CSS INSTITUT
WORKING PAPER SERIES

CSS INSTITUT WP 2026/01

ACCESS TO WHOM? RELATIONAL FRICTIONS AND HEALTHCARE DEMAND

LINN **HJALMARSSON**
NICOLAS **SCHREINER**

FEBRUARY 2026

CSSINSTITUT
WORKING PAPER SERIES

CSS INSTITUT WP 2026/01

ACCESS TO WHOM? RELATIONAL FRICTIONS AND HEALTHCARE DEMAND

LINN **HJALMARSSON** CSS INSTITUTE
NICOLAS **SCHREINER** CSS INSTITUTE & UNIVERSITY OF BASEL

FEBRUARY 2026

CSS INSTITUTE FOR EMPIRICAL HEALTH ECONOMICS
Tribschenstrasse 21, CH-6002 Lucerne
info@css-institut.ch
www.css-institut.ch

ABSTRACT

ACCESS TO WHOM? RELATIONAL FRICTIONS AND HEALTHCARE DEMAND*

How responsive is healthcare demand to the availability of a patient's specific physician, as opposed to the aggregate supply of providers? We exploit a novel natural experiment—the temporary holiday absences of a patient's regular general practitioner (GP). Using administrative claims data for 1.8 million individuals from Switzerland's largest health insurer over 2014–2024, we implement an event-study design that isolates plausibly exogenous shocks to provider availability. During GP absences, the weekly probability of any primary care visit falls by two-thirds. Substitution to other physicians is virtually nonexistent, while emergency department visits increase by a negligible amount. Total weekly healthcare expenditures fall by 19%. The results are statistically indistinguishable across markets with few versus many alternative GPs. Utilization reverts immediately to baseline upon the GP's return, with no catch-up demand and no increase in hospitalizations over the subsequent two months—suggesting that the foregone care was not acutely necessary over this horizon. Our findings establish that the patient–physician relationship is a first-order determinant of healthcare demand.

JEL CLASSIFICATION: I11, I12, I13, D12, D83

KEYWORDS: Primary Care, Healthcare Demand, Supplier-Induced Demand, Physicians, Relational Frictions, Discretionary Care, Low-Value Care

CORRESPONDING AUTHOR:
Nicolas Schreiner
CSS Institute for Empirical Health Economics
Tribschenstrasse 21
CH-4002 Lucerne

E-mail: nicolas.schreiner@css-institut.ch

*We thank Caroline Chuard-Keller, Alexa Dahinden, Philip Hochuli, Michael Gerfin, Fabienne Lötschter, Christian P.R. Schmid, Bianca Strasser, Roman Stutzer, Katrin Zocher, and conference participants participants of the EuHEA Conference 2024 in Vienna and the Croatian Health Economics Workshop 2025 in Split for their suggestions and comments.

1 Introduction

Policies aimed at containing healthcare costs or improving population health frequently hinge on improving and managing access to primary care (Baicker and Chandra, 2004). As the "front door" to the health system, General Practitioners (GPs) act as gatekeepers, care coordinators, and long-term managers of chronic disease. A central question for health economics and policy design is how responsive patient demand is to the availability of these providers. If access to the usual point of entry to the broader healthcare system becomes more restrictive, does demand shift to other local providers, is it deferred, or does it vanish? The answer has immediate implications for the efficiency of healthcare markets, the prevalence of low-value care, and our understanding of the non-monetary frictions that shape utilization.

Isolating the causal effect of primary care access on demand is empirically difficult. The supply and demand for healthcare are deeply endogenous—physicians choose where and how much to practice based on local disease prevalence and patient characteristics, while patient health status and health-seeking behaviors are themselves a function of the available supply of care (Chandra et al., 2016). Simple cross-sectional associations between provider density and healthcare utilization are therefore difficult to interpret causally. Alternative research designs that study the entry or exit of providers typically capture a complex set of both short-run and long-run equilibrium adjustments by market participants, making it difficult to isolate the immediate behavioral response of patients to access barriers (see, e.g., Finkelstein et al., 2021).

This paper overcomes these empirical challenges by introducing a novel source of identification: the temporary holiday absences of a patient's regular GP. Using administrative data from CSS Insurance, Switzerland's largest mandatory health insurer, covering 1.8 million individuals over 2014–2024, we exploit the fact that GP holiday schedules are determined by the physician's personal preferences, independent of short-term fluctuations in their patients' health needs. These absences—lasting one to three weeks—create a sharp, temporary reduction in access to the patient's usual entry point while leaving insurance coverage, cost-sharing, and the aggregate supply of alternative providers unchanged. We implement a two-step event-study design comparing weekly healthcare utilization before, during, and after a patient's GP is on holiday to a contemporaneous control group of similar patients whose GP is never absent in the respective period, thereby absorbing seasonal trends and local shocks.

Our main finding is stark: a GP's holiday absence reduces the weekly probability of any primary care visit by 5.4 percentage points, a 68% decline relative to the baseline of 7.9%. Substitution to other providers is virtually nonexistent. Despite the continuing availability of other GPs, specialists, and emergency departments in the same local market, specialist visits and prescription-drug spending *fall* significantly (by 10% and 24%, respectively) during the absence. Emergency department visits increase, but the effect is economically negligible—offsetting roughly 1% of foregone primary care volume. Altogether, total weekly healthcare expenditures decline by CHF 17.7, or about 19% per week of regular GP absence. Crucially, utilization and spending revert immediately to baseline upon the GP's return, with no evidence of a "catch-up" surge or a rise of inpatient hospitalizations in the subsequent two months.

We make three contributions. First, we quantify a *relational friction* in healthcare demand that is larger or of similar magnitude to documented price or other non-monetary frictions.[1] The existing literature has extensively documented the role of prices (Brot-Goldberg et al., 2017), e.g., with the RAND Health Insurance Experiment establishing a price elasticity of approximately -0.2 (Manning et al., 1987; Aron-Dine et al., 2013). By contrast, the temporary removal of a single provider—holding prices, coverage, and local supply fixed—reduces primary care utilization by roughly 70% and overall health expenditures by 19%. The friction is not explained by search or travel costs: exploiting geographic variation in provider density, we show that the demand collapse is statistically indistinguishable across localities with few (1–3) versus many (10 or more) available alternative GPs. This points to the specific patient–physician relationship—accumulated information and trust that resolve the fundamental asymmetry first articulated by Arrow (1963)—as the binding constraint, rather than supply capacity. When that relationship is disrupted, most patients prefer to forego care entirely rather than consult an unfamiliar provider.

Second, our results speak to the prevalence of effectively discretionary healthcare. A common policy goal in attempts to rein in rising healthcare expenditures is the reduction of "low-value" services—care where the cost exceeds the clinical benefit (Schwartz et al., 2014; Cutler et al., 2019). Unlike studies that rely on clinical guidelines to classify such treatments, we use a revealed-preference approach: care that is foregone without substitution, without catch-up, and without observable adverse health events in the subsequent two months provides a behavioral upper bound on the share of routine demand that is not driven by acute medical necessity. We stress that the absence of short-run health consequences does not rule out longer-run costs; the estimate should be interpreted as suggestive of the margin at which demand is most elastic to relational access.

Third, our findings illuminate the GP's role as a *de facto* gatekeeper and demand catalyst. The health economics literature has extensively debated the effectiveness of formal gatekeeping models (e.g., Health Maintenance Organizations), which contractually require a primary care referral for specialist access, with mixed evidence

[1]The role frictions and other behavioral aspects have been widely studied for health insurance choices see, e.g, Handel, 2013; Handel and Kolstad, 2015.

on their ability to contain costs. Separately, the theory of supplier-induced demand, rooted in the economics of information and credence goods, posits that physicians can use their informational advantage to shift patients' demand curves outward, potentially leading to overuse (Mcguire, 2000; Iizuka, 2012; Clemens and Gottlieb, 2014). The decline in downstream utilization—specialist visits, prescriptions, non-physician outpatient services—occurs even among patients in insurance plans that permit direct specialist access, and is in fact *largest* for this group. This suggests that informal, patient-driven "behavioral gatekeeping" through a trusted GP can be as binding as formal insurance-based referral requirements, and that a substantial share of downstream demand is catalyzed by the specific GP rather than by the generic availability of medical services. Our results thus provide a micro-foundation for supplier-induced demand as a relational phenomenon: when the specific inducing agent is removed, much of the induced demand dissipates rather than transferring to other suppliers.

The remainder of this paper proceeds as follows. Section 2 describes the Swiss institutional setting. Section 3 presents the empirical strategy. Section 4 describes the data and sample construction. Section 5 reports results, heterogeneity analyses, and a discussion of implications. Section 6 concludes.

2 Institutional Background

Switzerland's healthcare system is centered around mandatory health insurance (MHI), governed by the Federal Health Insurance Act (*KVG*) and operates under a managed competition framework. All residents must purchase coverage from one roughly 40 private, non-profit insurers. Switching insurers or plans is permitted once per year during the open enrollment period, which runs in the fall until the end of November. Insurers in MHI may not reject applicants for any reason, and all MHI plans must provide the same comprehensive benefits package, covering a standardized set of ambulatory, inpatient, and pharmaceutical services. The content of this benefits catalog is set at the federal level; insurers cannot add or exclude any treatments. Premiums are community-rated within each canton and broad age category (child, young adult, adult). To prevent cream-skimming, Switzerland employs a national risk-adjustment scheme that incorporates morbidity measures through 34 pharmaceutical cost groups (PCGs), as well as demographic and regional factors. Given this high degree of regulation, competition among insurers is largely confined to premiums and service quality.

Cost-sharing rules are standardized across all MHI plans. All adult enrollees choose an annual deductible (*Franchise*) between CHF 300 and CHF 2,500; for children, deductibles are typically zero. After the deductible is met, patients pay a 10% coinsurance rate until reaching an annual stop-loss of CHF 700 for adults and CHF 350 for children, after which coinsurance falls to zero. This uniform piecewise structure ensures that marginal out-of-pocket prices are identical across the population, conditional on cumulative annual spending. On the supply side, outpatient physicians are reimbursed on a fee-for-service basis according to the national *TARMED* tariff, while inpatient care is financed via diagnosis-related group (DRG) case-based payments. Prices for prescription drugs included in the MHI benefits are also set administratively by the federal government.

Every insurer must offer a ”standard” plan with unrestricted physician choice, which permits direct access to most outpatient specialists. In addition, insurers may offer regulated alternative insurance plans (AIPs) with lower premiums in exchange for restricting the patient's first point of contact. These include: (i) a general practitioner (GP) model, in which patients register with a specific GP who must be consulted before specialist care (though enforcement is imperfect); (ii) a Health Maintenance Organization (HMO) model, requiring patients to first visit physicians within a defined network; and (iii) a telemedicine model, which requires initial triage via a medical hotline for triage before in-person consultation, except in emergencies. All AIPs constrain the first point of contact but must still deliver the full, nationally uniform benefits package, ensuring that any observed differences in utilization across plan types reflect access restrictions, not differences in covered services. Even in the standard plan, a physician referral is typically required for most non-physician medical services to be reimbursed. Thus, plan choice effectively reflects a trade-off between lower premiums and reduced provider access, rather than differences in covered services.

Primary care in Switzerland is delivered predominantly by self-employed physicians working in solo or small-group practices. Little institutionalized coordination exists between different GPs, although some cantons designate a rotating regional emergency GP during non-working hours. Certain GPs opt to publish absence announcements with GP substitutes in the local municipality bulletin, however we cannot observe to what extent patients are aware of these announcements. Although patient registration is not mandated outside GP- or HMO-gatekeeping plans, over 90% of adults report having a regular GP (Dorn, 2023). The GP typically serves as the *de facto* coordinator of care, the main prescriber, and the principal source of both formal and informal referrals—even for patients with nominal free provider choice.

This institutional arrangement is central to our empirical strategy. A temporary absence of a patient's regular GP constitutes a salient, plausibly exogenous organizational shock to healthcare access. Crucially, such absences affect neither insurance coverage nor patient cost-sharing, allowing us to isolate the role of non-monetary access frictions in shaping healthcare utilization and expenditures.

3 Empirical Strategy

3.1 Identification

Understanding how access to primary care shapes healthcare demand is central to informing health policy, especially when weighing broad availability against the risk of overuse. Estimating this relationship is empirically challenging for two reasons. First, the location and intensity of healthcare provision are endogenous: providers choose where and how much to work based on local demand and patient characteristics, while area-level factors such as income, demographics, and proximity to medical and educational institutions influence both supply and treatment styles. Second, research designs based on provider entry or exit typically capture large, persistent changes that trigger both short- and long-run adjustments, making it difficult to isolate patient responses. Estimates from cross-sectional variation or persistent within-area changes in provider supply risk conflating access effects with unobserved differences in patient populations and health system dynamics, leading to biased conclusions about how primary care access affects demand.

Our design overcomes these challenges by exploiting the temporary holiday absences of a patient's regular general practitioner (GP) as a plausibly exogenous source of variation in primary care access. Such absences—lasting one to three weeks in our setting—remove the patient's usual entry point into the healthcare system without materially changing the overall supply of care in the region. Our identifying assumption is that the timing of GP holiday absences is orthogonal to short-run fluctuations in the health needs of the GP's patient pool. Because these absences affect neither insurance coverage nor patient cost-sharing, they allow us to isolate the role of non-monetary access frictions in shaping utilization and expenditures. By comparing healthcare utilization before, during, and after the absence to utilization among a contemporaneous control group whose regular GP is not absent, we can estimate the causal effect of primary care access frictions on healthcare demand and assess the extent and consequences of foregone care.

3.2 Design

The unit of observation is the individual-calendar-week. For each patient-year, we assign a regular GP and identify the weeks in which that GP is completely absent (see Sections 4.2 and 4.2 for details). We focus on a fixed seasonal window—weeks 20 to 39 in our baseline[2]—capturing the peak summer holiday period when GP absences are most common. This restriction ensures: (i) clear classifications of patients as ei-

[2] The design can be applied to any seasonal window; results are robust to this choice.

ther experiencing a single continuous absence spell or having uninterrupted access; (ii) comparability across years; and (iii) avoidance of spurious effects from unrelated within-year shocks. Individuals appearing in multiple years are treated as separate observational units to avoid cross-year contamination from different treatment exposure statuses.

Treated individuals are those whose regular GP has exactly one absence spell in the observation window; never-treated individuals have no such absences and serve as a control group to flexibly absorb seasonal variation in healthcare use. This is important because patient absences (e.g., travel abroad where we do not observe healthcare use) are also more frequent during peak holiday periods. Using contemporaneous never-treated controls helps account for such patterns and yields plausible counterfactuals for treated patients.

Our interest extends beyond only in the average treatment effect during the GP absence to the dynamic responses before and after patients experience access frictions. In our setting treatment timing is not only staggered but treatment also varies in its duration[3] and also ”switches off” again for all treated units. To accommodate this, we adopt a ”collapsed” event study design. To maintain comparability across treatment lengths, we code all weeks of GP absence as event time $k = 0$; $k = -1$ represents the last pre-absence week, and $k = 1$ refers to the first week post-return. Weeks with $k \neq 0$ are always one week long, ensuring well-defined pre/post-treatment dynamics across absence lengths.[4]

The choice of the omitted reference period requires balancing two considerations. On the one hand, patients who anticipate their GP's upcoming absence—particularly those requiring regular check-ups or chronic-disease management—may shift demand forward into the weeks immediately preceding the holiday. Retaining multiple pre-treatment weeks allows us to test for such anticipatory behavior, but it also means that weeks close to treatment onset may themselves be affected by the announcement of the absence. On the other hand, the reference period must be close enough to the treatment window to reflect the underlying trends in utilization.[5] We use $k = -4$ as the omitted baseline in our main specification: it is sufficiently removed from the onset of the absence to be plausibly unaffected by anticipation effects, yet close enough to capture the relevant seasonal trajectory. The event-study coefficients for $k \in \{-6, -5, -3, -2, -1\}$ then serve a dual purpose—testing for par-

[3]As we explain in Section 4, to ensure that treatment spells represent meaningful, temporary shocks rather than extended or potentially structural absences, we restrict the analysis to GP absences lasting no more than three full weeks.

[4]We later assess whether estimated effects differ systematically by the length of the absence.

[5]An additional complication arises from us defining an absent week only if the GP has zero patient visits during the week. It is therefore possible that a patient's GP is absent on all days of the week except for one and thus the actual access disruption already started earlier or will have ended only later than indicated by our week-based absence measure.

allel pre-trends and detecting any forward-shifting of demand. We show in Section 5 that pre-treatment coefficients are flat and close to zero, supporting this choice. Figure 1 illustrates the event time structure, showing examples of possible treatment paths and the alignment with control group observations.

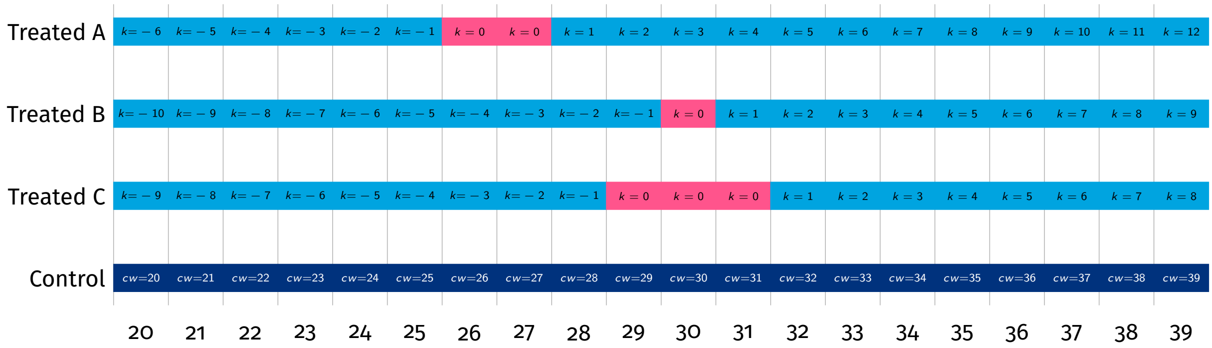

Figure 1: Setup of Treatment

3.3 Estimation

Our setting resembles a staggered difference-in-differences design. However, standard two-way fixed effects estimators are ill-suited given the associated concerns (see, e.g., Roth et al. (2023) for a overview) and existing staggered-treatment solutions, such as Callaway and Sant'Anna (2021) or Sun and Abraham (2021), do not directly accommodate our data structure. In particular, calendar weeks repeat across years, and pooling them as if they occurred in the same year would ignore year-specific shocks—such as varying school breaks, holiday timing, and epidemic waves—that are essential to control for. The large number of week-by-year fixed effects, relative period indicators, and outcomes of interest also poses computational challenges with these estimators.

We therefore implement a two-step residualization procedure, similar in concept to the approaches proposed by Borusyak et al. (2024) (imputation based estimation), Gardner et al. (2025) (two-stage difference-in-differences), or Liu et al. (2024). In step one, we use never-treated individuals to model untreated outcomes; in step two, we run an event-study regression on residualized outcomes for treated individuals.

Step 1: Predicting Untreated Outcomes

Let y_{it} denote an healthcare outcome for individual i utilized primary care services in year-calendar-week t. Using only never-treated individuals ($i \in \mathcal{N}$), we estimate:

$$y_{it} = \gamma_t + \boldsymbol{X}'_{it}\boldsymbol{\beta} + \varepsilon_{it} \tag{1}$$

where γ_t are calendar week-by-year fixed effects and $\boldsymbol{X}_{it}$ includes demographics, insurance characteristics, region, and morbidity indicators. Individual fixed effects are omitted in this stage to allow out-of-sample prediction for treated individuals.

Using the estimated parameters, predicted untreated potential outcome for treated individuals ($i \in \mathcal{T}$) in year-calendar-week t are $\hat{y}_{it}^{(0)} = \hat{\gamma}_t + \boldsymbol{X}'_{it}\hat{\boldsymbol{\beta}}$ and subsequent residualized outcomes $\tilde{y}_{it} = y_{it} - \hat{y}_{it}^{(0)}$.

Step 2: Event Study on Residualized Outcomes

For treated individuals we estimate:

$$\tilde{y}_{it} = \alpha_i + \sum_{k \in \mathcal{K} \setminus \{-4\}} \delta_k \cdot 1\{t - T_i = k | t - T_i < 0, t - R_i = k | t - R_i > 0\} + u_{it} \tag{2}$$

where α_i are individual fixed effects, T_i denotes the first week in which individual i's regular GP is absent, R_i denotes the last week the in which individual i's regular GP is absent, k indexes event time, and $k = -4$ is the omitted baseline. We include k from -6 to +8 to capture pre-trends, immediate effects, and potential post-absence ”catch-up” behavior. Standard errors are clustered at the individual-year level throughout.

4 Data

4.1 Data Source

Our analysis draws high-frequency administrative data from CSS Insurance (CSS), Switzerland's largest mandatory health insurer. In 2024 CSS insured roughly 1.5 million individuals in MHI—about one-sixth of the Swiss population—across all cantons, with a demographic composition broadly representative of the national population. The data span 2014–2024, excluding 2020 and 2021 to avoid distortions from the COVID-19 pandemic, which markedly altered provider availability and patient behavior. The dataset combines complete claims histories with individual demographics and insurance contract characteristics. Each claim records the exact service date, a unique provider identifier, the prescribing physician (where applicable), and both gross and patient co-payment amounts. Demographic fields include age, sex, nationality, place of residence; contract fields include deductible level, accident insurance coverage, and insurance plan. Crucially, the claims data cover virtually every healthcare service reimbursed under Switzerland's basic health insurance, enabling day-level measurement of healthcare use and spending. CSS's large market share and broad coverage make these data well-suited for studying population-wide patterns in care utilization.

4.2 Sample Construction

We begin by restricting to enrollees with at least one physician visit in the entire sample period, since identifying a "regular GP" requires at least one GP contact. Within each calendar year, we retain only individuals continuously enrolled for the full year, excluding newborns and those who die mid-year. This ensures complete observation windows for GP assignment and absence detection.

We exclude all claims billed to accident insurance and all maternity-related claims. Because the claims data do not cover employer-provided accident insurance for employed adults, we exclude the partial accident claims we do observe (for children, retirees, and non-employed adults) to ensure a comparable sample. Maternity episodes, while fully covered, involve predictable, high-intensity utilization that is orthogonal to our variation in GP availability but could dominate outcome measures in affected weeks while also exempt from co-payments.

Assigning the Regular GP

For each patient-year, we identify a "regular GP" using a data-driven algorithm reflecting Swiss primary care practice. Eligible providers are general internal medicine physicians, pediatricians, "practicing physicians" in GP residency, and any GP or group practice explicitly designated as a gatekeeper in the patient's insurance contract.

We count the number of distinct days in the year on which the patient had a visit, defined by TARMED code 00.0010 ("basic consultation, first 5 minutes"), which can be billed only once per patient interaction and thus reliably captures an in-person consultation. The provider with the highest count is assigned as the regular GP for that year, provided there is no tie. In the event of a tie, we select the provider listed in the insurance contract; if neither is listed, the patient-year is excluded. This procedure yields a regular GP for roughly 75% of patient-years can be assigned a regular GP under this procedure; the remainder are dropped to ensure an unambiguous, active primary care relationship.

Measuring GP Absences

We detect GP holiday absences directly from visit patterns. For each GP, we build a week-by-week panel and count unique patients—based on a full sample without any prior exclusions—seen under code 00.0010 (allowing at most one such visit per patient per day). A week with zero such visits is coded as "absent". Consecutive absent weeks are grouped into a single holiday spell, with the spell length equal to the number of consecutive absent weeks. We also record the total number of absent weeks per GP-year.

This approach avoids reliance on self-reports or scheduling data and is consistently defined across all providers and years. Focusing on basic, physical consultations ensures we do not misclassify remote services, follow-up calls, or administrative billings as of presence.

Patient–Week Panel

To align outcomes with the timing of GP absences, we transform the claims data into a patient–week panel. For each patient–week, we construct binary indicators and costs for multiple healthcare categories such as primary care, emergency department, specialist, inpatient, and pharmaceutical drugs encounters. The unfiltered panel covers 1.84 million unique individuals and 7.45 million patient-years (about 388 million patient–weeks). To ensure that absences reflect temporary holidays rather than structural low activity, we exclude patient-years where the GP (i) is absent eight or more weeks in total (suggestive of part-time practice or semi-retirement); (ii) records fewer than 100 visits over the entire year (≈600 visits per year in the total population, implausible for a full-time GP). After these restrictions, about 335 million patient–weeks remain in the analysis dataset.

Main Summer Analysis Window

We focus on the summer holiday period, when GP absences are both most frequent and most concentrated. Throughout most of the year, 5–10% of GPs are absent in a given week; between weeks 26 and 33 (late June to early August), this rises to 20–25% (Figure 2).[6] We define our main analysis window as weeks 20–39 (early/mid-May to mid/late-September), balancing a high incidence of absences with enough surrounding weeks to assess pre-trends and post-absence dynamics.[7] To ensure clean pre-treatment measurement, we exclude patient-years where the GP is absent in week 19 (which could make week 20 the GP's first week back). The summer window now initially contains approximately 129 million individual-week observations.

We classify individuals as ”treated” if their GP has at least one absence spell in weeks 20–39 and as ”never treated” otherwise. Roughly one-third fall into the never treated group. Among treated patients, we keep only those whose GP has exactly one absence spell in the window, eliminating 40% of treated cases with multiple spells to avoid overlapping anticipation and recovery periods in favor of a cleaner, single-episode design.

[6] This period coincides mostly with school summer break in Switzerland, which differs both in length as well as starting and ending points between cantons, however.

[7] Later robustness analyses show that other seasonal windows produce almost identical results as the summer period.

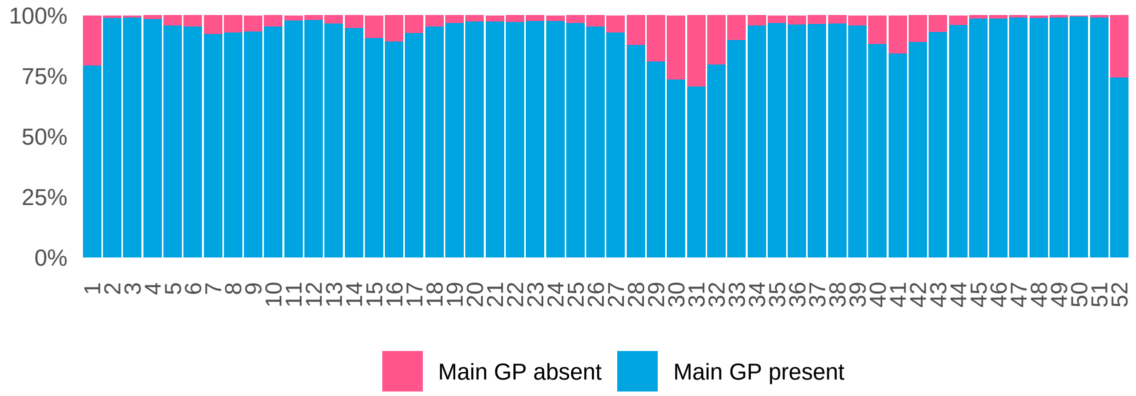

Figure 2: Share of Individuals with their GP Absent by Week

To sharpen identification, maximize comparability and dynamic coverage, we further restrict treated patients to those whose absence starts in weeks 26–33, yielding at least six pre-treatment and four post-treatment weeks within the window. we focus on the core high-absence period. We retain all never-treated clients and, among the treated, only those whose first absent week lies between weeks 26 and 33. This restriction leaves us with about 91 million observations. Finally, we remove absences longer than three weeks (about 5% of treated cases), as these are more likely to reflect sabbaticals or illness rather than typical holidays.

The resulting main summer sample comprises roughly 89 million patient-weeks: 50% never treated, 10% with a one-week absence, 25% with two weeks, and 15% with three weeks.

4.3 Descriptive Statistics

This section documents pre-treatment sample composition in the main summer window and presents unadjusted primary care utilization patterns around regular GP absences. We first summarize baseline characteristics to assess the comparability of treated (indexed by the first week of GP absence) and control individuals, then plot raw utilization data to offer a transparent preview of the main findings.

4.3.1 Baseline Summary Statistics

Table 1 reports summary (means and standard deviations) pre-treatment characteristics (weeks 20–25) by treatment timing. The groups are well-balanced on demographic indicators. Across all groups, the average age is in the mid-40s, approximately 54% of individuals are female, and roughly 80% hold Swiss citizenship. Baseline health status, measured by comorbidities—the number of pharmaceutical cost groups (PCGs) used in national risk adjustment—is comparable, albeit slightly better among the never-treated. The overall average of roughly 0.6 PCGs indicates most individuals have no chronic conditions, while a meaningful minority has multiple.

Modest differences emerge in insurance plan choices. Accident coverage inclusion in the MHI contract is more common in the treated groups.[8] Individuals in the never-treated control group select higher annual deductibles on average compared to their treated counterparts (CHF 820 versus CHF 617–677). This aligns with their greater tendency to enroll in HMO or telemedicine models, whereas treated individuals are more frequently in standard or GP-gatekeeping plans. While this pattern may reflect underlying differences in practice types, our two-step estimation procedure, which conditions on these plan characteristics, is designed to account for such compositional differences.

Baseline weekly healthcare utilization is relatively low and remarkably stable across groups, supporting the validity of our research design. In any given pre-treatment week, only 7–8% of patients visit a primary care physician, 2.6–3.6% see a specialist, and about 0.3% attend an emergency department. Contact with any reimbursed healthcare provider (also including pharmacies, laboratories, and therapists) occurs in 25–31% of weeks. Expenditure patterns mirror utilization: mean total weekly spending lies between CHF 86 and CHF 95, with pharmaceuticals accounting for CHF 20–23, primary care services for CHF 9–10, and total physician services for CHF 31–34.

4.3.2 Raw Primary Care Utilization Around GP Absences

To provide a non-parametric overview of the raw data, Figure 3 plots the weekly probability of a primary care visit for each treatment group and for the never-treated control group. The visualization offers compelling initial evidence that strongly motivates our main empirical strategy. Three patterns are immediately apparent.

First, in the pre-absence period (weeks 20–25), the visit probabilities of all groups track one another closely at about 7–8% per week. This visual evidence supports the parallel trends assumption crucial for our difference-in-differences design. Second, the onset of a GP's absence precipitates a sharp and substantial decline in primary care visit probability for the affected group. Weekly rates collapse to between 1% and 2%, a reduction of almost 70% relative to the baseline. Importantly, there is no discernible evidence of anticipatory care-seeking in the weeks immediately preceding the absence. Third, the recovery is as swift as the decline. Utilization rates for the treated groups revert to the level of the control group almost immediately upon the GP's return, typically within one to two weeks. The absence of any significant

[8] The generally large share of individuals not covered by their employer (i.e., children, retirees, and self- or non-employed adults) largely results from the sample's tilt towards younger and older patients due to the GP visit requirement.

Table 1: Summary Statistics for Summer Window Sample

	Never Treated	26	27	28	29	30	31	32	33
Variable		**Treatment Group (First GP Absence Week)**							
Client Demographics									
Age	43 (23)	50 (24)	45 (26)	44 (26)	43 (26)	43 (26)	43 (26)	43 (26)	44 (25)
Female	0.54	0.56	0.54	0.54	0.54	0.54	0.54	0.54	0.54
Swiss Nationality	0.78 (0.42)	0.81 (0.39)	0.79 (0.41)	0.81 (0.40)	0.79 (0.41)	0.79 (0.41)	0.78 (0.41)	0.78 (0.41)	0.76 (0.43)
Insurance Contract									
Deductible	820 (914)	655 (770)	658 (817)	641 (796)	626 (793)	620 (792)	627 (791)	617 (779)	677 (816)
Model									
... Standard	0.26	0.44	0.36	0.33	0.32	0.32	0.33	0.35	0.35
... GP	0.35	0.38	0.41	0.42	0.40	0.38	0.39	0.38	0.39
... HMO	0.25	0.08	0.13	0.16	0.19	0.21	0.19	0.19	0.17
... Telemedicine	0.137	0.097	0.102	0.090	0.091	0.090	0.087	0.082	0.087
Accident Insurance Inclusion	0.55 (0.50)	0.63 (0.48)	0.65 (0.48)	0.65 (0.48)	0.65 (0.48)	0.65 (0.48)	0.64 (0.48)	0.65 (0.48)	0.62 (0.48)
Comorbidity									
Number of PCGs	0.57 (0.97)	0.79 (1.10)	0.69 (1.10)	0.67 (1.10)	0.65 (1.00)	0.64 (1.00)	0.65 (1.00)	0.67 (1.10)	0.66 (1.10)
Weekly HC Utilization (Binary Visit Outcomes)									
Primary Care Physician	0.073 (0.26)	0.082 (0.28)	0.079 (0.27)	0.080 (0.27)	0.079 (0.27)	0.079 (0.27)	0.079 (0.27)	0.080 (0.27)	0.078 (0.27)
Emergency Department	0.0031 (0.056)	0.0029 (0.054)	0.0035 (0.059)	0.0031 (0.056)	0.0033 (0.057)	0.0034 (0.058)	0.0034 (0.058)	0.0034 (0.059)	0.0036 (0.060)
Specialist Physician	0.026 (0.16)	0.036 (0.19)	0.032 (0.18)	0.029 (0.17)	0.029 (0.17)	0.029 (0.17)	0.030 (0.17)	0.032 (0.17)	0.032 (0.18)
Any Healthcare Provider	0.25 (0.44)	0.31 (0.46)	0.29 (0.45)	0.28 (0.45)	0.27 (0.44)	0.27 (0.44)	0.27 (0.45)	0.28 (0.45)	0.28 (0.45)
Weekly HC Expenditures (CHF)									
Total	86 (574)	106 (642)	95 (707)	93 (604)	90 (572)	90 (574)	92 (584)	92 (551)	93 (560)
Pharmaceutical Drugs	20 (260)	27 (256)	23 (204)	23 (260)	22 (240)	22 (262)	22 (232)	22 (242)	21 (194)
Primary Care Services	9.4 (40)	9.3 (36)	9.5 (38)	9.0 (36)	9.2 (37)	9.4 (38)	9.6 (39)	9.9 (40)	9.7 (39)
Physician Services	31 (131)	36 (154)	34 (143)	32 (135)	32 (138)	32 (137)	33 (142)	34 (143)	34 (140)
GP Absence Length									
... 0 Weeks	1	0	0	0	0	0	0	0	0
... 1 Week	0	0.25	0.20	0.13	0.12	0.14	0.24	0.31	0.59
... 2 Weeks	0	0.42	0.42	0.46	0.48	0.53	0.51	0.55	0.29
... 3 Weeks	0	0.33	0.38	0.41	0.40	0.33	0.26	0.14	0.12
Individuals	905,332	47,483	95,110	198,076	310,149	401,404	343,595	145,724	52,920
Individual-Years	2,207,289	59,948	121,433	272,911	446,875	613,321	494,855	175,418	58,330

Notes: Values stem from weeks 20 to 25, before any of the treated units experienced an absence of their regular GP.

"catch-up" spike suggests that the care foregone during the GP's holiday is not simply deferred.

The raw patterns strongly suggest that temporary access barriers lead to a net reduction in healthcare encounters. Our upcoming formal event-study analysis below formalize these patterns and extend them to other utilization categories and spending outcomes, while conditioning on high-dimensional time and individual structure.

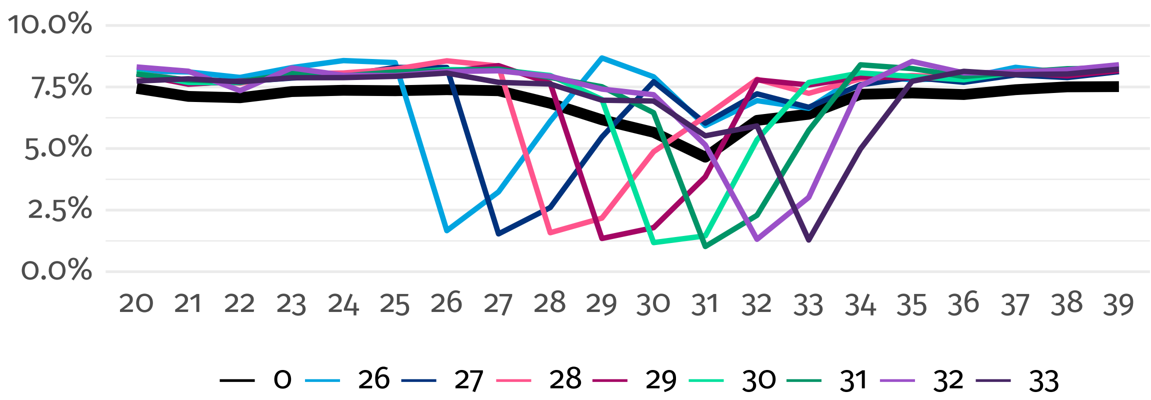

Figure 3: Weekly Probability of a Primary Care Visit by Treatment Group

5 Results

This section quantifies how temporary disruptions to continuity of care—patients' regular GP being on holiday—affect healthcare use and spending, and what these impacts imply for the broader question: when the usual "front door" is briefly shut, do patients substitute to other channels, postpone care, or simply forego it? We first document the main effects on various forms of care-seeking and their associated costs based on event-study estimates for the main summer window (weeks 20–39). We then perform robustness checks and exploit heterogeneity in local market conditions, patient morbidity, absence duration, and insurance design to probe the underlying mechanisms. The evidence consistently demonstrates that relational continuity—the established bond between a patient and their specific GP—is the principal driver of our findings, a friction more powerful than raw provider capacity or nominal insurance plan design.

5.1 Main Effect in the Summer Window

We estimate dynamic effects on weekly outcomes using the two-step procedure described in Section 3.3. For all main estimations in the summer window, the first stage uses 44.15 million never-treated observations; the second stage (which we report the results for) uses 35.01 million treated observations.

Utilization

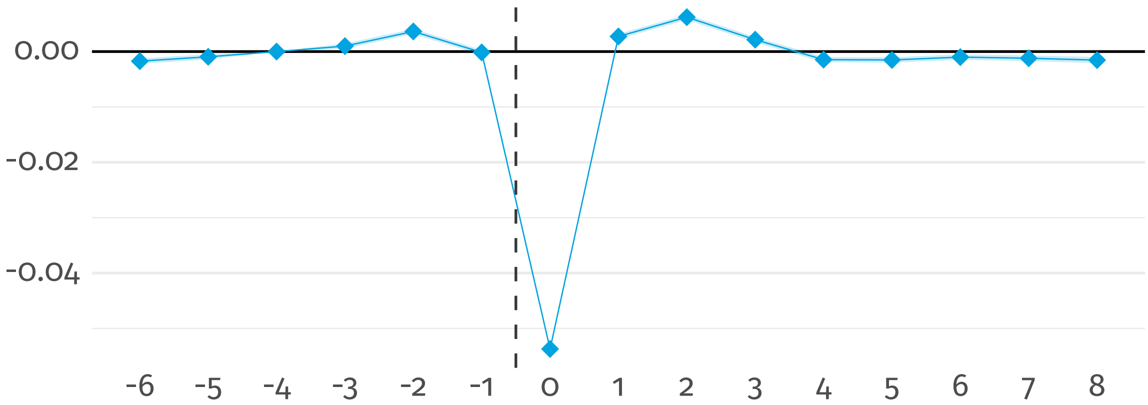

Figure 4: Effect on any Primary Care Physician Visit

Note: Results from the second stage regression with 35.01 million observations and with 44.15 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) is -0.054. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 7.93%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

Primary care visits collapse when the regular GP is away and normalize immediately on return.[9] The event-study estimates in Figure 4 reveal a 5.4 percentage point decline in the weekly probability of any primary care visit during the absence weeks of a patient's regular GP ($k = 0$). Relative to the pre-treatment baseline probability of 7.9% among the treated group, this constitutes a 68% reduction in primary care encounters. The pre-treatment trend is flat, supporting the parallel trends assumption of our research design, and utilization reverts to its baseline level immediately upon the GP's return. The absence of any "catch-up" surge in the subsequent weeks provides the first piece of evidence that the care is foregone, not merely postponed.

A critical question is whether this shock to the primary access point is absorbed elsewhere in the healthcare system. Our findings suggest that substitution is limited and, in most cases, negative. Rather than increasing, specialist visits decrease by a statistically significant 0.3 percentage points from a 3.0% baseline, a 10% relative drop (Figure 5). This implies that the regular GP serves as a crucial, albeit often informal, gateway for specialist referral. In their absence, this channel of care is also constricted, even for patients in plans with unrestricted provider choice.

While we observe a statistically significant increase in Emergency Department (ED) visits of 0.07 percentage points (Figure 6), this effect is economically small. It

[9]As the patient's regular GP is not available, any visits to primary care physicians must be at a still present GP (outside of any hospital setting). Figure A1 in the Appendix shows the evolvement of visits at the respective regular GPs around their absence.

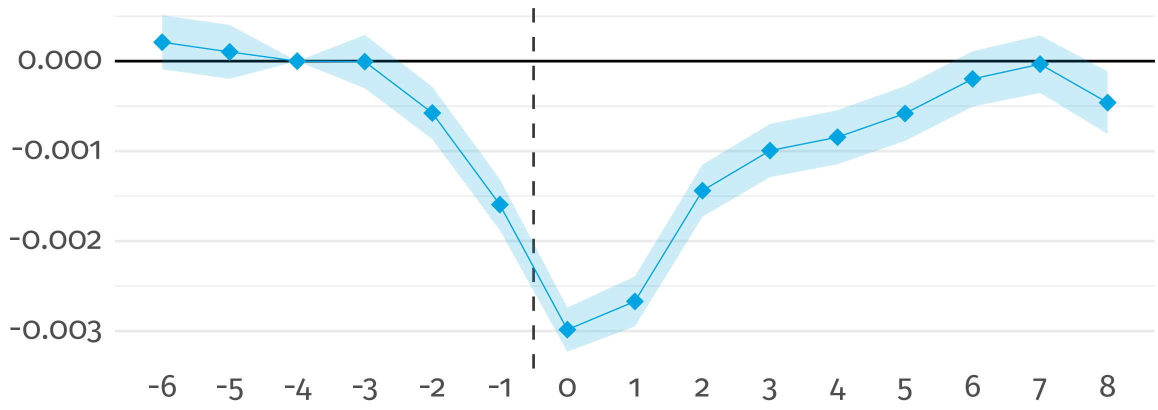

Figure 5: Effect on any Specialist Physician Visit

Note: Results from the second stage regression with 35.01 million observations and with 44.15 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any ED visit in a week) is -0.003. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 3.0%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

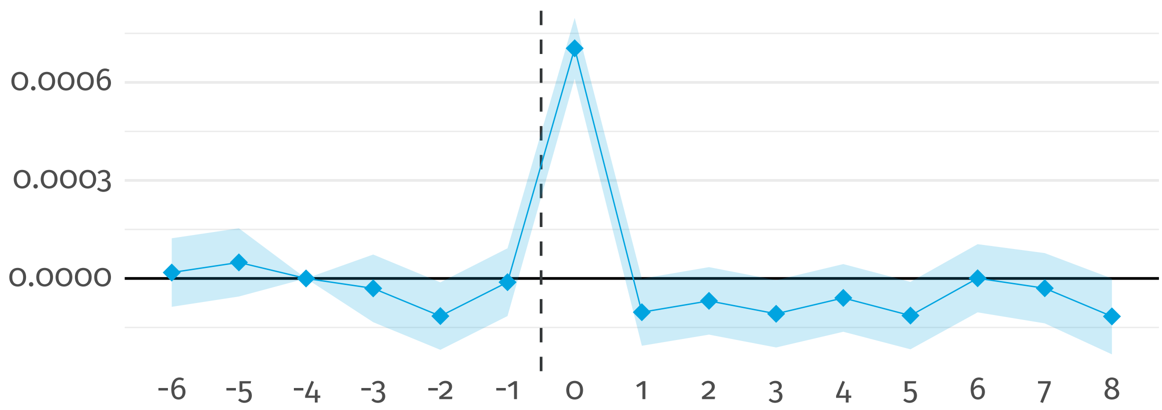

Figure 6: Effect on any Emergency Department Visit

Note: The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any ED visit in a week) is 0.0007. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 0.33%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

is nowhere near large enough to compensate for the massive drop in primary care visits, suggesting that only a tiny fraction of the foregone care was of an urgent nature that demanded immediate alternative attention. This is further corroborated by Figure A2 in the Appendix showing emergency visits at physicians outside of hospitals to be less likely when the regular GP is away than when present. Indeed, when combining ED visits with other physician emergency services, we find a net decline in urgent care encounters (Appendix Figure A3).

Crucially, we find a precise null effect on the start of inpatient hospital stays (Appendix Figure A4). This result provides strong evidence against the hypothesis that the foregone care leads to a deterioration of health that requires acute, high-cost intervention. The shock to primary care access is therefore not absorbed by the system but leads to a net reduction in patient-provider interactions, as shown by the 6.7 percentage point (a 25% relative decline) in the probability of visiting *any* healthcare provider in Figure 7 in the Appendix.

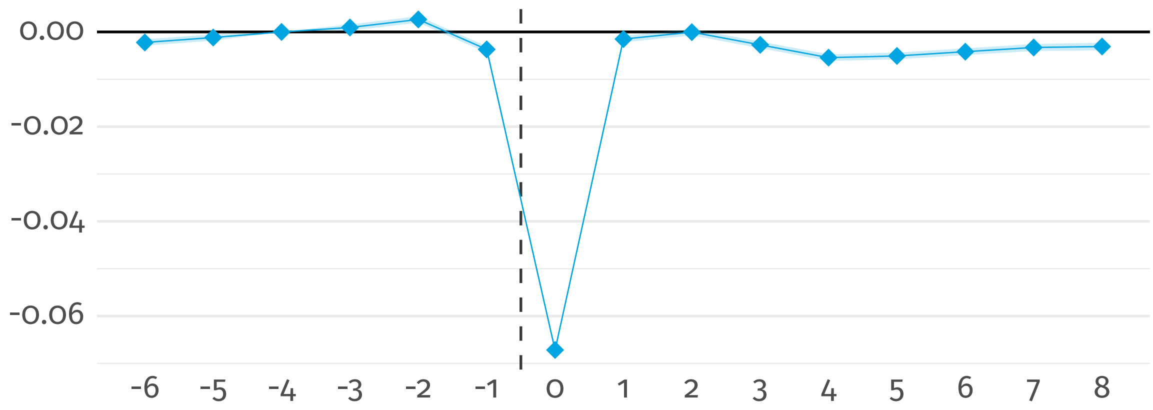

Figure 7: Effect on any Healthcare Provider Visit

Note: Results from the second stage regression with 35.01 million observations and with 44.15 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any healthcare provider visit in a week) is -0.067. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 27.4%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

Healthcare Expenditures

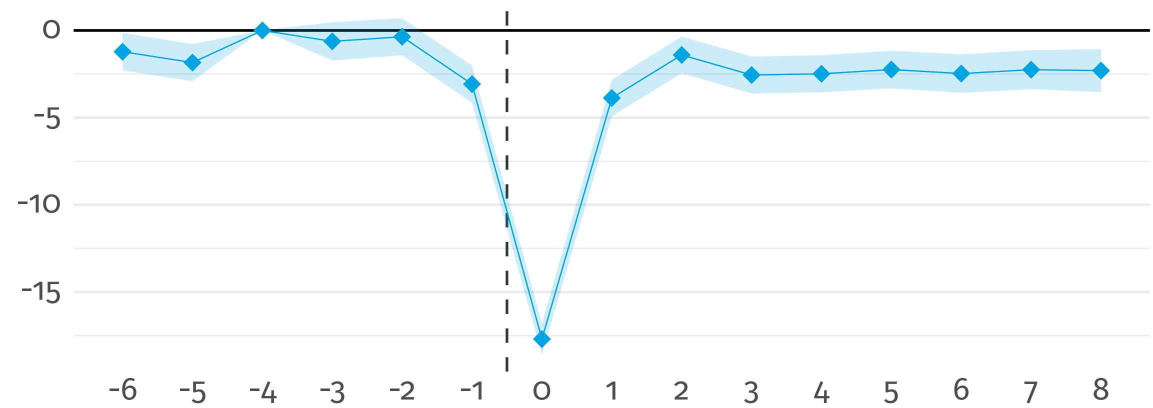

Figure 8: Effect on total costs

Note: Results from the second stage regression with 35.01 million observations and with 44.15 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on total healthcare expenditures in a week) is CHF -17.7. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 91.8. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

The sharp decline in utilization translates directly into a significant reduction in healthcare expenditures. As shown in Figure 8, total weekly healthcare expenditures fall by CHF 17.7, a 19% reduction from a baseline of CHF 91.8 for the treated group. This reduction is broad-based and not confined to primary care services alone. Total physician costs fall by CHF 7.7 (a 24% relative decline, Figure 9), while prescription drug costs decrease by CHF 5.3 (again a 24% relative decline, Figure 10). The fall in pharmaceutical spending is particularly telling, as it indicates that the GP's role as a prescriber is not easily substitutable, and that this entire chain of care delivery is disrupted alongside the consultation itself. Non-physician outpatient costs (e.g., labs, physiotherapy) also contract by CHF 2.6 (once more a 24% relative reduction, Appendix Figure A7). Consistent with the utilization patterns, there is no evidence of a post-absence cost surge, indicating that the temporary access friction leads to a net system-wide savings without subsequent financial payback through more severe health episodes due to the forgone care.

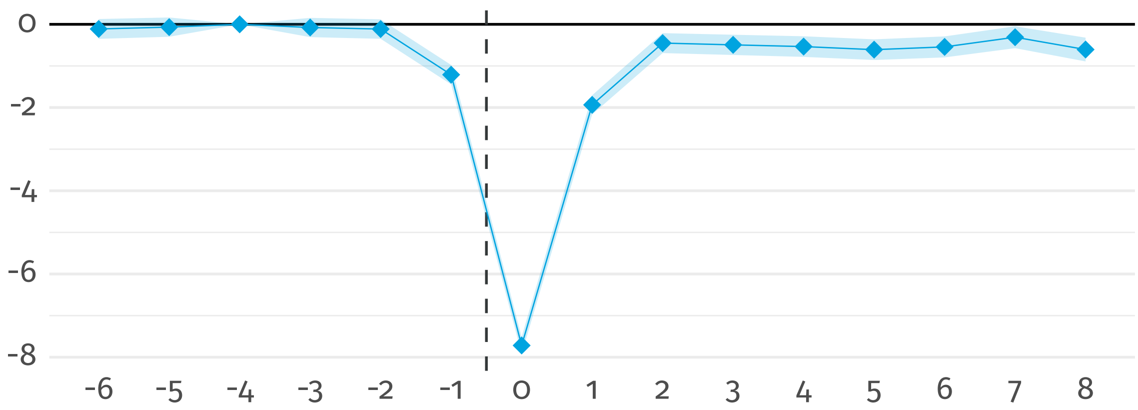

Figure 9: Effect on physician costs

Note: Results from the second stage regression with 35.01 million observations and with 44.15 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on physician expenditures in a week) is CHF -7.7. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 32.5. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

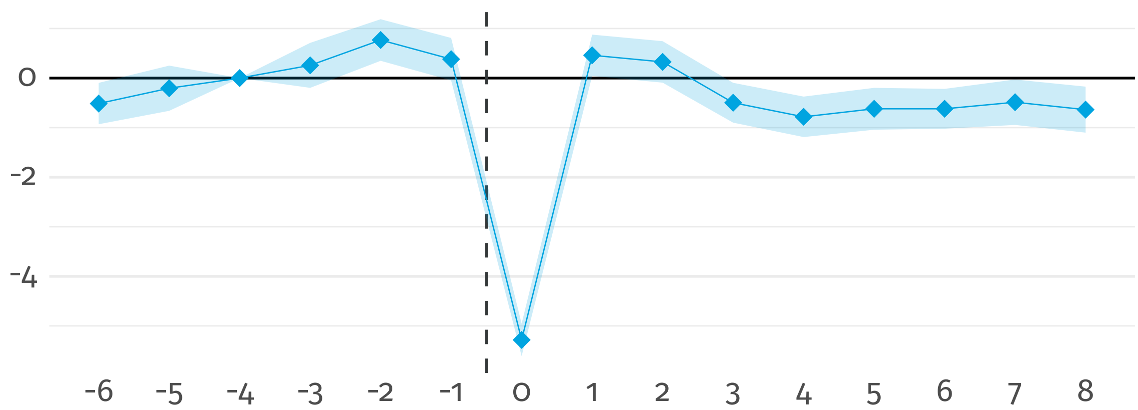

Figure 10: Effect on prescription drug costs

Note: Results from the second stage regression with 35.01 million observations and with 44.15 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on prescription drug expenditures in a week) is CHF -5.3. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 22.1. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

Robustness Checks

Appendix Figures A8 and A9 reprises the outcomes in the main summer window for primary care and total healthcare expenditures on a fully balanced support, with all treated units being observed throughout each of the $k \in \{-6, -5, ..., 7, 8\}$ weeks. The results remain highly similar to those from unbalanced sample around the event window.

We apply the same design as in the summer window to spring (weeks 3–25) and fall (weeks 35–51). These exercises yield the qualitatively same results for primary care visits and total costs as in the main summer sample as indicated by Appendix Figures A10, A11, A12, and A13.

5.2 Heterogeneity

To disentangle the mechanisms behind these effects, we explore heterogeneity across several key dimensions. The evidence consistently points toward the central importance of the established doctor-patient relationship—what we term relational continuity—over general market capacity or insurance-based channeling.

5.2.1 Regular GP Absence Length

The per-week effect on utilization and spending is stable regardless of whether the GP is absent for one, two, or three weeks (Figures 11 and 12). This implies that the decision to forego care is based on the immediate unavailability of the GP, and that the total, cumulative amount of foregone care and associated savings scales directly with the length of the disruption.

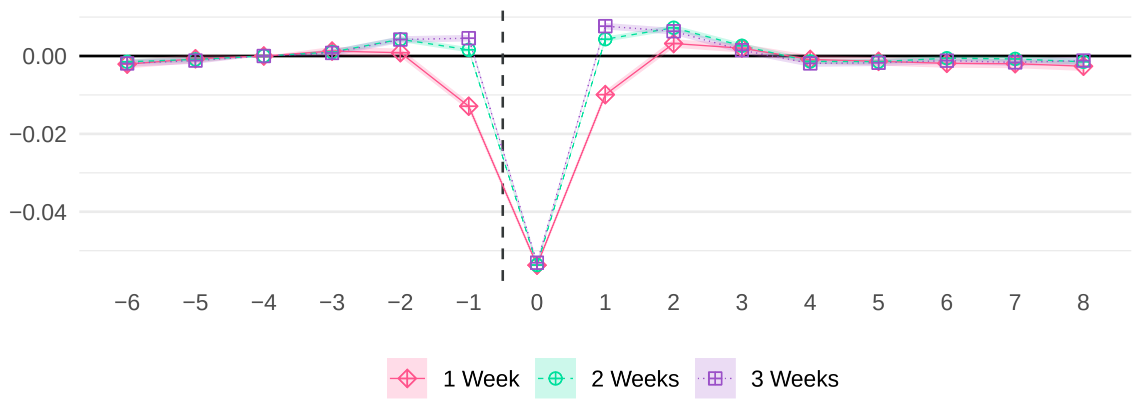

Figure 11: Effect on any Primary Care Visit by GP absence length

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) are -0.054, -0.054, and -0.053. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 7.97%, 7.97%, and 7.85%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

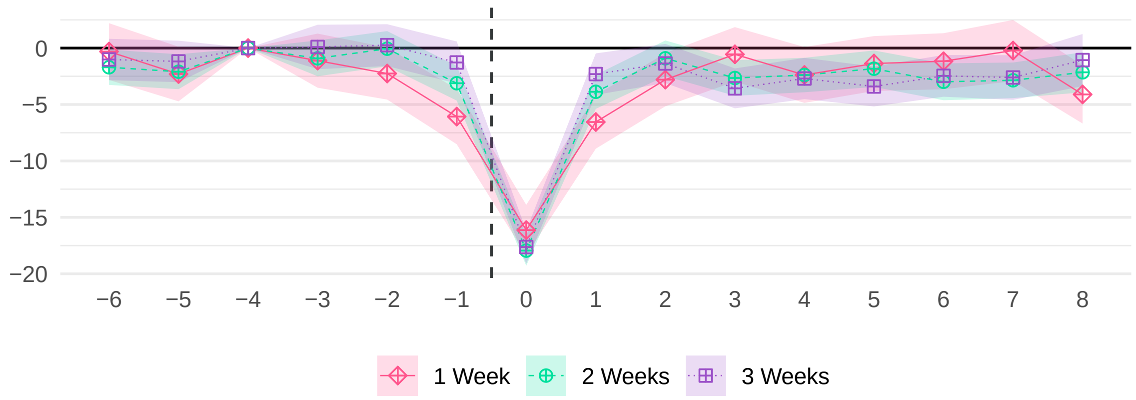

Figure 12: Effect on Total Health Care Expenditures by GP absence length

Note: The aggregated coefficient are (average treatment effect of regular GP absence on total healthcare expenditures in a week) are CHF -16.1, CHF -17.9, and CHF -17.6. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 91.7, CHF 91.2, and CHF 92.7. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

Regardless of absence length, we do not observe any noticeable increase in utilization or costs across the nearly two-month period after the GP's return. If anything, health care demand persistently remains lower the longer an absence lasted. Even with a three-week absence of a patient's regular GP (and thus a substantial amount of forgone primary care visits) no negative consequences for patient health are evident. The lack of any detrimental health ramifications of GP absences becomes even more surprising given the substantially greater decline in probability to visit a specialist physician for three-week absences, as evident in Appendix Figure A14. This finding once again indicates the importance of GPs as entry points to the healthcare system. Without access to primary care patients appear to—rather than seek substitutes—forego even more care at other providers, amplifying the net cost consequences of access to the regular GP.

5.2.2 Number of Locally Available GPs

Regional GP scarcity represents a policy concern for many healthcare systems. We test whether the negative demand effect is attenuated by greater local supply, a proxy for lower search and travel costs. We stratify our sample by the density of other GPs in the patient's local market and find that the reduction in primary care visits is statistically indistinguishable across areas with few (1-3), moderate (4-10), or many (more than 10) GPs (Figures 13 and 14).

Likewise, the overall drop in total spending remains stable across these strata (Figure 14). Figure A15 in the appendix shows that the same applies to the for trivial increase in ED use. The presence of more alternative providers does not mitigate the effect of one's usual GP being unavailable. This powerful null result suggests the friction is not a standard supply constraint, but is instead rooted in the specific,

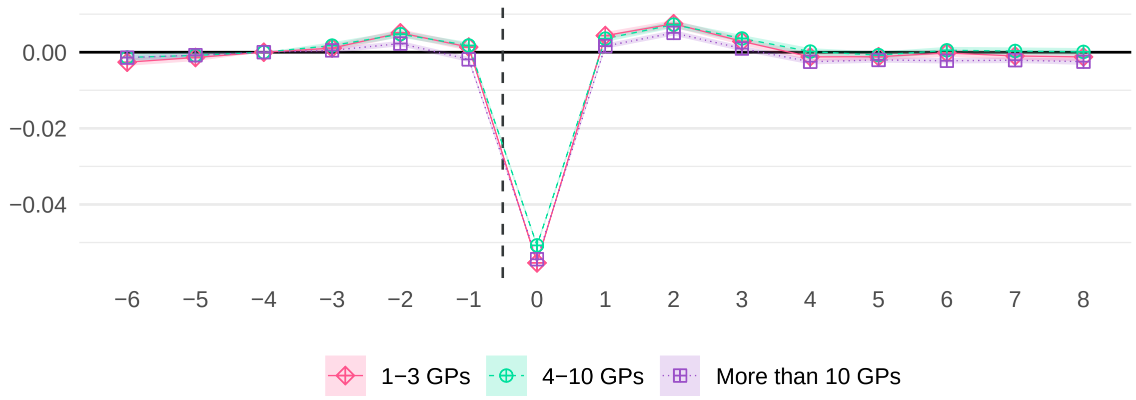

Figure 13: Effect on any primary care visit by Number of Local GPs

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) are -0.055, -0.051, and -0.054. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 7.84%, 7.85%, and 8.03%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

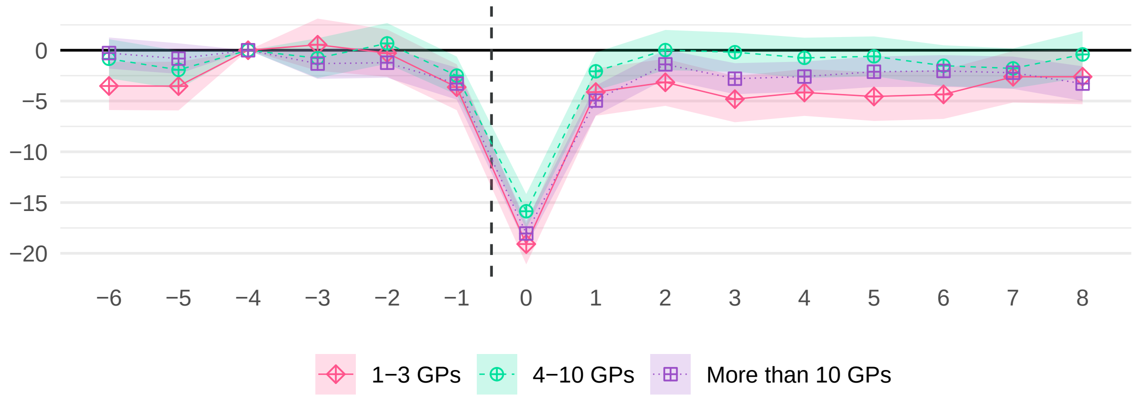

Figure 14: Effect on Total Health Care Expenditures by Number of Local GPs

Note: The aggregated coefficient are (average treatment effect of regular GP absence on total healthcare expenditures in a week) are CHF -19.1, CHF -15.9, and CHF -18.0. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 93.5, CHF 89.1, and CHF 92.5. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

non-transferable nature of the patient-provider relationship. Policies that marginally alter total primary care capacity margin may thus only have limited impact—at least in settings where patients rarely visit their GP for medically urgent care as it appears to be the case in Switzerland.

5.2.3 Chronic Conditions

Next, we explore how the response varies with patient health, a proxy for the baseline need for care. Figures 15 and 16 reveal that the effects are largest in absolute terms for patients with more comorbidities. Individuals with two or more chronic conditions—whose baseline utilization is highest—experience a primary care visit reduction of 9.3 percentage points and a spending decrease of CHF 37.8 per week, far exceeding

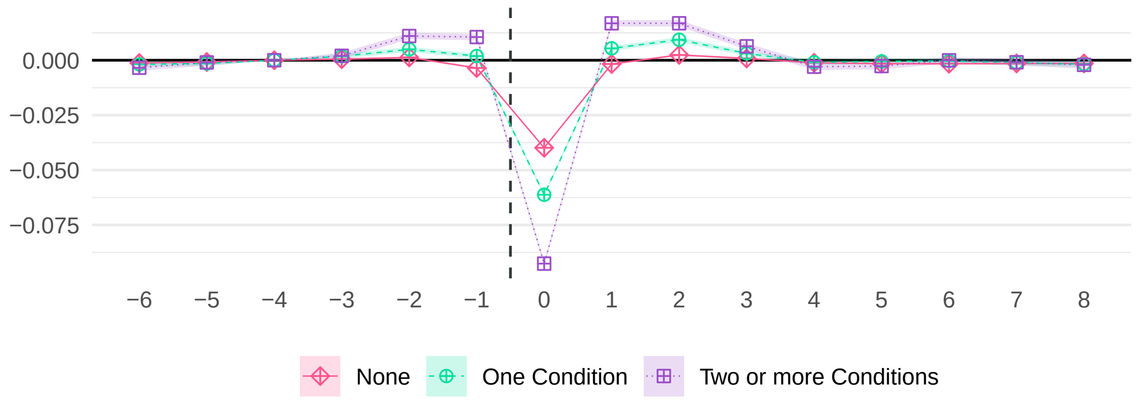

Figure 15: Effect on any primary care visit by Number of Chronic Conditions

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) are -0.040, -0.061, and -0.093. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 6.04%, 9.15%, and 13.26%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

the declines seen among healthier individuals. This finding presents a key insight regarding relational frictions: the patients with the most pressing objective need for care are no more likely to seek a substitute when their GP is unavailable. This suggests that for patients with complex histories, the informational and trust-based capital embodied in the existing relationship with their GP creates a switching cost so high that it dominates their greater medical need. Nonetheless, while we observe some indication of minor increases in anticipatory and catch-up visits for sicker patients, the probability to forgo otherwise sought care is roughly the same irrespective of chronic health status.

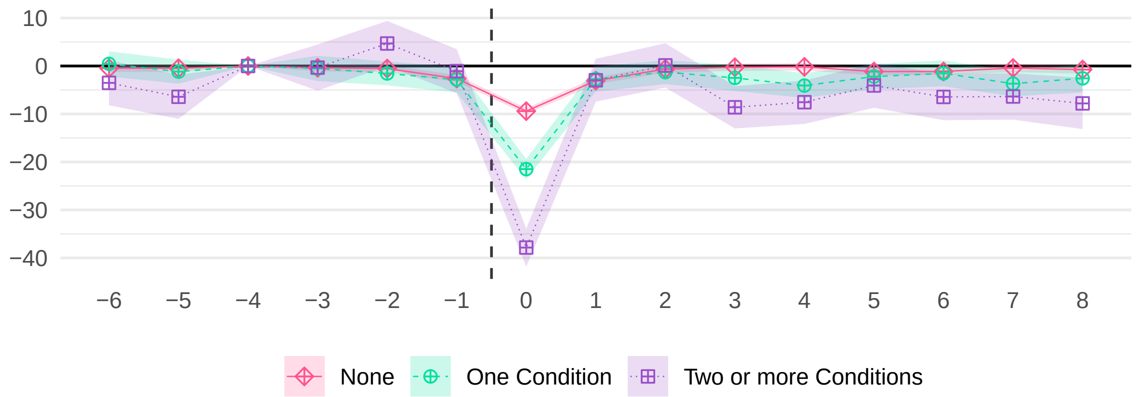

Figure 16: Effect on Total Health Care Expenditures by Number of Chronic Conditions

Note: The aggregated coefficient are (average treatment effect of regular GP absence on total healthcare expenditures in a week) are CHF -9.4, CHF -21.5, and CHF -37.8. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 39.3, CHF 113.7, and CHF 252.1. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

5.2.4 Insurance Model

Finally, our analysis by insurance plan type (Figures 17 and 18 as well as A16 in the Appendix) shows that while patients in more structured HMO and telemedicine plans have lower baseline use and slightly smaller absolute reductions, they still forego the vast majority of care. Interestingly, patients in "standard" plans with unrestricted choice exhibit the largest absolute declines. This suggests that in the absence of formal network constraints, patients form even stronger de facto dependencies on their chosen regular GP, making substitution less likely when that specific provider is unavailable.

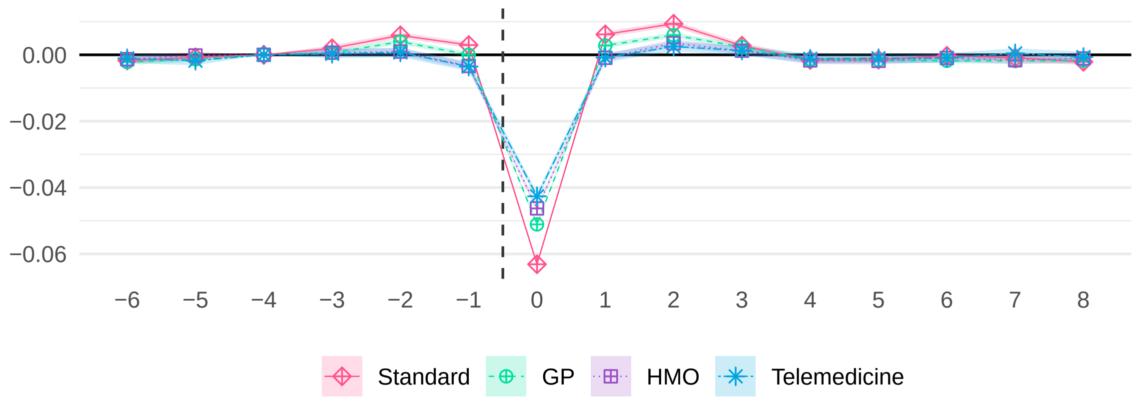

Figure 17: Effect on any Primary Care Visit by Insurance Model

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) are -0.063, -0.051, -0.046, and -0.043. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 9.06%, 7.70%, 7.36%, and 5.96%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

In addition, the effect is larger in weeks when patients face a higher marginal price (deductible not yet met). Appendix Figure A17 shows a -4.0 percentage point decline under full co-pay (baseline 3.9%) versus -6.8 percentage points under lower marginal price (baseline 10.9%). These estimates have to be interpreted cautiously given the markedly differing baselines across these strata. Nonetheless, while the pattern is consistent with the marginal price of seeking care amplifying the continuity shock, the non-monetary friction seems to be more important than the financial incentives built into MHI designed to prevent overuse of the healthcare system.

In sum, our results show that a temporary, planned disruption to a patient's regular GP causes a large, immediate, and net reduction in healthcare use and spending. This effect is not driven by a lack of local provider capacity but by a breakdown in relational continuity, a non-monetary friction that proves paramount. Patients do not meaningfully substitute to other providers, even in markets with high physician density; instead, they simply forego care. This pattern is most pronounced in absolute terms among chronically ill patients, for whom continuity is arguably most critical. The lack of any subsequent "catch-up" in visits or costs, combined with a null effect on

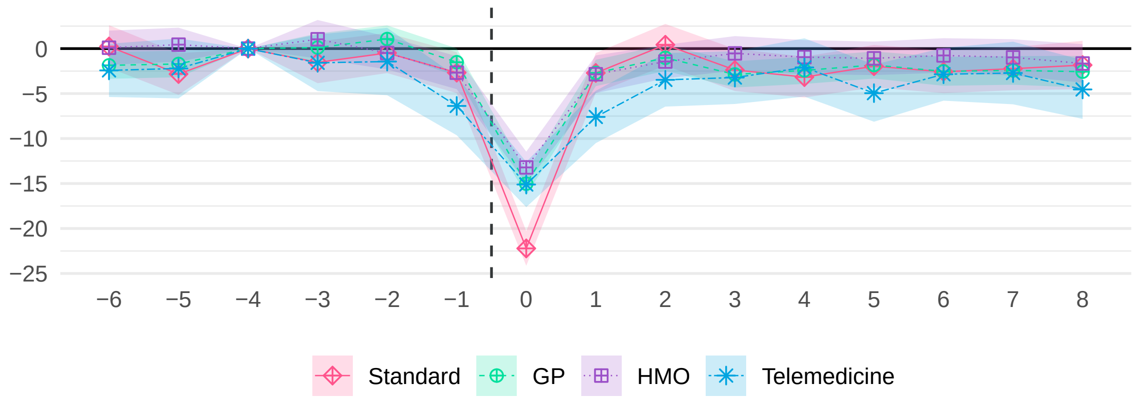

Figure 18: Effect on Total Health Care Expenditures by Insurance Model

Note: The aggregated coefficient are (average treatment effect of regular GP absence on total healthcare expenditures in a week) are CHF -22.2, CHF -15.0, CHF -13.2, and CHF -15.1. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 125.3, CHF 81.2, CHF 69.0, and CHF 61.5. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year-level.

inpatient admissions, strongly suggests that the majority of these foregone services were not medically urgent and could be skipped without observable, adverse short-term health consequences. This research highlights that in understanding healthcare demand, the specific, established relationship between a patient and their physician can be a more binding constraint than either price or the raw supply of providers.

5.3 Discussion

Our results point to three sets of implications: for the role of non-monetary frictions in healthcare demand, for measuring discretionary care, and for understanding the GP's function in the healthcare system.

Relational frictions as a first-order determinant of demand

The magnitude of the demand response we document—a roughly two-thirds reduction in primary care visits from the temporary unavailability of a single provider—exceeds, by a wide margin, estimated demand responses to price variation. The RAND Health Insurance Experiment, the canonical reference, finds a price elasticity of approximately -0.2 (Aron-Dine et al., 2013). A separate literature documents significant but more modest effects of travel time and distance (see for example, Sabety et al., 2023). Our estimates imply that the implicit cost of consulting an unfamiliar provider is, for most patients, prohibitively high relative to the perceived benefit of the visit.

The nature of this friction is clarified by our heterogeneity analysis. If the friction were primarily about search or travel costs, it should dissipate in markets with more alternative providers. The null interaction with local GP density rules out this channel and points instead to the non-transferable, relationship-specific capital between patient and provider. This capital—comprising accumulated clinical knowledge, mutual familiarity, and trust—helps patients resolve the information asymmetry inherent in healthcare (Arrow, 1963). When that specific relationship is disrupted, patients appear to judge the cost of rebuilding it with an unknown provider as too high to justify the visit, at least for routine care.

The finding that chronically ill patients experience the largest absolute declines in utilization deepens this interpretation. For patients with complex medical histories, the informational component of relational capital is most valuable and most costly to reconstruct. That these patients are *less* likely to substitute, despite having the greatest clinical need, is consistent with relationship-specific capital that is an increasing function of medical complexity.

Implications for discretionary and low-value care

The absence of catch-up utilization or spending, combined with a null effect on hospitalizations in the two months following the GP's return, suggests that the vast majority of foregone care was not acutely necessary—at least over this horizon. This provides a behavioral complement to the clinical-guidelines-based approach to measuring low-value care. Rather than relying on ex ante clinical judgment, our design lets patients themselves reveal which services they are willing to forego when faced with a relatively modest access barrier.

We emphasize two important caveats. First, the absence of short-run adverse events does not establish the absence of longer-run health costs. Deferred preventive care or chronic-disease management could have consequences that materialize over months or years, beyond our observation window. Our estimates should therefore be interpreted as an upper bound on the share of routine care that can be foregone *without short-term health consequences*, not a welfare statement about the value of that care. Second, patient willingness to forego care does not imply that the care had zero value *ex ante*. Patients may rationally choose to wait for their trusted GP rather than seek lower-expected-quality care from an unfamiliar provider, even for visits with positive expected value.

The GP as gatekeeper and demand catalyst

A striking feature of our results is that the GP's absence reduces not only primary care but the entire downstream chain: specialist visits, prescriptions, and non-physician outpatient services. This occurs even among patients whose insurance permits direct specialist access, and the effect is largest for this group. The pattern suggests that patients rely on their GP not only for direct treatment but for the information, validation, and referrals that initiate demand for other services. In the absence of formal network constraints, this self-imposed "behavioral gatekeeping" appears to be at least as binding as contractual gatekeeping models.

This finding refines the theory of supplier-induced demand. Classical SID models (e.g., Mcguire, 2000; Iizuka, 2012; Clemens and Gottlieb, 2014) emphasize the physician's ability to shift patient demand outward. Our results suggest that this demand-inducing capacity is relationship-specific: it is the *specific*, *trusted* GP who catalyzes demand, and when that agent is removed, the induced demand largely disappears rather than being redirected to other available suppliers. This has implications for policies focused on aggregate provider supply—such as increasing the number of GPs in underserved areas—which may have smaller effects on total utilization than models assuming fungible providers would predict. Our finding that local GP density does not buffer the access shock is consistent with this interpretation.

6 Conclusion

We exploit temporary holiday absences of general practitioners in Switzerland to estimate the causal effect of primary care access on healthcare demand. When a patient's regular GP is away for one to three weeks, the weekly probability of a primary care visit falls by roughly two-thirds, total healthcare expenditures decline by 19%, and substitution to other providers is virtually nonexistent. Utilization reverts to baseline immediately upon the GP's return, with no evidence of catch-up demand or adverse short-run health events.

Three findings stand out. First, the patient–physician relationship is a first-order determinant of demand, generating a friction that far exceeds documented responses to prices or travel costs. The effect is not attenuated by local provider density, confirming that what matters is not the availability of *a* doctor but the availability of *the patient's* doctor. Second, the volume of care that can be foregone without short-term consequences—roughly three-quarters of weekly primary care visits and one—fifth of total spending—provides a revealed-preference benchmark for the share of routine care that is effectively discretionary. Third, the GP serves as a de facto gatekeeper and demand catalyst whose temporary removal constricts the entire care pathway, including downstream specialist and pharmaceutical services, even in insurance plans with unrestricted provider choice.

Several limitations deserve emphasis. Our post-period spans approximately two months, so we cannot rule out adverse health consequences at longer horizons. Our setting is Switzerland—a wealthy country with high physician density, universal insurance, and fee-for-service reimbursement—and the magnitude of relational frictions may differ in systems with formal patient registration (e.g., the UK's NHS) or more fragmented insurance markets (e.g., the United States). Finally, our administrative data do not reveal the clinical reasons for foregone visits, limiting our ability to assess whether the skipped care was preventive, chronic-disease management, or acute. Combining this design with clinical data from electronic health records is a natural next step.

Despite these caveats, our results demonstrate that the patient-physician relationship is not a peripheral feature of healthcare markets but a fundamental institution governing demand. Policies aimed at improving system efficiency should therefore contend with the powerful role of relational capital—and consider how to support its formation, maintenance, and, when necessary, efficient transfer.

References

Aron-Dine, Aviva, Liran Einav, and Amy Finkelstein, "The RAND Health Insurance Experiment, Three Decades Later," *Journal of Economic Perspectives*, 2013, *27* (1), 197–222.

Arrow, Kenneth J., "Uncertainty and the Welfare Economics of Medical Care," *American Economic Review*, 1963, *53* (5), 941–973.

Baicker, Katherine and Amitabh Chandra, "Medicare Spending, the Physician Workforce, and Beneficiaries' Quality of Care: Areas with a High Concentration of Specialists Also Show Higher Spending and Less Use of High-Quality, Effective Care.," *Health Affairs*, 2004, *23* (Suppl1), W4–184–W4–197.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad, "What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics*," *The Quarterly Journal of Economics*, 2017, *132* (3), 1261–1318.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-In-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson, "Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector," *American Economic Review*, 2016, *106* (8), 2110–2144.

Clemens, Jeffrey and Joshua D. Gottlieb, "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?," *American Economic Review*, 2014, *104* (4), 1320–1349.

Cutler, David, Jonathan S. Skinner, Ariel Dora Stern, and David Wennberg, "Physician Beliefs and Patient Preferences: A New Look at Regional Variation in Health Care Spending," *American Economic Journal: Economic Policy*, 2019, *11* (1), 192–221.

Dorn, Michael, "Erfahrungen der Wohnbevölkerung ab 18 Jahren mit dem Gesundheitssystem – Situation in der Schweiz und im internationalen Vergleich. Analyse des International Health Policy (IHP) Survey 2023 der amerikanischen Stiftung Commonwealth Fund (CWF) im Auftrag des Bundesamtes für Gesundheit (BAG)," Technical Report, Schweizerisches Gesundheitsobservatorium, Neuchâtel 2023.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Place-Based Drivers of Mortality: Evidence from Migration," *American Economic Review*, 2021, *111* (8), 2697–2735.

Gardner, John, Neil Thakral, Linh T Tô, and Luther Yap, "Two-Stage Differences in Differences," *Unpublished*, 2025.

Handel, Benjamin R, "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

Handel, Benjamin R. and Jonathan T. Kolstad, "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 2015, *105* (8), 2449–2500.

Iizuka, Toshiaki, "Physician Agency and Adoption of Generic Pharmaceuticals," *American Economic Review*, 2012, *102* (6), 2826–2858.

Liu, Licheng, Ye Wang, and Yiqing Xu, "A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data," *American Journal of Political Science*, 2024, *68* (1), 160–176.

Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, Arleen Leibowitz, and M. Susan Marquis, "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment.," *American Economic Review*, 1987, *77*, 251–277.

Mcguire, Thomas G., "Physician Agency," in Anthony J. Culyer and Joseph P. Newhouse, eds., *Handbook of Health Economics*, Elsevier, 2000, pp. 461–536.

Roth, Jonathan, Pedro H.C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-In-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Sabety, Adrienne, Jonathan Gruber, Jin Yung Bae, and Rishi Sood, "Reducing Frictions in Health Care Access: The ActionHealthNYC Experiment for Undocumented Immigrants," *American Economic Review: Insights*, 2023, *5* (3), 327–346.

Schwartz, Aaron L., Bruce E. Landon, Adam G. Elshaug, Michael E. Chernew, and J. Michael McWilliams, "Measuring Low-Value Care in Medicare," *JAMA Internal Medicine*, 2014, *174* (7), 1067.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A Appendix A

A.1 Additional Results

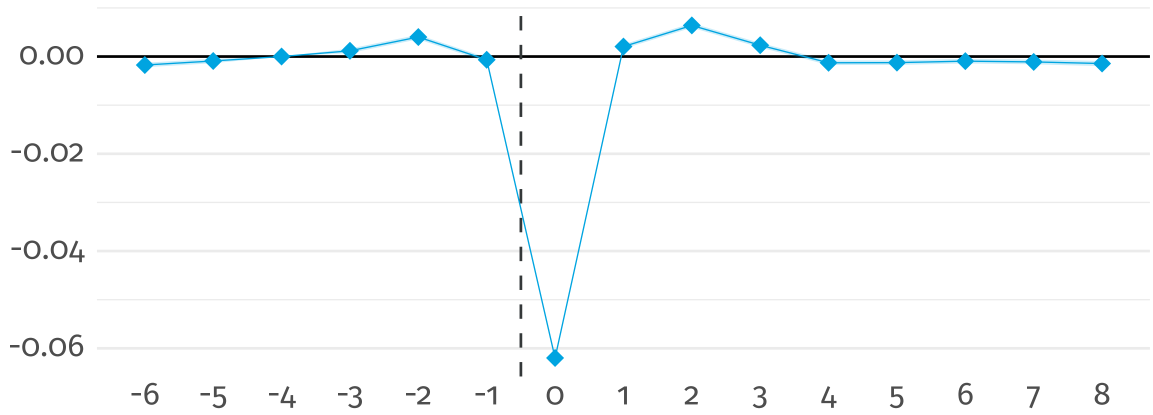

Appendix Figure A1: Effect on any Visit at the Regular GP

Note: The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any primary care or ED visit in a week) is -0.053. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 8.18%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

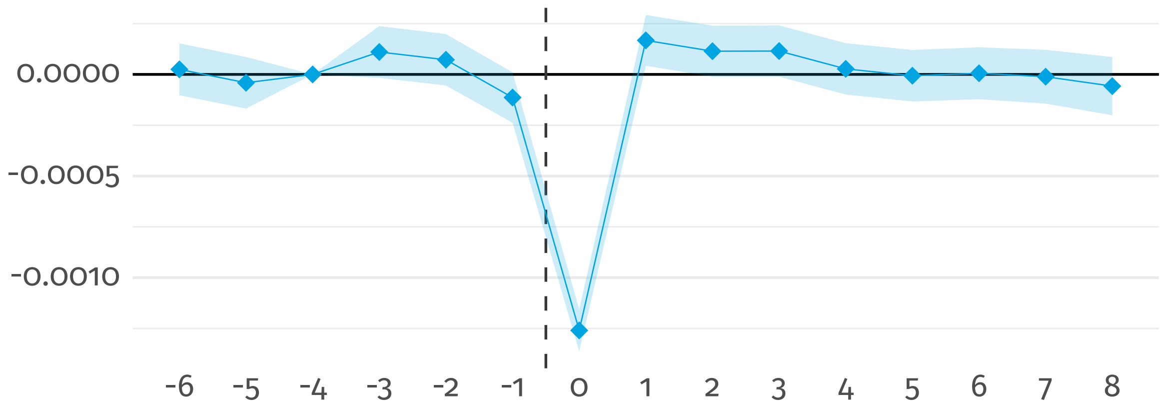

Appendix Figure A2: Effect on Emergency Visit at any Physician Outside the ED

Note: The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any emergency physician visit outside of the ED in a week) is -0.0013. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 0.49%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

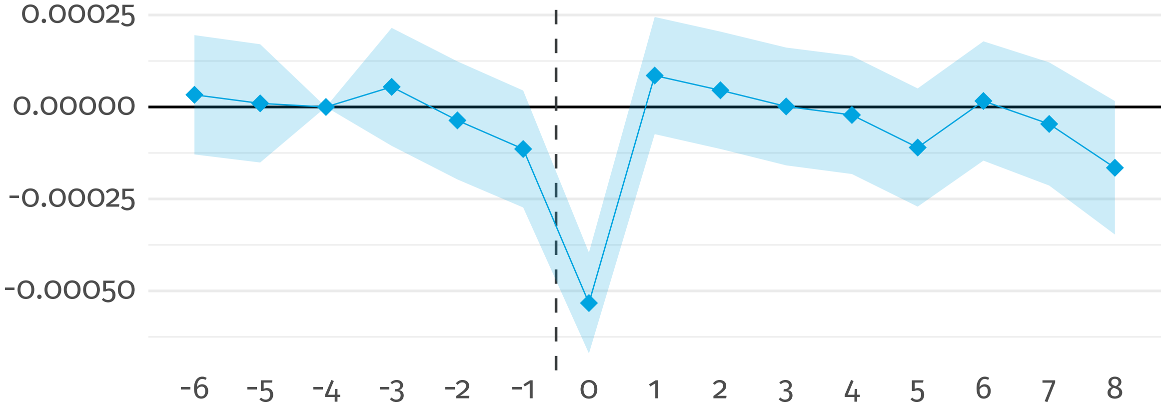

Appendix Figure A3: Effect on Emergency Visit at Primary Care Physicians or ED

Note: The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any emergency physician or ED visit in a week) is -0.0005. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 0.80%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

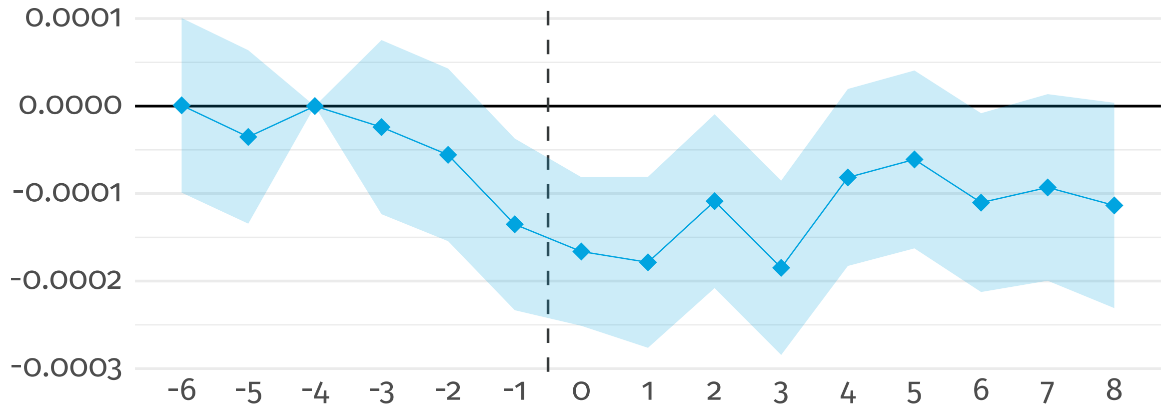

Appendix Figure A4: Effect on Start of a Inpatient Stay

Note: The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of the start of a inpatient stay in a week) is -0.0002. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 0.30%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

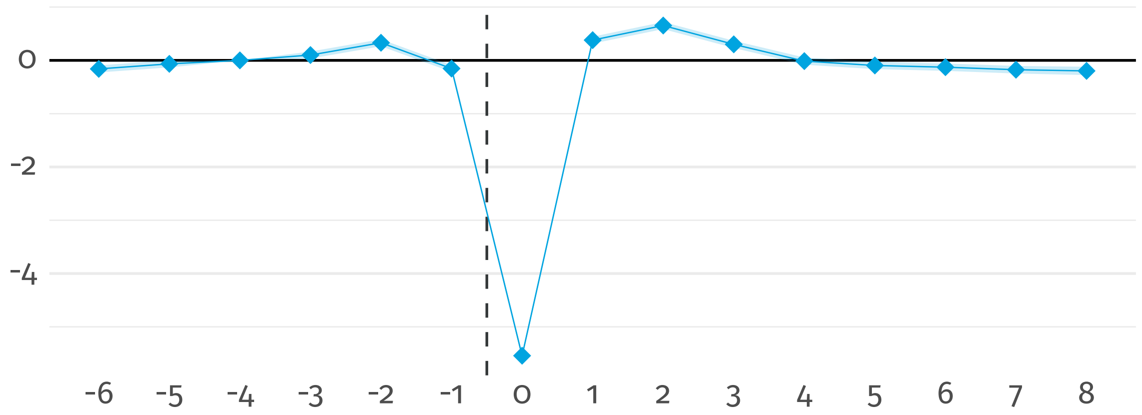

Appendix Figure A5: Effect on Primary Care Costs

Note: The aggregated coefficient (average treatment effect of regular GP absence on primary care physician expenditures in a week) is CHF -5.5. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 9.4. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

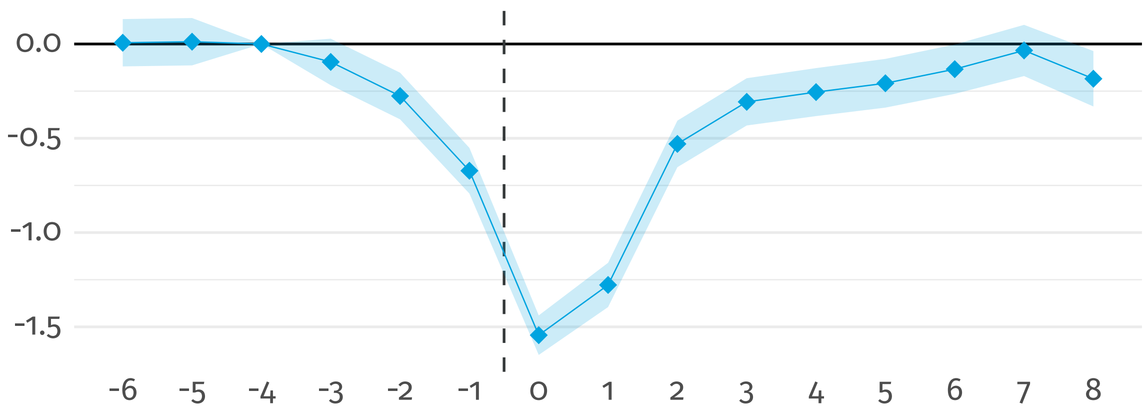

Appendix Figure A6: Effect on Specialist Physician Costs

Note: The aggregated coefficient (average treatment effect of regular GP absence on specialist physician expenditures in a week) is CHF -1.5. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 11.2. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

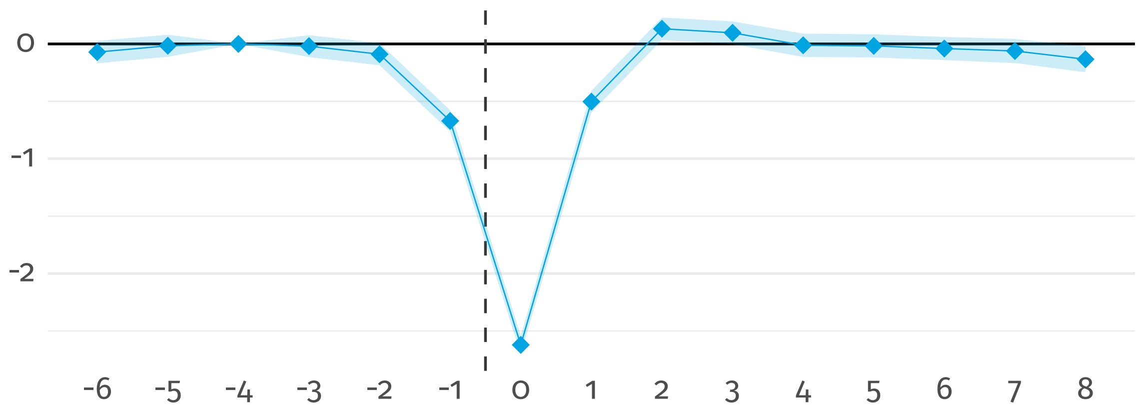

Appendix Figure A7: Effect on Non-Physician Outpatient costs

Note: The aggregated coefficient (average treatment effect of regular GP absence on non-physician/drug outpatient expenditures in a week) is CHF -2.6. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 10.7. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

A.2 Robustness

A.2.1 Fully Balanced

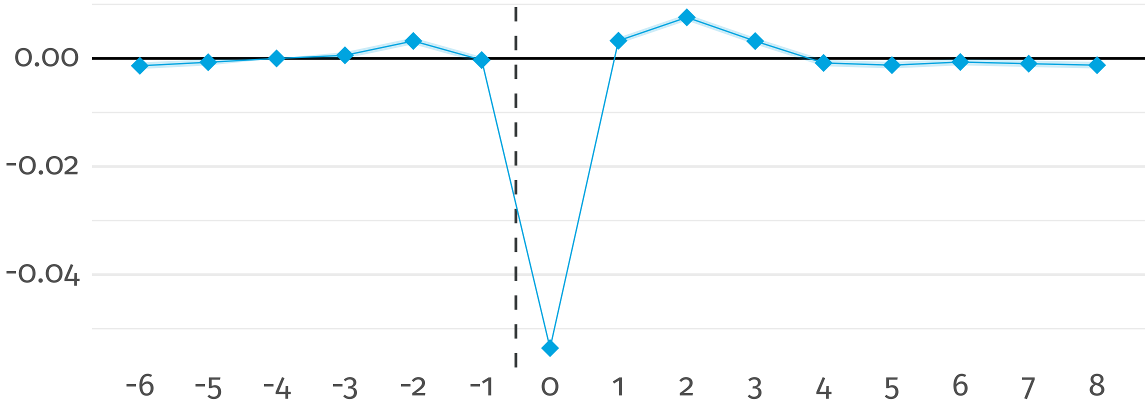

Appendix Figure A8: Effect on any Primary Care Physician Visit with Fully Balanced Sample

Note: Results from the second stage regression with X million observations and with 44.1 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) is -0.054. The baseline probability of the outcome in weeks 20 to 25 among the treated population is 7.95%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

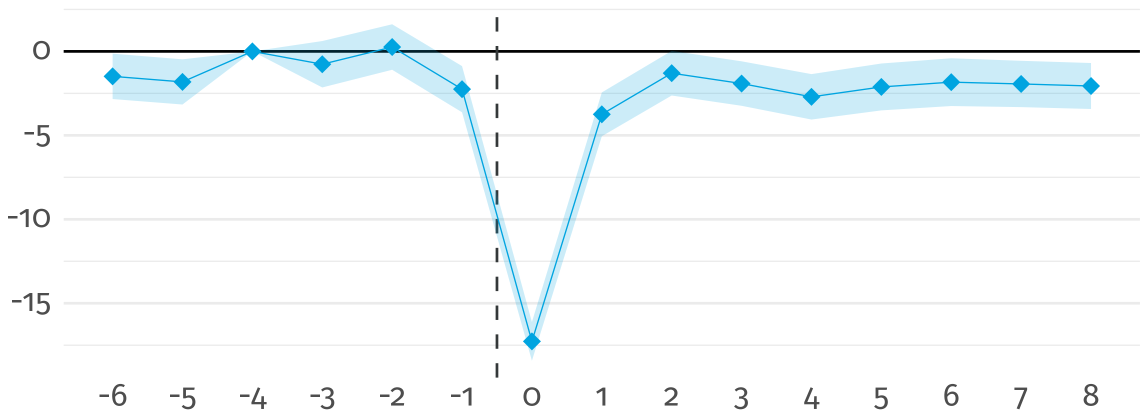

Appendix Figure A9: Effect on Total Costs with Fully Balanced Sample

Note: The aggregated coefficient (average treatment effect of regular GP absence on total healthcare expenditures in a week) is CHF -17.3. The baseline costs of the outcome in weeks 20 to 25 among the treated population are CHF 91.3. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

A.2.2 Spring Window (Weeks 3 to 25)

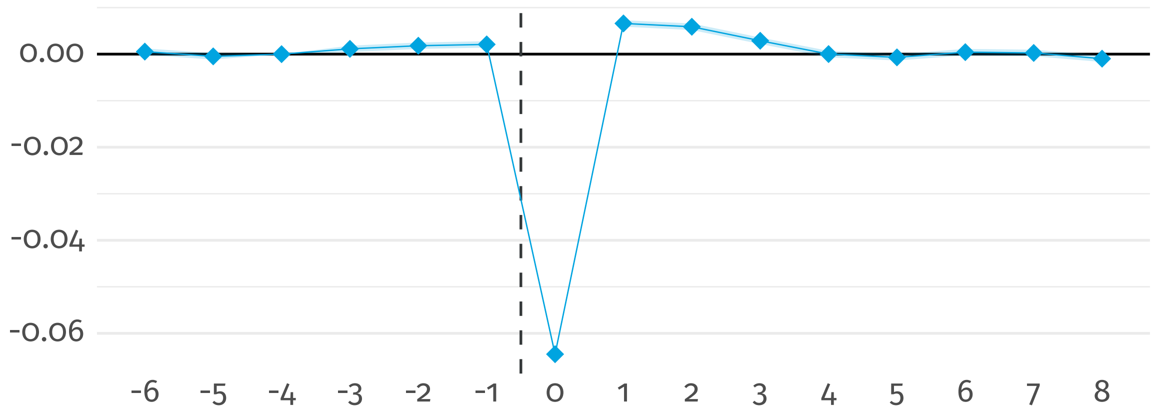

Appendix Figure A10: Effect on any Primary Care Physician Visit in Spring Window

Note: Results from the second stage regression with 17.7 million observations and with 64.0 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) is -0.064. The baseline probability of the outcome in weeks 3 to 8 among the treated population is 8.93%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

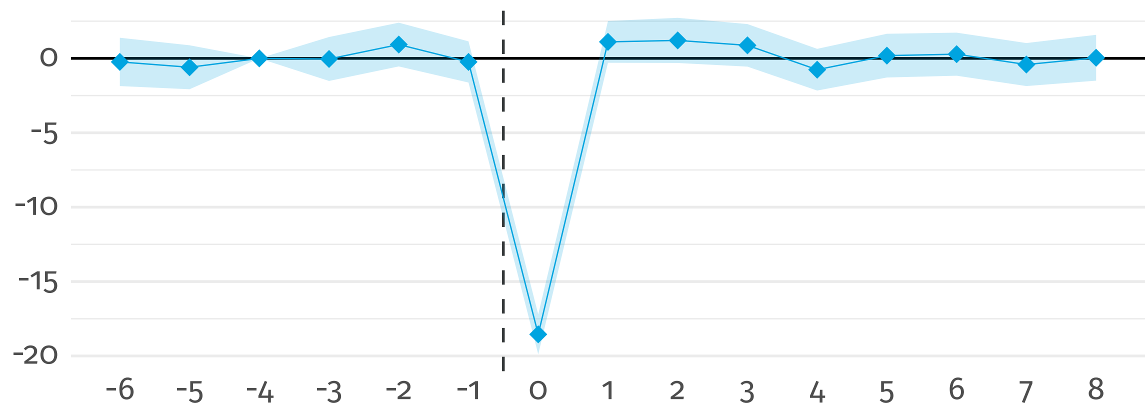

Appendix Figure A11: Effect on Total Costs in Spring Window

Note: The aggregated coefficient (average treatment effect of regular GP absence on total healthcare expenditures in a week) is CHF -18.5. The baseline costs of the outcome in weeks 3 to 8 among the treated population are CHF 94.7. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

A.2.3 Fall Window (Weeks 35 to 51)

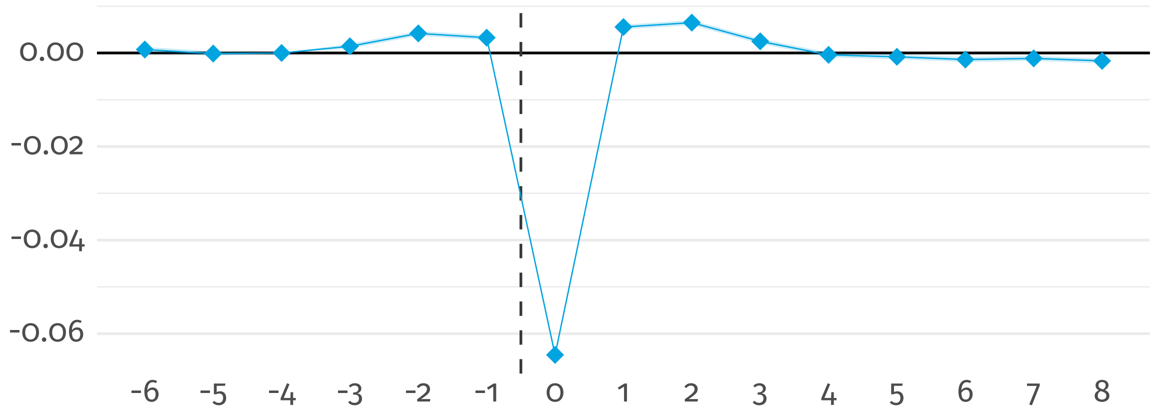

Appendix Figure A12: Effect on any Primary Care Physician Visit in Fall Window

Note: Results from the second stage regression with 30.7 million observations and with 52.6 million control observations in the first stage regression. The aggregated coefficient (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) is -0.065. The baseline probability of the outcome in weeks 35 to 39 among the treated population is 8.06%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

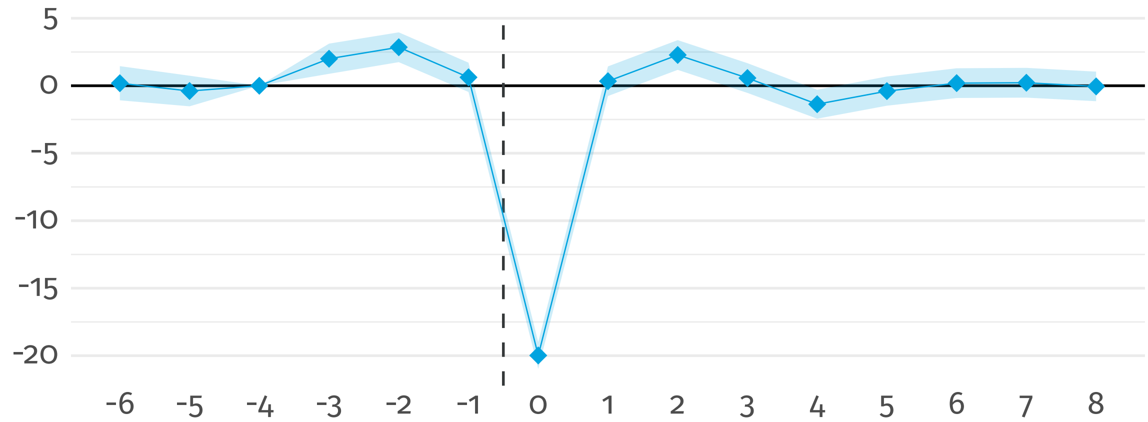

Appendix Figure A13: Effect on Total Costs in Fall Window

Note: The aggregated coefficient (average treatment effect of regular GP absence on total healthcare expenditures in a week) is CHF -20.0. The baseline costs of the outcome in weeks 35 to 39 among the treated population are CHF 96.6. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

A.3 Heterogeneity

A.3.1 GP Absence Length

A.3.2 Number of Locally Available GPs

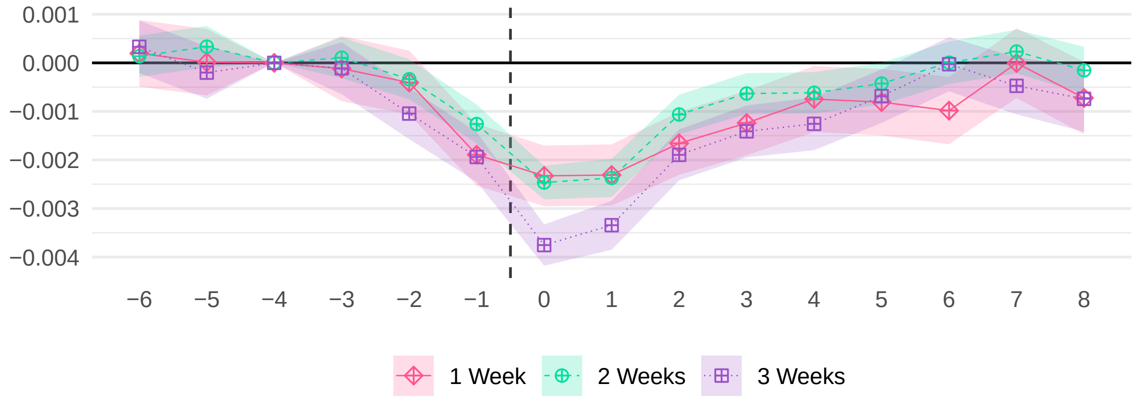

Appendix Figure A14: Effect on any Specialist Physician Visit by GP Absence Length

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any specialist physician visit) are -0.002, -0.002, -0.004. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 2.95%, 2.96%, and 3.07%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

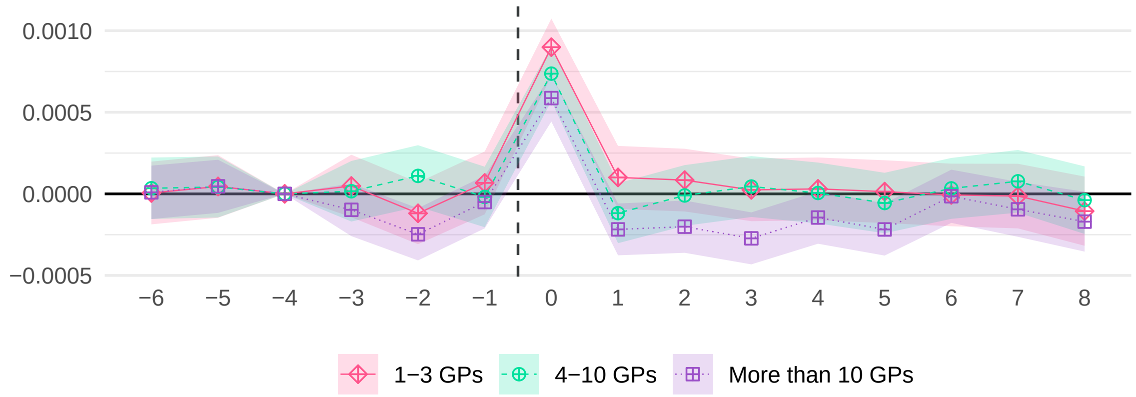

Appendix Figure A15: Effect on Any Emergency Department Visit by Number of Local GPs

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any ED visit in a week) are 0.0009, 0.0007, and 0.0006. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 0.28%, 0.31%, and 0.38%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

A.3.3 Insurance Model

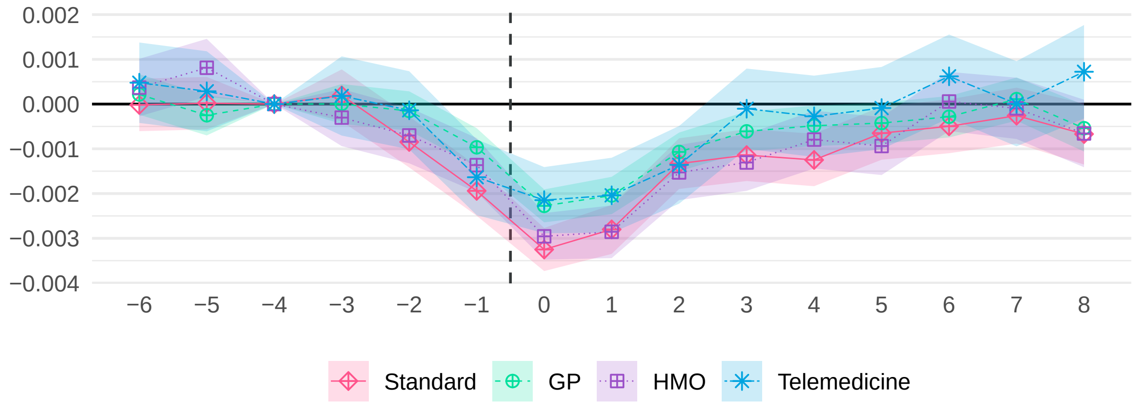

Appendix Figure A16: Effect on any Specialist Physician Visit by Insurance Model

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any specialist physician visit in a week) are -0.0033, -0.0023, -0.0030, and -0.0021. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 3.88%, 2.63%, 2.53%, and 2.37%. The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.

A.3.4 Co-Pay Status

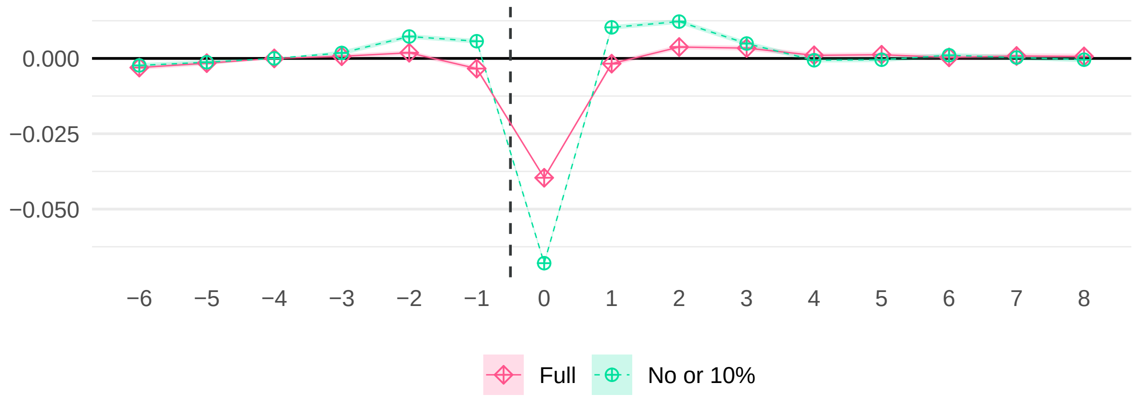

Appendix Figure A17: Effect on any Primary Care Visit by Co-Payment Status in Week 19

Note: The aggregated coefficients (average treatment effect of regular GP absence on the likelihood of any primary care visit in a week) are -0.040 and -0.068. The baseline probabilities of the outcome in weeks 20 to 25 among the treated population are 3.90% and 10.9% The shaded area surrounding the coefficients (shown as diamonds) represents the 99% confidence interval of the estimate, with standard errors clustered at the individual-year level.