municipalities, but nurse visits decreased in the treated municipalities. The effects on GP visits, in contrast, are small or zero, and there may be a small decreasing pre-trend in GP visits in the treated areas relative to the comparison municipalities (Figure A8).[18]

We also estimate dynamic event-study regressions using the stacked data (Model (1.1)), comparing the evolution of outcomes between the treated and unexposed municipalities in Fig. 3. Consistent with the trend plots, nurse visits decreased in the treated municipalities after the copayment adoption compared to the comparison municipalities. The potential effects on GP visits are negative but close to zero. The reduction in nurse visits is larger at the bottom 40% of the income distribution in absolute terms. However, we do not find such difference in relative terms using logarithmic outcomes (Figure A9, post-blind).

Main difference-in-differences results. Table 1 reports static treatment effect estimates of the impact of introducing the nurse visit copayment on annualized nurse and GP visits per capita using DD and under the PTA (Model (1)). The estimates show statistically significant decreases in both nurse and GP visits in both the bottom 40% and the top 40% of the income distribution. The nurse copayment reduced the number of nurse visits per capita by 0.13 ($SE = 0.03$), equivalent to a 9.3% reduction, in the bottom 40% and 0.05 ($SE = 0.02$) visits per capita, equivalent to a 8.0% reduction, in the top 40%. The estimated impacts on GP visits per capita are closer to zero than on nurse visits, in both absolute and relative terms. GP visits decreased by 0.06 ($SE = 0.02$), equivalent to a 3.9% reduction, in the bottom 40% and by 0.03 visits ($SE = 0.01$), equivalent to a 4.7% reduction, in the top 40%.

[18] For GP visits, we consequently prefer the stacked specification that allows for a linear pre-trend difference or the CS estimator with the never-treated units as comparisons that assumes parallel trends only from the last pre-treatment period on. These estimates should be closer to zero than the estimates from our baseline stacked specification, assuming parallel trends in every period.

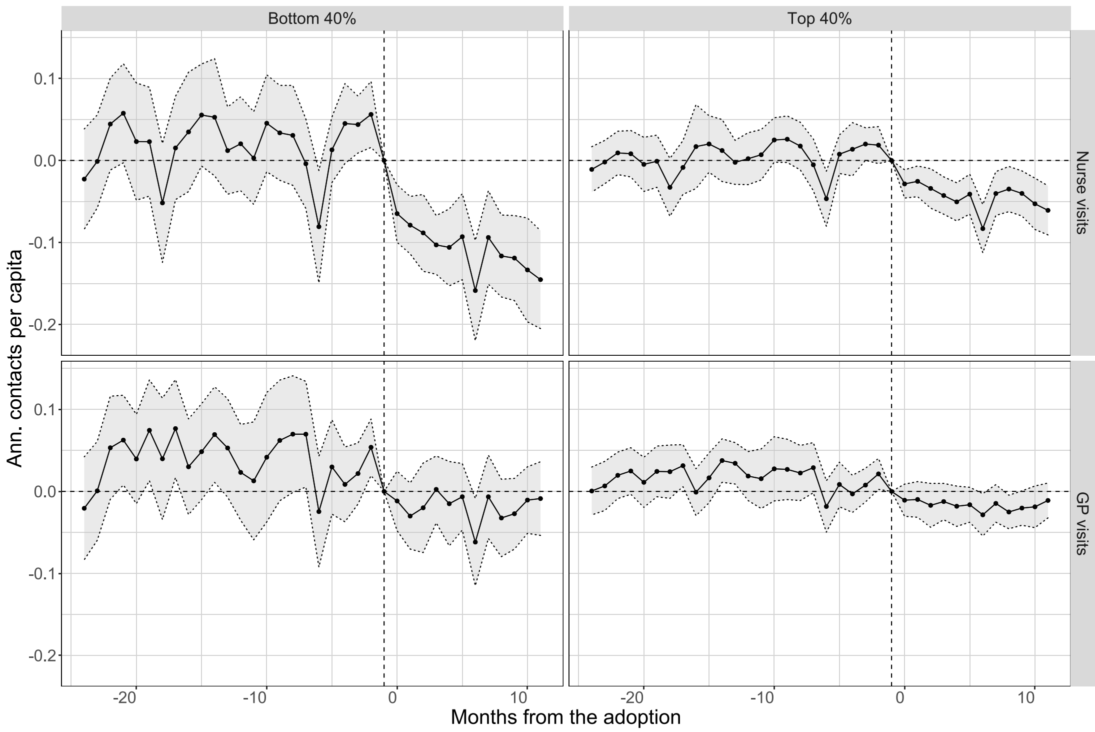

Fig. 3. Adoption: Event-study plot on primary care visits with stacked data. *Notes:* The estimates are from Model (1.1), representing how much the annualized number of contacts per capita of the treatment group differs from the unexposed municipalities as a function of time relative to the copayment adoption. We compute the estimates separately for individuals at the bottom 40% and the top 40% of the distribution of equivalized family disposable income. The dataset is stacked, and event-specific datasets balanced. Due to heterogeneity in municipality size, we weight by population size. Standard errors are clustered by municipality.

These effect estimates are robust to using an alternative model that allows for a linear trend difference (Model (1.2)), but as the estimated effects on GP visits somewhat attenuate, the difference in point estimates between nurse and GP visits slightly increase.

With respect to potential heterogeneity in effects by income, the previous estimates appear heterogeneous in absolute terms — the lower end of the income distribution responds more. However, given the much higher baseline utilization of low-income individuals, there is no substantial and consistent heterogeneity in relative terms. Our PAP did not explicitly specify which of the two dimensions of heterogeneity, absolute or relative, is more relevant. Accordingly, we highlight the difference in estimates.

Effect heterogeneity by income level. We apply a triple difference (DDD) model (Model (2)) using the stacked data to formally test treatment effect heterogeneity by income, comparing the evolution of public primary care use at the bottom 40% of the income distribution to that at the top 40% both in the treatment and comparison areas. In Table 2, we estimate that the nurse visit copayment reduced the number of annualized nurse visits per capita by 0.07 to 0.08 visits more in the bottom 40% than in the top 40%. The difference is statistically significant and equivalent to a reduction of 5.3% to 5.6% when compared to the baseline in the bottom 40%. GP visits decrease more at the bottom 40%, but the difference is statistically insignificant. In relative terms (logs), all the estimated differences in effects between the income groups are statistically insignificant.

Complementarily, we plot the effects of the nurse visit copayment on public primary care use by income decile in Fig. 4 (post-blind) using stacking with balanced event-specific datasets (Model (1)). The pattern is clear in absolute terms: the effect estimate attenuates as income increases. The bottom 10% is an exception. A plausible explanation is that the decile has the highest share of individuals receiving social assistance, a last-resort benefit that can be used to cover copayments. Apart from the first decile, there may be heterogeneity also in relative terms: the lower end of the income distribution is somewhat more sensitive than the top.

Estimates for all individuals (post-blind). We also provide the DD results for the entire sample population (Online Appendix Section A.1) . To summarize these findings, we estimate that the nurse visit copayment reduced the annual number of nurse visits per capita by 0.09–0.10 visits, equivalent to a reduction of 9%–10%, using stacking and a one year follow-up. The CS estimator, that averages over all group-time ATEs and thus has a follow-up longer than a year for most units, yields larger estimated reductions: 0.13–0.16 visits, or 13%–17% with logarithmized outcomes. Stacked effect estimates show a reduction of 3%–5% in GP visits, but the corresponding CS estimates are closer to zero and statistically insignificant.

5.2. Supplementary analyses

Estimates on social assistance use. Adopting the nurse visit copayment could have caused financial hardship for the poorest, increasing the need for social assistance which is a means-tested last-resort benefit. As secondary pre-registered outcomes, we

Table 1
Adoption: Difference-in-differences estimates of effects of nurse visit copayment on primary care visits.

Metric	Nurse visits		GP visits	
	Bottom 40%	Top 40%	Bottom 40%	Top 40%
Baseline mean	1.373	0.603	1.450	0.705
A. Baseline estimates (Model (1))				
Effect	−0.127	−0.048	−0.056	−0.033
SE	0.032	0.018	0.021	0.010
p-value	0.000	0.008	0.007	0.001
Relative effect (%)	−9.252	−8.020	−3.879	−4.669
B. Alternative estimates with linear trends (Model (1.2))				
Effect	−0.131	−0.057	−0.048	−0.026
Relative effect (%)	−9.565	−9.490	−3.298	−3.706
Events	17	17	17	17
Treated areas	152	152	152	152
All areas	245	245	245	245

Notes: The table contains our baseline stacked difference-in-differences results for the impact of introducing the nurse visit copayment, estimated separately for the bottom 40% and the top 40% of the distribution of equivalized family disposable income. Outcomes are the annualized number of curative nurse and GP visits per capita. Baseline means are the average of pre-treatment outcomes of the treated. Relative effects are calculated by dividing effect estimates by baseline means and multiplying by 100. The dataset is stacked, balanced, and contains a one-year follow-up. Our baseline model is Model (1), but we also use its extension (Model (1.2)) allowing for a differential linear pre-trend. Due to heterogeneity in municipality size, we weight by population size. Standard errors are clustered by municipality.

Table 2
Adoption: Triple differences estimates assessing heterogeneity in effects of nurse visit copayment between bottom 40% and top 40% of income distribution.

Metric	Balanced datasets		Unbalanced datasets	
	Nurse visits	GP visits	Nurse visits	GP visits
A. Annualized contacts per capita				
Difference in effects	−0.073	−0.025	−0.078	−0.021
SE	0.017	0.013	0.017	0.013
p-value	0.000	0.058	0.000	0.115
Baseline mean (Bottom 40%)	1.373	1.450	1.386	1.444
Relative difference (%)	−5.304	−1.718	−5.635	−1.462
Events	17	17	19	19
Treated areas	152	152	175	175
All areas	245	245	264	264
B. Logarithmized annualized contacts per capita				
Difference in effects	−1.544	0.711	−1.575	0.818
SE	1.559	0.903	1.523	0.925
p-value	0.322	0.431	0.301	0.377
Events	17	17	19	19
Treated areas	126	135	175	175
All areas	209	225	264	264

Notes: The table contains our stacked triple differences results from Model (2), formally testing for a difference between the bottom 40% and the top 40% of the distribution of equivalized family disposable income in effects of introducing the nurse visit copayment. Outcomes are the annualized number of curative nurse and GP visits, or their logarithm. In Panel A, the baseline mean is an average of pre-treatment outcomes of individuals residing in the treated areas and being at the bottom 40% of the income distribution. Relative difference compares the DDD estimand to this baseline mean. In Panel B, estimates for logarithmized outcomes are multiplied by 100. The dataset is stacked, and event-specific datasets are either balanced or unbalanced, containing a one-year follow-up. Due to heterogeneity in municipality size, we weight by population size. Standard errors are clustered by municipality.

examined the impacts of the copayment adoption on (1) the share of individuals living in a family where someone received social assistance (extensive margin) and on (2) the annual sum of social assistance received by the family (intensive margin). To summarize, the effects are inconclusive: the point estimates are positive at the intensive margin and negative at the extensive margin but, in both cases, not statistically different from zero (Online Appendix Section A.1).

Potential health effects (post-blind). We examine the heterogeneity of the effects with respect to having received a drug prescription in 2018–2019 with an anatomical therapeutic chemical (ATC) code referring to diabetes or hypertension (A10, C02–C03, and C07–C09), proxying a diagnosis of these conditions. Those with a prescription for diabetes or hypertension responded more strongly in absolute terms to the copayment adoption, but we found no difference in relative terms (Table A2, post-blind). Online Appendix Section A.1 presents the results on emergency department (ED) visits and unplanned hospitalizations for ambulatory care sensitive conditions (ACSC), both outcomes being post-blind. The estimated effects – more exploratory than confirmatory – suggest increases but are imprecisely estimated and statistically insignificant for both outcomes. To summarize, these findings do not offer sufficient evidence either for the existence or nonexistence of health effects.

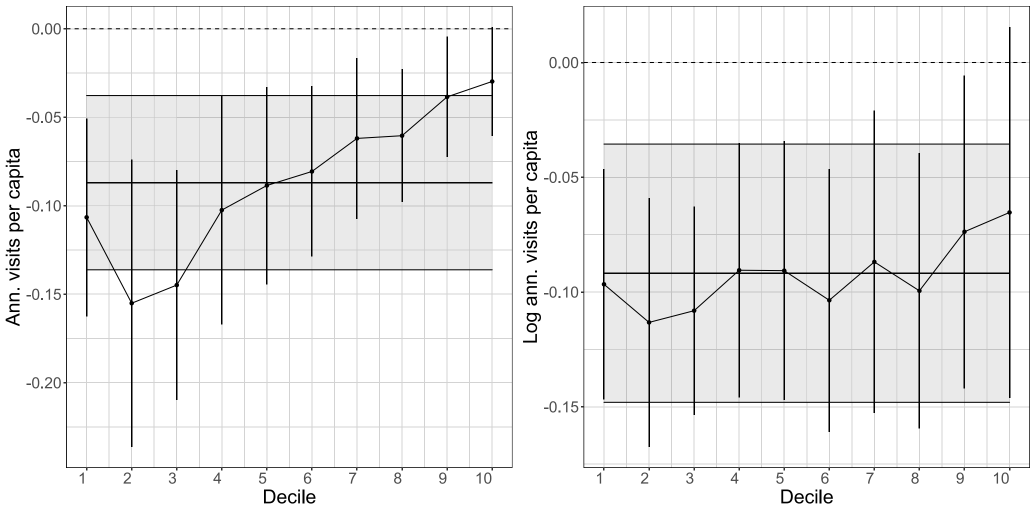

Fig. 4. Adoption: Estimates on nurse visits by income decile. *Notes:* This figure was not pre-registered and is post-blind. It shows our stacked difference-in-differences results from Model (1) for the impact of adopting the nurse visit copayment on the annualized number of curative nurse visits (and its logarithm) by deciles of the distribution of equivalized family disposable income. The dataset is stacked and balanced. Due to heterogeneity in municipality size, we weight by population size. Standard errors are clustered by municipality. The gray block, centered at the black horizontal line, shows the ATT estimate for the whole population and its confidence interval.

Robustness of Main Difference-in-Differences Results. We conducted an extensive set of analyses to test the robustness of the DD estimates of Table 1 on the impacts of the nurse visit copayment on public primary care use. Here, we summarize the findings (the details are in Online Appendix Section A.1). First, the stacked estimates on primary care use are robust to logarithmic outcomes and to unbalanced event-specific datasets that have more municipalities and observations than the balanced datasets in the main analysis. Second, the main findings are qualitatively robust to weighting municipalities uniformly instead of using population weights. Third (post-blind), the estimated reductions in nurse visits are somewhat larger when a longer follow-up of two years (instead of one year) is used with the stacking estimator.

Fourth, we use the CS estimator as an alternative to stacking. The CS estimator produces a static DD estimate by averaging over all group-time ATTs, weighted by group size, and thus having for most units a longer follow-up than the one year used with stacking. The CS estimates differ from the stacked estimates in two ways: (i) The effects on nurse visits are larger, partially explained by the longer follow-up. We estimate that the nurse copayment reduced annualized nurse visits per capita by 0.18–0.23 visits, or 13%–17% when using logarithmized outcomes, in the bottom 40% of the income distribution and by 0.08–0.10 visits, or 13%–16% when using logarithmized outcomes, in the top 40%. (ii) The CS estimates on the impacts on GP visits are all negative, but close to zero and statistically not different from zero.

Fifth (post-blind), we also assume quasi-random treatment timing among the later-treated municipalities, excluding the never-treated, and use the estimator proposed by Roth and Sant'Anna (2023a). The nurse visit copayment reduced nurse visits but had no effect on GP use, mostly in line with our main results. However, there is a caveat to our main heterogeneity findings: the differences in the effects in absolute terms are not that large or clear using the RS estimator. Moreover, the relative effects appear even larger for the top 40% of the income distribution.

6. National copayment abolition

Primary care utilization in 2020–2022 was affected by the COVID-19 pandemic. Curative nurse visits had not recovered to the pre-pandemic levels by May 2022 (Figure A10). This pattern is plausibly explained by the supply-side factors: nurses had been allocated to test, trace, and vaccinate. Consequently, primary care use was likely more supply-driven and gatekeeping stricter than during pre-pandemic times, attenuating the impact of the copayment. Second, the parallel trends assumption is less credible *ex ante* for the copayment abolition as (i) the treatment group has an overrepresentation of rural areas, and (ii) the epidemiological situation varied regionally throughout the pandemic, potentially correlated with the imbalanced baseline features. As shown below, the pre-treatment patterns in outcomes observed *ex post* do suggest that there is considerable uncertainty about the validity of the PTA.

Trend plots. Fig. 5 plots the trends in curative nurse visits for the bottom 40% and the top 40% of the distribution of equivalized family disposable income in the treatment and comparison municipalities, relative to the copayment abolition. There are large fluctuations in nurse use between the treated municipalities and their comparison group during the observation window. Nurse visits were increasing more in the treated areas for both income groups after the copayment abolition, compared to the periods immediately prior to the reform, although there was not much difference (if any) between the income groups in absolute terms. However, we do not view these patterns as causal and emphasize caution in interpreting them. The difference in nurse

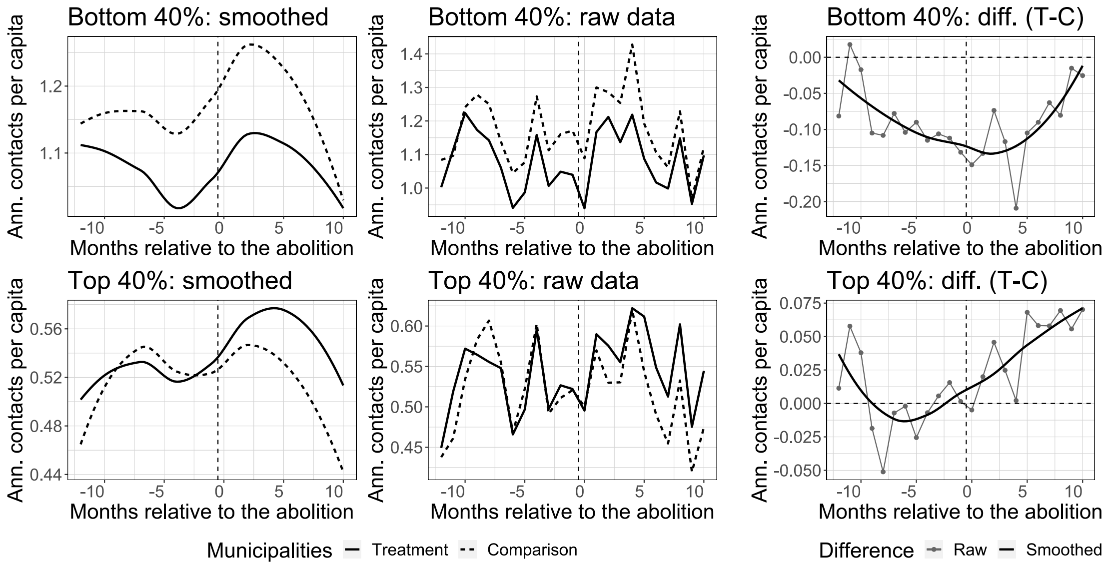

Fig. 5. Abolition: Evolution in nurse visits. *Notes:* The figure shows monthly evolution in the annualized number of curative nurse visits per capita separately at the bottom 40% and the top 40% of the distribution of equivalized family disposable income. The nurse visit copayment was abolished at time 0 in relative time. We divide the number of visits by individuals in income group G in municipality M in month T by the number of individuals in income group G living in municipality M, multiply the monthly rates by 12 to get annualized utilization, and average over municipalities by weighting by population. Sample sizes: 200 treated and 41 unexposed municipalities. The left column contains smoothed conditional means, fitted with local linear regression. The raw data are illustrated in the middle column, while the difference between treatment and comparison areas is depicted in the right column.

visits fluctuated considerably before the copayment abolition, which makes it plausible that there could be fluctuations of similar magnitude also during the post-treatment periods even in the absence of the policy change. We observe largely similar pre-treatment fluctuations in the difference of GP use between the policy groups as for nurse visits, but GP visits do not appear to increase after the nurse visit copayment abolition (Figure A11).

Fig. 6 (post-blind) illustrates that even small violations of parallel trends, which we consider very likely, lead to major uncertainty about the effects of the abolition. In essence, the abolition analyses are imprecise and provide evidence neither for the existence nor the absence of effects. The figure reports confidence intervals (CIs) at the 5% level as proposed by Rambachan and Roth (2023), bounding the maximum post-treatment violation of parallel trends between consecutive periods by $\bar{M}$ times the maximum pre-treatment violation of parallel trends. When naively assuming parallel trends, the estimates in the CIs are predominantly positive (they do include zero) for both the bottom 40% and the top 40% of the income distribution. The CIs noticeably widen when allowing for small post-treatment violations of parallel trends. For instance, the CI for the bottom 40% includes estimates from −10% to +20% when post-treatment violations of parallel trends are at maximum (only) 0.15 times the maximum pre-treatment violation of parallel trends. In contrast to the PAP, the rest of the pre-specified regression analysis is in Online Appendix Section A.2 due to considerable uncertainty about the PTA.

7. Discussion

We find that the introduction of a 10-euro copayment reduced curative nurse visits by 9%–10% in public primary care during a one-year follow-up among Finnish adults aged 25 or older. The finding that patients respond to copayments by adjusting their health care utilization is in accordance with studies conducted in several countries, such as the Nordic countries (Nilsson and Paul, 2018; Johansson et al., 2019; Landsem and Magnussen, 2018; Olsen and Melberg, 2018; Johansson et al., 2023; Kruse et al., 2022), Ireland (Nolan and Layte, 2017; Ma and Nolan, 2017), the United States (Chandra et al., 2010), Japan (Iizuka and Shigeoka, 2022), and Taiwan (Han et al., 2020). Most of the literature cited above focuses on children or adolescents, often using age-specific eligibility cutoffs for identification, while there is less or mixed evidence for moderate copayments reducing health care utilization in the general adult population. In fact, several studies report small or no effects on doctor visits in Germany, Sweden, or Ireland (Kunz and Winkelmann, 2017; Farbmacher and Winter, 2013; Johansson et al., 2023; Layte et al., 2009), with Chandra et al. (2010), based on US data, being an exception. Our study's age range encompasses both working-age individuals and pensioners, whereas previous studies often focus exclusively on either working-age adults or the elderly. These studies also primarily concentrate on doctor visits and not on primary care nurse visits, except for Johansson et al. (2023), who differentiate between non-physician and physician visits.

Our effect sizes for Finnish adults, defined as semi-arc elasticities following Brot-Goldberg et al. (2017), are reasonable when compared with previous findings concerning Swedish children and adolescents. The semi-arc elasticities represent changes in the

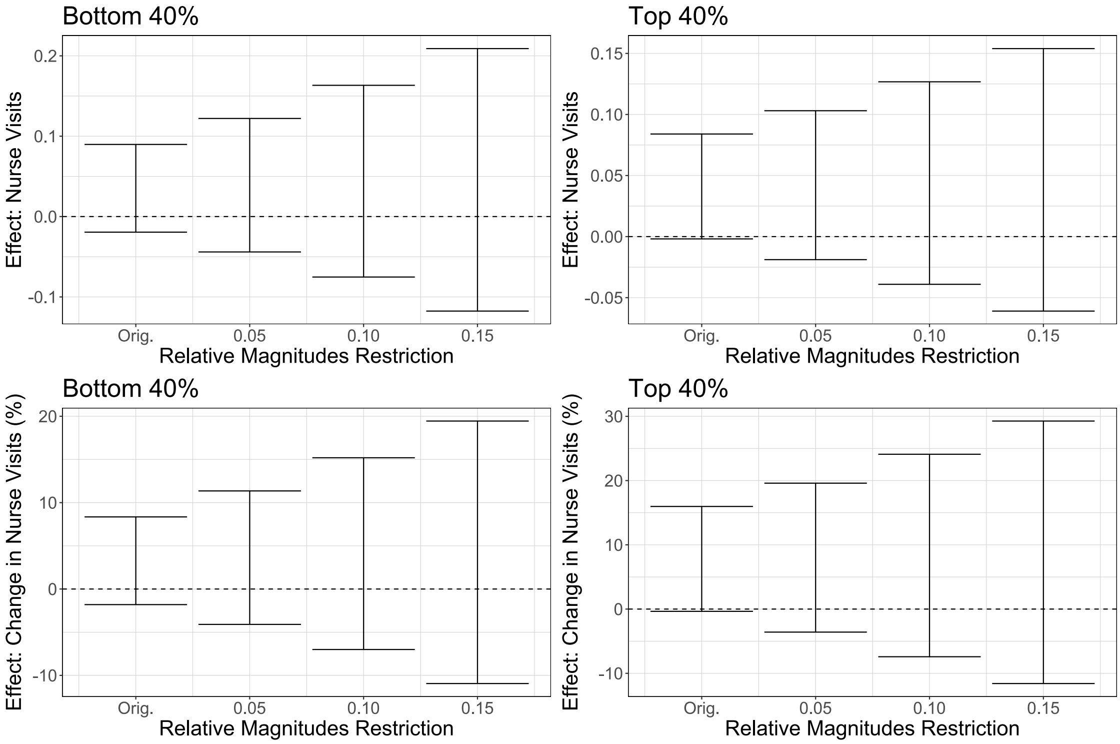

Fig. 6. Abolition and nurse visits: Bounding violations of parallel trends. *Notes:* This figure was not pre-registered and is post-blind. We apply the method proposed by Rambachan and Roth (2023) to construct confidence intervals for the estimated effect of the copayment abolition at the 5% level by bounding violations of parallel trends. The effects represent the estimated change in the number of annualized nurse visits during an 11-month follow-up. The change in percentage terms compares the estimate to the pre-treatment mean in the treated municipalities. Bottom 40% and top 40% are based on the distribution of equivalized family disposable income. First, we estimate a population-weighted event study specification that includes dynamic treatment indicators for the treated municipalities (D^l_{mt}), normalized at time $t = -1$, and municipality (m) and time (t) fixed effects: $y_{mt} = \alpha_m + \gamma_t + \sum_{l=-12,l\neq-1}^{10} \mu_l D^l_{mt} + \varepsilon_{mt}$. Standard errors are clustered by municipality. Next, we use the "relative magnitudes" restriction, $\Delta^{RM}(\bar{M})$, and construct conditional-least favorable hybrid confidence sets (C-LF) for the average of the estimated post-treatment effects using the R package *HonestDiD*. $\Delta^{RM}(\bar{M})$ bounds the maximum post-treatment violation of parallel trends between consecutive periods by $\bar{M}$ times the maximum pre-treatment violation of parallel trends.

number of nurse visits (q), normalized by the baseline, divided by the price change (p): $\frac{(q_1-q_0)/(q_1+q_0)}{(p_1-p_0)/2}$. Following Nilsson and Paul (2018), we define the price as the share of out-of-pocket costs of the total cost of the visit. The elasticity is sensitive to the selected parameters, listed in Online Appendix Section A.5. Therefore, we provide two estimates: the baseline (−0.41) and a large estimate (−1.24).[19] In Sweden and for GP visits, Nilsson and Paul (2018) report a semi-arc elasticity of −0.88 at the 20th birthday and −0.55 at the 7th birthday for a 10–15-euro copayment, and the estimates of Johansson et al. (2019) map to an elasticity of −1.11 at the 20th birthday for a 10-euro copayment.[20] All these Nordic estimates are smaller than the elasticities of −2.11 and −2.26 that Brot-Goldberg et al. (2017) calculate based on the RAND Health Insurance Experiment.

We find heterogeneity in treatment effects by income in absolute terms, with low-income individuals being more responsive. However, such heterogeneity is much weaker and statistically insignificant in relative terms (percentage changes) when the baseline differences in utilization are accounted for.[21] In contrast, previous studies on children and adolescents have found heterogeneity by income in both absolute and relative terms for GP copayments in Sweden (Johansson et al., 2019; Nilsson and Paul, 2018), but a study focusing on people in their mid-eighties found no evidence of heterogeneity (Johansson et al., 2023). In Denmark, the impacts on psychologist treatment were larger in relative terms for adolescents from low-income families (Kruse et al., 2022). These findings suggest that income-based heterogeneity in the impacts of cost sharing may depend on age, with the strongest evidence for heterogeneity being for children and adolescents. Based on our findings, we recommend specifying *ex ante* whether heterogeneity is defined in absolute or relative terms, given that baseline means often differ significantly between the groups of interest. In the context of our study, both dimensions are relevant, as argued in Section 1.

[19] The corresponding elasticity of the nurse visit copayment on GP visits is close to zero as we estimate only small reductions in GP use.

[20] We use the estimates for all individuals from their Table 1 and use a copayment of SEK 100 and the total cost of SEK 1500 per visit.

[21] There is one caveat to these heterogeneity results. In a robustness check using the Roth and Sant'Anna (2023a) estimator and later-treated areas as controls, added after registering our pre-analysis plan, the differences in effects in absolute terms are not clear-cut. Moreover, the relative effects appear even larger for the top 40%. However, the results from other robustness checks are in line with the finding of treatment effect heterogeneity by income in absolute terms.

The triage system in Finnish public primary care enabled us to examine whether the introduction of a copayment for nurse visit influenced not only the utilization of nurse visits but also that of GP visits. Our findings on GP visits are mixed: the baseline stacking analysis implies that the nurse visit copayment reduced the number of GP visits by 3%–5%, but robustness checks produce effect estimates that are not statistically different from zero. The lack of large effects on GP use should be a reassuring finding for those concerned about the potential health effects of copayments, given that in Finland GP visits proxy a nurse-assessed need for diagnosis and treatment. Previous research is limited, with the most relevant comparison coming from Sweden: Johansson et al. (2023) report that individuals delayed their primary care visits prior to the elimination of a copayment at age 85, affecting both physician and non-physician visits. Given that this delay was attributed to reductions in non-physicians visits, planned visits and follow-up visits, the authors conclude that the delay was primarily due to non-urgent care.

The limitations of our study include the following. First, we do not address the health and broader welfare implications of cost sharing. Gaining confirmatory insights into these effects in further studies is crucial, especially concerning low-income households. Other studies, such as those by Nilsson and Paul (2018) and Johansson et al. (2023), provide more comprehensive analyses of the impacts by type of visit, serving as a useful proxy for the severity of the underlying health condition. Second, we focus on the impacts on public primary care use only due to limitations in data availability, assuming that potential spillovers to other sectors, such as occupational or private healthcare, are small. Ideally, these potential spillovers should have been examined empirically. Third, our analysis does not account for the fact that the copayment was often charged only for the first three visits annually. This characteristic creates a promising opportunity to study how dynamic incentives (i.e., the spot price is reduced to zero after the third visit) affect primary care utilization. The dynamic incentives in cost sharing have received growing attention in the recent literature (see, e.g., Farbmacher and Winter, 2013; Aron-Dine et al., 2015; Cabral, 2017; Simonsen et al., 2021).

The strengths of our study lie in its research design and our PAP-based workflow. First, our staggered design is based on several events, reducing the risk of external shocks biasing the estimates. Moreover, the timing of the copayment adoption conditional on adopting is plausibly arbitrary. Second, our DD estimands differ from those used in age-based regression discontinuity designs (RDD), which are common in the literature. Age-based RDDs tend to produce estimates that are rather local to a specific birthday, and the potential discontinuities in service use may capture not only persistent impacts but also intertemporal substitution effects (Johansson et al., 2023). In contrast, we estimate effects for a broad adult population (aged 25 or older) over a one-year follow-up in a setting where potential anticipation effects are small. Third, we provide an example of using a pre-analysis plan and analysis blinding for retrospective observational data. Although the firewall between planning and analysis is not fraud-proof in our workflow, the PAP is a useful self-control tool and provides additional transparency to readers.

8. Conclusion

We analyze the effects of a staggered adoption of a nurse visit copayment (approximately 10 euros) on public primary care use of Finnish adults. Moreover, we provide an example of how a pre-analysis plan and analysis blinding can be used retrospectively in nonexperimental hypothesis testing. We find that the copayment reduced curative nurse visits by 9%–10% in public primary care during a one-year follow-up. There is statistically significant heterogeneity by income in absolute terms: the decrease in the number of visits is more than two times larger for individuals at the bottom 40% of the income distribution than at the top 40%. However, such heterogeneity is much weaker and statistically insignificant in relative terms (percentage changes). The estimates for GP visits are negative but closer to zero, ranging from −3% to −5%. However, the point estimates on GP visits are no longer statistically different from zero when using alternative estimators in robustness checks. Moreover, we analyze the subsequent nationwide abolition of the nurse visit copayment (also pre-registered) but were unable to draw causal conclusions from this reform. This inability stemmed from the lack of credibility in the parallel trend trends assumption required for such analyses.

Funding

This work is supported by the Finnish Institute for Health and Welfare, the Finnish Ministry of Social Affairs and Health, and Yrjö Jahnsson Foundation (research grant No. 20197209).

CRediT authorship contribution statement

Tapio Haaga: Conceptualization, Formal analysis, Writing – original draft, Writing – review & editing. **Petri Böckerman:** Conceptualization, Supervision, Writing – review & editing. **Mika Kortelainen:** Conceptualization, Supervision, Writing – review & editing. **Janne Tukiainen:** Conceptualization, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Pre-analysis plan and earlier working paper versions: https://osf.io/skuv9/. Replication codes: https://github.com/tapiohaa/ASMA3.

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jhealeco.2024.102866.

References

Aron-Dine, A., Einav, L., Finkelstein, A., Cullen, M., 2015. Moral hazard in health insurance: Do dynamic incentives matter? Rev. Econ. Stat. 97, 725–741. http://dx.doi.org/10.1162/REST_a_00518.

Baker, A.C., Larcker, D.F., Wang, C.C.Y., 2022. How much should we trust staggered difference-in-differences estimates? J. Financ. Econ. 144, 370–395. http://dx.doi.org/10.1016/j.jfineco.2022.01.004.

Brot-Goldberg, Z.C., Chandra, A., Handel, B.R., Kolstad, J.T., 2017. What does a deductible do? The impact of cost-sharing on health care prices, quantities, and spending dynamics. Q. J. Econ. 132, 1261–1318. http://dx.doi.org/10.1093/qje/qjx013.

Cabral, M., 2017. Claim timing and ex post adverse selection. Rev. Econom. Stud. 84, 1–44. http://dx.doi.org/10.1093/restud/rdw022.

Callaway, B., Sant'Anna, P.H.C., 2021. Difference-in-differences with multiple time periods. J. Econometrics 225, 200–230. http://dx.doi.org/10.1016/j.jeconom.2020.12.001.

Cengiz, D., Dube, A., Lindner, A., Zipperer, B., 2019. The effect of minimum wages on low-wage jobs. Q. J. Econ. 134, 1405–1454. http://dx.doi.org/10.1093/qje/qjz014.

Chandra, A., Flack, E., Obermeyer, Z., 2023. The Health Costs of Cost-Sharing. NBER Working Paper No. 28439, http://dx.doi.org/10.3386/w28439.

Chandra, A., Gruber, J., McKnight, R., 2010. Patient cost-sharing and hospitalization offsets in the elderly. Amer. Econ. Rev. 100 (1), 193–213. http://dx.doi.org/10.1257/aer.100.1.193.

de Chaisemartin, C., D'Haultfœuille, X., 2023. Two-way fixed effects and differences-in-differences estimators with several treatments. J. Econometrics 236, 105480. http://dx.doi.org/10.1016/j.jeconom.2023.105480.

Farbmacher, H., Ihle, P., Schubert, I., Winter, J., Wuppermann, A., 2017. Heterogeneous effects of a nonlinear price schedule for outpatient care. Health Econ. 26, 1234–1248. http://dx.doi.org/10.1002/hec.3395.

Farbmacher, H., Winter, J., 2013. Per-period co-payments and the demand for health care: evidence from survey and claims data. Health Econ. 22, 1111–1123. http://dx.doi.org/10.1002/hec.2955.

Gardner, J., 2021. Two-stage differences in differences. http://dx.doi.org/10.48550/arXiv.2207.05943.

Gormley, T.A., Matsa, D.A., 2011. Growing out of trouble? Corporate responses to liability risk. Rev. Financ. Stud. 24, 2781–2821. http://dx.doi.org/10.1093/rfs/hhr011.

Haaga, T., Böckerman, P., Kortelainen, M., Tukiainen, J., 2022. Cost sharing and primary care use: Evidence from staggered copayment adoption and later abolition. a pre-analysis plan.. http://dx.doi.org/10.17605/OSF.IO/FV2GA.

Haaga, T., Böckerman, P., Kortelainen, M., Tukiainen, J., 2023. Do adolescents from low-income families respond more to cost-sharing in primary care? Version 2. URL https://osf.io/vmuzf/.

Haaga, T., Böckerman, P., Kortelainen, M., Tukiainen, J., 2024. Does abolishing a copayment increase doctor visits? A comparative case study. B.E. J. Econ. Anal. Policy 24 (1), 187–204. http://dx.doi.org/10.1515/bejeap-2023-0056.

Han, H.-W., Lien, H.-M., Yang, T.-T., 2020. Patient cost-sharing and healthcare utilization in early childhood: Evidence from a regression discontinuity design. Am. Econ. J.: Econ. Policy 12, 238–278. http://dx.doi.org/10.1257/pol.20170009.

Iizuka, T., Shigeoka, H., 2022. Is zero a special price? Evidence from child health care. Am. Econ. J.: Appl. Econ. 14, 381–410. http://dx.doi.org/10.1257/app.20210184.

Johansson, N., de New, S.C., Kunz, J.S., Petrie, D., Svensson, M., 2023. Reductions in out-of-pocket prices and forward-looking moral hazard in health care demand. J. Health Econ. 87, 102710. http://dx.doi.org/10.1016/j.jhealeco.2022.102710.

Johansson, N., Jakobsson, N., Svensson, M., 2019. Effects of primary care cost-sharing among young adults: varying impact across income groups and gender. Eur. J. Health Econ. 20, 1271–1280. http://dx.doi.org/10.1007/s10198-019-01095-6.

Kruse, M., Olsen, K.R., Skovsgaard, C.V., 2022. Co-payment and adolescents' use of psychologist treatment: Spill over effects on mental health care and on suicide attempts. Health Econ. 31, 92–114. http://dx.doi.org/10.1002/hec.4582.

Kunz, J.S., Winkelmann, R., 2017. An econometric model of healthcare demand with nonlinear pricing. Health Econ. 26, 691–702. http://dx.doi.org/10.1002/hec.3343.

Landsem, M.M., Magnussen, J., 2018. The effect of copayments on the utilization of the GP service in Norway. Soc. Sci. Med. 205, 99–106. http://dx.doi.org/10.1016/j.socscimed.2018.03.034.

Lavaste, K., 2023. Private health insurance in the universal public healthcare system: The role of healthcare provision in Finland. Health Policy 132, 104820. http://dx.doi.org/10.1016/j.healthpol.2023.104820.

Layte, R., Nolan, A., McGee, H., O'Hanlon, A., 2009. Do consultation charges deter general practitioner use among older people? A natural experiment. Soc. Sci. Med. 68 (8), 1432–1438. http://dx.doi.org/10.1016/j.socscimed.2009.02.014.

Ma, Y., Nolan, A., 2017. Public healthcare entitlements and healthcare utilisation among the older population in Ireland. Health Econ. 26, 1412–1428. http://dx.doi.org/10.1002/hec.3429.

Maier, C.B., Aiken, L.H., 2016. Task shifting from physicians to nurses in primary care in 39 countries: a cross-country comparative study. Eur. J. Public Health 26, 927–934. http://dx.doi.org/10.1093/eurpub/ckw098.

Nilsson, A., Paul, A., 2018. Patient cost-sharing, socioeconomic status, and children's health care utilization. J. Health Econ. 59, 109–124. http://dx.doi.org/10.1016/j.jhealeco.2018.03.006.

Nolan, A., Layte, R., 2017. The impact of transitions in insurance coverage on GP visiting among children in Ireland. Soc. Sci. Med. 180, 94–100. http://dx.doi.org/10.1016/j.socscimed.2017.03.026.

Nosek, B.A., Ebersole, C.R., DeHaven, A.C., Mellor, D.T., 2018. The preregistration revolution. Proc. Natl. Acad. Sci. 115, 2600–2606. http://dx.doi.org/10.1073/pnas.1708274114.

OECD, 2021. OECD health statistics 2021, accessed 26/06/2023. URL https://stat.link/m5nfxa.

Olden, A., Møen, J., 2022. The triple difference estimator. Econom. J. 25, 531–553. http://dx.doi.org/10.1093/ectj/utac010.

Olken, B.A., 2015. Promises and perils of pre-analysis plans. J. Econ. Perspect. 29, 61–80. http://dx.doi.org/10.1257/jep.29.3.61.

Olsen, C.B., Melberg, H.O., 2018. Did adolescents in Norway respond to the elimination of copayments for general practitioner services? Health Econ. 27, 1120–1130. http://dx.doi.org/10.1002/hec.3660.

Rambachan, A., Roth, J., 2023. A more credible approach to parallel trends. Rev. Econom. Stud. http://dx.doi.org/10.1093/restud/rdad018.

Rinne, H., Blomgren, J., 2023. Use of outpatient healthcare services before and after the onset of unemployment: A register-based propensity score matched study from Finland. PLoS One 18 (8), 1–13. http://dx.doi.org/10.1371/journal.pone.0288423.

Roth, J., Sant'Anna, P.H., 2023a. Efficient estimation for staggered rollout designs. J. Political Econ. Microecon. http://dx.doi.org/10.1086/726581.

Roth, J., Sant'Anna, P.H., 2023b. When is parallel trends sensitive to functional form? Econometrica 91, 737–747. http://dx.doi.org/10.3982/ECTA19402.

Simonsen, M., Skipper, L., Skipper, N., Christensen, A.I., 2021. Spot price biases in non-linear health insurance contracts. J. Public Econ. 203, 104508. http://dx.doi.org/10.1016/j.jpubeco.2021.104508.

Winkelmann, R., 2004. Co-payments for prescription drugs and the demand for doctor visits – evidence from a natural experiment. Health Econ. 13, 1081–1089. http://dx.doi.org/10.1002/hec.868.