WP 26/11

Primary care connectedness and Healthcare spending in Switzerland

Igor Francetic

August 2026

http://www.york.ac.uk/economics/postgrad/herc/hedg/wps/

Primary care connectedness and healthcare spending in Switzerland

Igor Francetic*

August 10, 2026

Abstract

Objectives: In healthcare systems with high provider fragmentation, integrated care could curb spending and reduce inefficiency. This paper evaluates whether primary care network topology and the connectedness of providers affect patient-level healthcare expenditures and utilization across six spending categories.

Methods: Using Swiss mandatory health insurance claims from 2015 to 2024 structured into three triennia (2015–2017, 2017–2019, 2022–2024), I construct patient-sharing networks among primary healthcare providers over baseline two-year windows. Communities are detected using a spatially weighted Leiden algorithm under the Constant Potts Model objective function. Patients are assigned to communities via their main primary care provider. Imposing community stability in the first two years of the triennium, patients switching communities in year three are identified as network movers. To identify supply-side network effects, I implement a within-movers ANCOVA design evaluating the change in community degree centralisation, the change in (log) degree centrality of the main primary care provider, and the change in local clustering, across both extensive (propensity to incur spending) and intensive (log expenditure) margins.

Results: Within-mover estimates point in the same direction at both scales: greater connectedness is associated with lower utilisation and spending. For communities comprised of 3 or more providers, a one-standard-deviation increase in community degree centralisation lowers the probability of incurring expenditure in every category examined, by between 0.5 percentage points (outpatient care) and 3.6 percentage points (emergency care), and reduces conditional spending on laboratory tests by 14.3% and on prescription drugs by 9.7%. Switching to a main primary care provider with higher degree centrality (i.e. one who collaborates more with other providers) likewise reduces spending across all expenditure classes, with the largest drops observed for laboratory tests (-1.1 p.p. extensive, -5.0% intensive per doubling of the provider's degree). Higher local clustering around the main primary care provider reduces the propensity to incur expenses on prescription drugs and emergency care and reduces conditional drug spending, but leaves the remaining categories unaffected.

Discussion: Connectedness in primary care, whether measured by the structural position of a patient's own provider or by the coordination architecture of the community that provider belongs to, is associated with lower spending, concentrated in the diagnostic and pharmaceutical categories where duplication is easiest to avoid. Whether network centralisation is a valid cost-containment instrument should be confirmed by further accompanying evidence on quality, for example using health outcomes to distinguish lower spending from under-provision.

Keywords: Switzerland; healthcare; collaboration; connectedness; integrated care.

JEL codes: I11, I18, D85.

Acknowledgments: I am grateful to Ludovica Alesci and Sam Khavandi for excellent research assistance, and to the data provider for their availability and flexibility. The work benefitted from the feedback of participants to the 2025 workshop on "Health Economics research using administrative data in Switzerland" (HEADS) in Lugano and the 2026 Swiss Health Economics Workshop (SHEW) in Lucerne, with special thanks to Prof. Simon Wieser and Dr. Ana Cecilia Quiroga Gutiérrez for their rich feedback. Generative AI was used in this work to create `tikz` figures, to spellcheck text, and to clean and review blocks of software code. This work was supported by the Swiss National Science Foundation, through its Ambizione programme (grant nr. 216326).

Data and code: This work builds on proprietary insurance claims data. Software code is available from the author.

*Igor Francetic: University of Applied Sciences of Southern Switzerland (SUPSI), Department of Business Economics, Health and Social Care (DEASS), Via Violino 11, 6928 Manno (Switzerland) and University of Manchester, School of Health Sciences, Manchester M13 9PL (UK), Email: igor.francetic@supsi.ch.

1 Introduction

Care integration is increasingly seen as a necessity to provide adequate care to multi-morbid patients, and concurrently as an efficiency-increasing policy that helps to curb the growth of healthcare costs (Damery et al., 2016; Struckmann et al., 2018; Zonneveld et al., 2018; Stokes et al., 2021; Zhang et al., 2025; Baxter et al., 2018; Wulsin et al., 2006). However, the evidence on the effects of integrated care on healthcare costs is far from conclusive (Rocks et al., 2020; Goldzahl et al., 2022; Peikes et al., 2009; Ehlert and Oberschachtsiek, 2014). Two challenges in this literature have been the lack of an agreed-upon definition of care integration and recognised measurement tools thereof (Bautista et al., 2016; Suter et al., 2017). Despite the lack of a coherent definition, the literature seems to converge in identifying collaboration between providers as a crucial dimension of care integration (Zonneveld et al., 2018; Wulsin et al., 2006).

There are various reasons why collaboration in healthcare delivery is crucial (Kee et al., 2023). Firstly, healthcare is a multidisciplinary endeavour (Yeager, 2005; Matulis et al., 2017): taking care of patients, let alone those with complex needs, requires a team of various healthcare professionals. Secondly, collaboration smooths the flow of information, which remains crucial even in the era of freely exchangeable electronic health records (Morley and Cashell, 2017). Collaboration also fosters improved communication between providers, for example between a family doctor and relevant specialists in primary care settings (Foy et al., 2010; Katon et al., 2002). The appeal of strengthening collaboration among healthcare providers appears particularly strong in primary care, given the varied nature of the conditions treated in this setting (Valentijn et al., 2013). Despite this, evidence on the effects of (multidisciplinary) collaboration in primary care remains inconclusive (Saint-Pierre et al., 2018; Bouton et al., 2023). One specific manifestation of collaboration across primary care providers is patient sharing (Kesternich and Rank, 2022; Shin, 2023).

In many healthcare systems, primary care physicians (PCPs) act (formally or informally) as gatekeepers and coordinators of specialised care. This is particularly relevant in Switzerland, where over 90 percent of the adult population relies on a single main primary care provider (Merçay, 2016). Prior literature demonstrates that provider practice styles significantly shape patient care pathways (Bischof and Kaiser, 2021; Hjalmarsson et al., 2023; Albertini et al., 2025; Badinski et al., 2023; de Graaf-Ruizendaal et al., 2018). Consequently, patients experience the consequences of provider collaboration directly through their interactions with their main PCP. However, whether and how provider collaboration affects healthcare expenditure across different network topologies and connectedness with the broader healthcare market remains an open empirical question.

This paper addresses these gaps by evaluating how primary care network topology and provider connectedness affect patient-level healthcare utilization and expenditure. Using administrative Swiss health insurance claims from 2015 to 2024 covering three triennia, I construct patient-sharing networks and detect natural primary care provider communities using a spatially weighted Leiden algorithm under the Constant Potts Model objective. I operationalise network topology and connectedness at three distinct levels: macro-level community degree centralisation, micro-level main PCP degree centrality, and meso-level local provider clustering. To identify supply-side network effects, I implement an ANCOVA within-movers design that evaluates patients switching provider communities or main PCPs on the extensive (outpatient care, inpatient care, laboratory tests, prescription drugs, physiotherapy, and emergency care) and intensive (outpatient care, laboratory tests, prescription drugs and physiotherapy) margins of healthcare utilisation and spending.

My findings reveal that greater connectedness reduces healthcare utilisation, and that this holds at both scales of analysis. At the micro level, switching to a main PCP with higher degree centrality consistently reduces utilisation and spending across all categories, with the largest reductions observed for laboratory testing (-1.1 percentage points on the extensive margin, -5.0% on the intensive margin per doubling of the provider's degree). At the macro level, moving to a one standard deviation more centralised community (i.e. hub-and-spoke) lowers the propensity to consume care across all six categories by between 0.5 and 3.6 percentage points, and reduces conditional spending on laboratory tests (by 14.3%) and prescription drugs

(by 9.7%). Meso-level local clustering around the main primary care provider (i.e. a tight-knit collaborative environment in the immediate neighbourhood of the gatekeeper) reduces the likelihood of incurring spending on drugs and emergency care and lowers conditional drug spending, but leaves the other categories unaffected. These results relate to communities with three or more providers (where centralisation is defined), and not to moves from very small to larger communities of provision.

The remainder of the paper is structured as follows: Section 2 outlines the institutional background of the Swiss healthcare system. Section 3 describes the data, network construction, and topological metrics. Section 4 presents the empirical strategy and mover design. Section 5 details the main results across extensive and intensive margins, and Section 6 concludes with a discussion of policy implications.

2 Brief institutional background

In Switzerland, around 90% of the resident population are estimated to have a consistent primary care provider, often a private general practitioner (GP) (Merçay, 2016). In Switzerland, GPs are usually independent and operate individually or in group practices (Bischof and Kaiser, 2021). Patients need to register with a GP and are responsible for changing their GP if they wish. Depending on their insurance plan, individuals can either directly see any doctor, including specialists (if they chose a more expensive standard plan with free choice of provider), or they might be required to first consult their main primary care provider or another gatekeeper (in the case of managed care plans like family doctor plans, HMOs, and telemedicine).

While basic health insurance covers the expenses for primary healthcare, patients are required to pay out-of-pocket up to a deductible (ranging between CHF 300 and CHF 2'500, or about 320 to 2'700 Euros as of 2026) which they also choose alongside the insurance plan (Schmid et al., 2018). Plans with a high deductible (i.e. lower coverage) are cheaper than plans with low deductible (i.e. high coverage); the premium difference between the cheapest and the most expensive plans can be up to 30-40%. After crossing the deductible, everyone is subject to a 10% co-payment for healthcare paid through mandatory health insurance.

Providers are reimbursed according to a number of different tariff system, depending on setting and type of service. For care provided in inpatient setting, hospitals are reimbursed according to Swiss DRG, a prospective payment system introduced in 2012 and originally based on the diagnosis related groups system introduced in Germany (G-DRG). For care activity in outpatient settings, between 2004 and 2025, providers were reimbursed according to a retrospective fee-for-service tariff system known as TARMED [1].

3 Data and Methods

3.1 Data

My study is based on data from a large health insurance company in Switzerland, covering 15-20% of the Swiss population depending on year and canton. Whilst reflecting a non-random sample of the population, it is broadly representative of the general population over time. Data collected includes billing data for all healthcare that is paid through the mandatory health insurance but not care that is paid privately (i.e. with private funds or with private health insurance coverage), through voluntary insurance, accident insurance, disability insurance, and military insurance. It includes detailed bill-level expenditure data, basic patients' sociodemographic characteristics (age, gender, an in-house developed income index), insurance contract data (insurance company, model, deductible, premium, receipt of subsidy), as well as basic healthcare provider and referral information (i.e. the provider who referred the patient to another provider, when available).

[1] More here: https://eurohealthobservatory.who.int/monitors/health-systems-monitor/updates/hspm/switzerland-2015/revision-of-the-reimbursement-scheme-of-ambulatory-care-(tarmed). TARMED has been updated and renamed TARDOC since January 2026.

These patient-level records also include the patient's geographic area of residence in Switzerland at the level of MedStat regions. MedStat regions are 706 postal code-based geographic zones in Switzerland. Each region includes large enough population to ensure the anonymous reporting of hospital patient residences[2]. To avoid identification, provider locations are available at the same aggregate geographical level. However, the data provided include a matrix reporting the geodesic distance (in kilometers) between each pair of providers.

3.2 Data preparation and sample selection

I construct my analytical panel by merging administrative health insurance contract records, monthly healthcare expenditures, inpatient claims, emergency department use flags, and outpatient care bills across three triennia (2015–2017, 2017–2019, and 2022–2024). After applying sample selection and eligibility criteria I am left with about 672 thousand unique patients observed across the three triennia.

Sample Selection and Eligibility Rules To ensure a consistent sample of insured individuals, I apply specific demographic and contract filters. First, I restrict the sample to adult patients aged 19 and older who maintain full 12-month calendar-year insurance coverage. Second, I exclude individuals who died at any point during the observation window to prevent end-of-life health shocks from distorting healthcare expenditure patterns. Third, to eliminate contract-switching noise, I drop patient-years involving mid-year deductible changes or non-standard deductible options. Finally, I exclude obstetric/maternity-related care episodes.

Inpatient claims and clinical severity I use details on inpatient care bills to calculate annual hospital length of stay (LOS, in days) per patient. For patients with an inpatient stay, I link each inpatient stay to the corresponding official annual SwissDRG tariff catalogue (2015-2024) to extract clinical severity levels derived directly from the fourth digit of the SwissDRG code.

Expenditure categories and emergency utilisation I aggregate monthly expenditure claims into five annual spending categories: outpatient care, inpatient care, pharmaceuticals, laboratory tests, physiotherapy. For these five spending categories, I retain the monetary amount (in CHF). Healthcare expenditure features a specific structure that requires distinct margins modelling. For the extensive margin ($\Pr(Y > 0)$), I consider six utilization domains: outpatient visits, inpatient admissions, prescription drugs, laboratory tests, physiotherapy, and emergency department visits. Inpatient care is evaluated strictly on the extensive margin, as acute hospitalizations are low-frequency, high-cost shocks largely outside primary care gatekeeping influence. For the intensive margin, I consider only high frequency outpatient domains (outpatient medical, drugs, lab, and physiotherapy spending).

3.3 Characterising communities of primary care provision

Characterising what in this paper is referred to as primary care "collaboration and connectedness" is one of the innovations of this work. The approach proposed here builds on tools developed in the field of network theory and analysis (Jackson et al., 2017; Burns et al., 2022). This subsection defines the key terms and outlines the process followed to obtain my measures of connectedness. All algorithms and network measurements described in this section have been conducted using the statistical software R, version 4.6.

[2]See more here: https://www.versorgungsatlas.ch/en/p/regions

Key terminology: To relate my setting to common terminology in the network literature, I start by defining key terms. A network is characterised by two key entities: nodes and edges, which are also known as vertices and ties (Burns et al., 2022). In the context of this paper, a **node** i is a healthcare provider. In a given network, there are H providers. An **edge** e_{ij} linking two providers i and j is represented by a patient p shared by the two providers. Here I do not consider the directionality of edges (i.e. distinguishing a referral from provider i to provider j and one from j to i); a patient-sharing relationship generates a bidirectional connection between two providers $i \neq j$. A **network** $\mathcal{N}(\tau)$ is the collection of nodes and edges for the H providers observed across the sample in period τ. I define each period as a triennium $\tau \in (2015-2017, 2017-2019, 2022-2024)$, with years in a triennium $t_1, t_2, t_3 \in \tau$ (for example, in triennium 2015-2017, $t_1 = 2015$). So, in my setting, $\mathcal{N}(\tau)$ represents the collection of providers and their shared patients in Switzerland over a given triennium τ. Practically, $\mathcal{N}(\tau)$ can be represented by a symmetric matrix with rows/columns representing providers, and cells representing the number of shared patients between providers i and j. Finally, let me define $\text{dist}_{i,j}$ as the distance (in km) between providers i and j. This information is useful to account for the influence of geography in determining patients sharing relationships.

Defining nodes, i.e. healthcare providers: The paper focuses demand responses to provider connectedness and integration. Whilst care integration can also be conceived in terms of vertical integration (e.g. primary care and secondary/tertiary care), my focus here is on horizontal integration in relation to care delivered in ambulatory outpatient settings. To achieve this, I consider only patient-sharing relationships amongst physicians (i.e. generalists and specialists) in ambulatory outpatient settings (i.e. not specialised care delivered in hospitals outpatient settings) billed according to TARMED. Only providers falling within these categories are considered as nodes, so that my full networks represent patient-sharing relationships between GPs/family doctors and specialists in ambulatory settings. This set of criteria aim to align with the goal of characterising horizontal coordination and collaboration at the primary care level, and not vertical integration (e.g. between primary and secondary/tertiary care levels).

Defining edges, i.e. shared patients: I define an edge $e_{i,j}$ as a patient receiving care from both providers i and j in the same period (and hence linking providers i and j by the fact that they are both treating the same patient. Whilst capturing explicit patient referrals $i \leftrightarrow j$ is a more direct and unambiguous way of identifying active collaboration between providers, missing data on referrals poses additional empirical challenges. To strengthen the definition of edges and avoid impromptu collaborations, I only consider edges between providers that share a patient for at least 3 distinct insurance claims in a given period. This approach follows the literature focused on identifying networks of healthcare providers from administrative and insurance data (Landon et al., 2013; Hietapakka et al., 2025). Sensitivity analyses where I increase the threshold for the minimum number of claims required to define a shared patient, and where I only consider explicit referrals between providers as edges, are included in the appendix C. An edge $e_{i,j}$ takes value 0 (representing providers who share zero or less than 3 patients in a given period) or the positive count of patients shared by providers i and j.

Detection of communities: This study uses data spanning 10 years (2015-2024) linking providers through episodes of care (i.e. insurance claims) for patients insured through the data provider. For a given observed network in period τ, if care was sought completely at random from providers, one would observe networks with no obvious clustering structure. However, most real-world networks feature patterns of clustering. Reasons for clustering among healthcare providers in the delivery of care in the Swiss setting include:

- Spatial proximity (i.e. lower transportation and opportunity cost driving collaboration) between providers, especially in rural or sparsely inhabited regions, or more broadly patients' preferences
- Personal social connections leading to collaboration between providers (e.g. a family doctor and a specialist who shared their foundation training in medicine)
- Insurance-related integration through managed care plans

The first step to measure provider connectedness is to identify clusters (henceforth referred to as *communities*) of providers that work tightly together based on their patterns of patient sharing. The goal is to define which community $c^\tau \in C^\tau$ a provider i belongs to in a given period τ. To reflect medium/long term equilibrium provider behaviour (as opposed to short term fluctuations), I use data across three-year periods (henceforth *triennia*) τ to identify communities of tight primary care collaboration. The dynamic evolution of care communities over time is interesting in its own right (e.g. *when a core player of a community exits the market, how does the community reconfigure?*). Nevertheless, this paper is focused on the responses to more or less connected and integrated communities of care delivery.

To net out the influence of spatial proximity in identifying communities, I use weighted edges $w_{i,j}$ which incorporate dynamic exponential spatial decay as follows:

$$w_{i,j} = \mathrm{c}_{i,j} \times \exp(-\lambda_\tau \cdot \mathrm{dist}_{i,j}), \tag{1}$$

where $\lambda_\tau = \ln(2)/\mathrm{dist}_{\mathrm{p50},\tau}$ is dynamically calibrated per triennium such that collaboration ties at the patient-weighted median distance ($\mathrm{dist}_{\mathrm{p50}}$) retain 50% of their weight.

Algorithm definition and calibration: To identify distinct clusters of primary care providers who frequently co-manage patients, I apply the Leiden algorithm Traag et al. (2011, 2019) using the Constant Potts Model (CPM) objective function. Formally, the algorithm partitions providers in the spatially-weighted patient-sharing network (Wang et al., 2021, 2022) into a set of non-overlapping communities $\mathcal{C}$ by maximising the objective function Q:

$$Q = \sum_{c \in \mathcal{C}} \left[\sum_{\substack{i,j \in c \\ i<j}} w_{i,j} \; - \; \gamma_\tau \binom{H_c}{2} \right], \tag{2}$$

where $w_{i,j}$ is the spatially decayed patient-sharing weight between providers i and j in community c, H_c represents the total number of providers in community c, and γ is the resolution parameter acting as a density penalty. Intuitively, two providers are assigned to the same community if their internal patient-sharing tie weight exceeds the expected density threshold defined by γ_τ. Each provider i in period τ belongs to one of these communities. Communities are not restricted to lie within cantons, but the spatial decay implies that most communities are naturally self-contained within cantons.

The resolution parameter γ_τ is calibrated separately for each triennium $\tau \in T$ by scanning a grid of resolution parameters and executing the Leiden algorithm across $N = 5$ random seeds per candidate value under the Constant Potts Model (CPM) objective. Further details are provided in appendix A, where I also report some plausibility checks for the algorithm's specificity. For example, in the Swiss institutional setting, cantons represent the natural macro-level unit for healthcare regulation, hospital planning, and provider collaboration. Consequently, evaluating community partitions across varying resolution scales confirms that the network construction accurately captures underlying spatial and institutional patient flows without introducing algorithmic randomness.

Running the algorithm I obtain, for each triennium, a partition of providers into $\mathcal{C}$ communities. Zooming from $\mathcal{N}(\tau)$ into a specific community $c \in \mathcal{C}$, one can also derive an induced subgraphs $G(c, \tau)$, i.e. treat the community (with its providers and patient-sharing relationships) as if it was a full network. This will be helpful below when defining integration, connectedness and collaboration metric the different communities. The number of communities identified in each canton over the triennia is reported in Table 1; there is relatively limited variation between triennia, which is indicative of algorithm stability.

Table 1: Number of communities identified across cantons over time

Canton	1517	1719	2224
AG	102	81	131
AI	15	13	18
AR	20	17	23
BE	150	115	170
BL	57	59	83
BS	45	43	71
FR	87	68	120
GE	94	77	129
GL	23	16	23
GR	56	51	70
JU	20	21	38
LU	63	53	81
NE	33	33	42
NW	24	22	37
OW	26	27	33
SG	66	57	79
SH	25	23	30
SO	71	64	90
SZ	59	51	75
TG	37	35	49
TI	64	49	71
UR	17	17	25
VD	103	88	162
VS	76	57	85
ZG	47	38	61
ZH	178	126	198

Main primary healthcare provider: Contrary to settings like the English NHS (where residents are required to register with a local General Practice), defining the main provider of primary healthcare for a given patient in Switzerland is not trivial. The main primary healthcare provider is typically a GP, a specialist, or a group practice. Patients can also switch between one main primary healthcare provider and another voluntarily or in response to exits of their previous main primary care provider. These latter switches have been shown to trigger changes in healthcare patterns in Switzerland (Hjalmarsson et al., 2023; Bischof and Kaiser, 2021). Absent an agreed-upon set of rules, I currently follow an algorithm that blends frequency of interactions and provider's "type" [3]. In short, for each patient p in year t, I filter outpatient (TARMED) claims to those billed by primary care provider (PCP) specialties, including general internal medicine, general practitioners, paediatricians, and group practices [4]. I then count claims registered by each patient with each provider and assign the main primary healthcare provider i for patient p in year t as p's healthcare provider that accounts for at least 50% of the patient's annual primary care claims. I impute gaps in this information at the level of patients (e.g. if a patient had no TARMED claims in one year, making it impossible to define their main primary care provider for that same year, but with information for other years of a same continuous contract stint with the insurer) based on the main primary healthcare provider identified in the closest year. Patients with no recorded TARMED-billed spending are excluded from the analysis.

In my study, identifying the main PCP serves multiple purposes. Conceptually, as the first point of contact with the healthcare system, care and referral decisions taken by the patient's main PCP impact patient's care and spending downstream. Empirically, I define a patient's membership to a specific community of primary care delivery based on the membership of their main PCP.

[3] This algorithm may be refined depending on the availability of the data provider. Nevertheless, similar definitions have been used in previous work based on the same data, e.g. Bischof and Kaiser (2021); Hjalmarsson et al. (2023)

[4] That is one of the following categories: 0100: Ärzte und Ärztinnen, 0105: Innere Medizin, 0116: Kinder- und Jugendmedizin, 0121: Allgemeine innere Medizin, 0153: Praktischer Arzt/Ärztin, 0175: Gruppenpraxen

3.4 Measuring care coordination, connectedness and collaboration

To capture network structure at both the community level and the "micro" provider level, I use three complementary metrics related to the provider patient-sharing networks described above. Each metric isolates a distinct network-structural mechanism influencing care coordination, provider collaboration, and healthcare utilisation. Each metric is computed separately for each triennium.

Community degree centralisation: To discriminate between different coordination topologies of patient-sharing communities, I calculate Freeman's degree centralisation on the community subgraph $G(c, \tau)$:

$$C_D(c) = \frac{\sum_{j \in H_c} (d_c^* - d_{j,c})}{(H_c - 1)(H_c - 2)}, \tag{3}$$

where H_c denotes the total number of providers belonging to community c, $d_{j,c}$ is the within-community degree of provider j (the count of non-spatially-weighted edges linking provider j to peers within community c), and d_c^* is the maximum within-community degree observed in c. The metric is bounded s.t. $C_D(c) \in [0, 1]$.

The denominator $(H_c - 1)(H_c - 2)$ in (3) is zero for communities of fewer than three providers; this is not a marginal case in my partition, as reported in appendix A.4. I therefore define $C_D(c)$ only for communities with $H_c \geq 3$ and restrict the community-level analysis to patients whose origin and destination communities both satisfy this condition. Because singleton and paired communities are by construction the ones almost no patients are assigned to, those retained cover 98.9% to 99.3% of patient-community assignments. The restriction is close to costless in terms of patients, and it removes a large tied mass at zero that would otherwise dominate any discretisation of the measure. The implications of moves from solo to networked communities (or vice versa), a different margin of connectedness, are explored in Appendix C .

Conceptual Interpretation: From an organisational and health policy perspective, degree centralisation reflects the structural hierarchy and coordination architecture of a primary care community. A high centralisation score with $C_D(c) \to 1$ indicates a structured, "hub-and-spoke" care delivery model where patient-sharing pathways are organised around a few dominant providers. This hierarchical architecture facilitates centralised gatekeeping, standardised clinical pathways, and clear referral channels, but creates a high concentration of structural influence and reliance on a small number of key central providers. Conversely, a low centralisation score with $C_D(c) \to 0$ represents a decentralised, horizontal network where patient-sharing relationships are evenly distributed across providers rather than anchored around a central coordinating hub. Such flat architectures allow clinical information and informal best practices to diffuse peer-to-peer through non-mediated professional interactions. Figure 1 shows two exemplars of extreme configurations of patient sharing networks.

Provider degree centrality: Similar to Shin (2023), to capture overall provider prominence and collaborative reach across the broader healthcare system, I measure the total degree of a patient's main primary care provider:

$$d(i) = \sum_{j \neq i} e_{ij}, \tag{4}$$

where e_{ij} is the raw (non-spatially-weighted) count of shared patients between a provider i and other providers $j \neq i$. In empirical models, to address right-skewness, this metric enters in log-transformed form:

$$\text{ldeg(i)} = \ln(d(i)), \tag{5}$$

representing the log degree centrality provider i in year t.

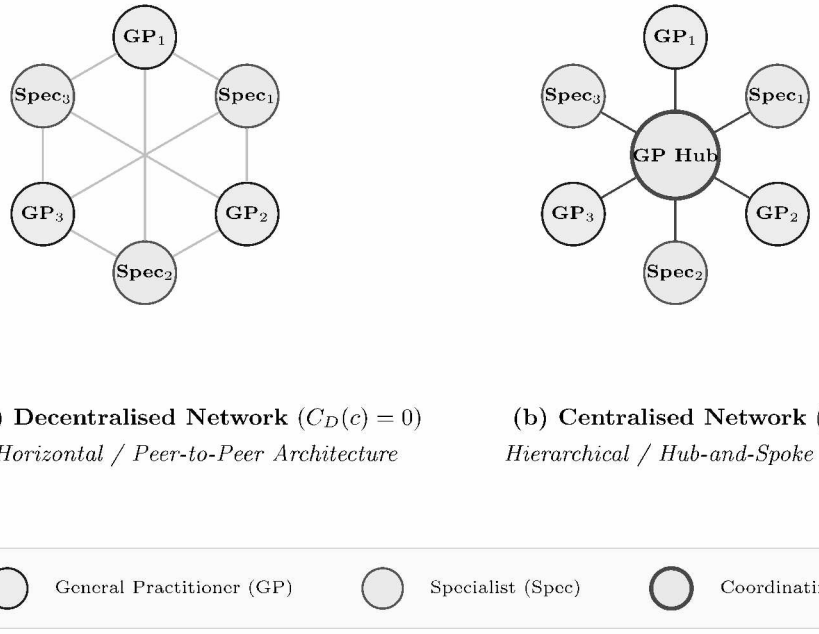

(a) Decentralised Network ($C_D(c) = 0$)
Horizontal / Peer-to-Peer Architecture

(b) Centralised Network ($C_D(c) = 1$)
Hierarchical / Hub-and-Spoke Architecture

Note: Panel (a) depicts a decentralised 3-regular network with an equal distribution of patient-sharing ties across primary care physicians (GP) and specialists (Spec), yielding $C_D(c) = 0$. Panel (b) illustrates a centralised star network where a primary care hub mediates all patient-sharing flows between peers and specialists, yielding maximum centralisation ($C_D(c) = 1$).

Figure 1: Extreme topologies of community degree centralisation $C_D(c)$

Conceptual Interpretation: Provider degree centrality measures a PCP's systemic prominence and connectedness (in terms of patient-sharing relationships) across the entire healthcare system (Jackson et al., 2017). High-degree PCPs are very prominent and play a systemic role as integrators, maintaining multiple referral links across various medical specialties. Practically, this broad connectedness could provide patients with easier access to specialised expertise and diagnostic resources. However, while high connectivity reflects clinical prominence, extreme degree centrality may also impose cognitive or management constraints on a provider.

Local clustering coefficient: Analogously to Shin (2023), to assess micro-level density around a patient's primary care provider (PCP), I calculate Watts and Strogatz's (Watts and Strogatz, 1998) local clustering coefficient C_i for each provider i:

$$C(i) = \frac{2\,k_i}{d(i)(d(i)-1)}, \tag{6}$$

where $d(i)$ is the overall degree of provider i in a network $\mathcal{N}(\tau)$, and $k_i = |\{(j,k) \in K : j,k \in N(i)\}|$ represents the number of realised patient-sharing links among the immediate neighbourhood $N(i)$ of provider i. For isolated nodes, $C(i)$ is set to zero. Importantly, since $C(i)$ is inversely proportional to $d(i)$, analyses where (6) is used as regressor should be adjusted for $d(i)$.

Conceptual Interpretation: Local clustering captures triadic closure and informal peer density around a PCP. High clustering means that a provider's collaborators are also actively co-managing patients with one another, forming a tightly knit clinical "clique." This structural density fosters strong social capital, informal peer monitoring, mutual trust, and alignment of practice styles among co-referring clinicians. Such dense peer environments facilitate swift communication during complex patient management and possibly reinforce shared norms against low-value or redundant interventions. Figure 2 represent configurations with extreme opposite local clustering values.

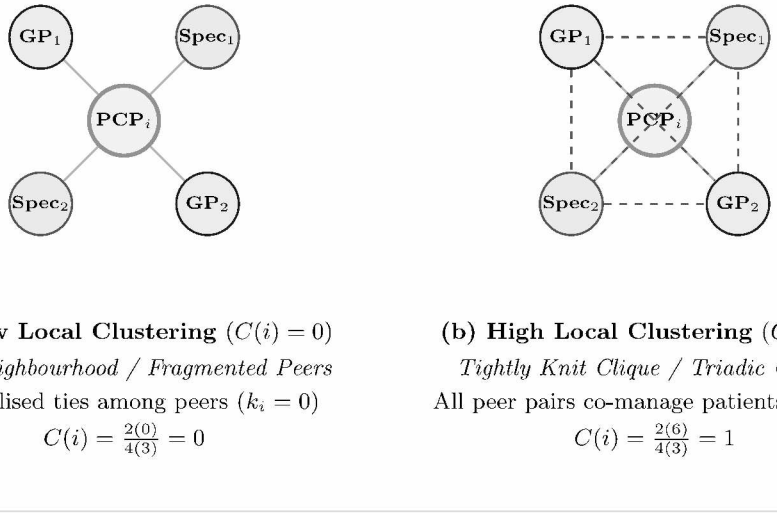

(a) Low Local Clustering ($C(i) = 0$)
Star Neighbourhood / Fragmented Peers
No realised ties among peers ($k_i = 0$)
$C(i) = \frac{2(0)}{4(3)} = 0$

(b) High Local Clustering ($C(i) = 1$)
Tightly Knit Clique / Triadic Closure
All peer pairs co-manage patients ($k_i = 6$)
$C(i) = \frac{2(6)}{4(3)} = 1$

Note: Both panels evaluate provider i with degree $d(i) = 4$. Panel (a) illustrates an isolated star neighbourhood ($C(i) = 0$), where none of provider i's co-referring peers co-manage patients with one another ($k_i = 0$). Panel (b) depicts maximum triadic closure ($C(i) = 1$), where all neighbouring clinicians form active pairwise patient-sharing links ($k_i = 6$), constructing a dense clinical clique.

Figure 2: Extreme local clustering $C(i)$ configurations around a PCP

4 Empirical Strategy

4.1 Baseline specification

My primary empirical objective is to evaluate whether the connectivity structure of the primary care network around a patient's main PCP, and its connectedness position in the network, influence the patient's healthcare utilisation (extensive margin) and spending (intensive margin) patterns. To isolate how network topology influences healthcare outcomes at both the macro (community) and micro (provider) levels, my three baseline specifications read as follows.

Role of community topology To test whether the community topology/architecture drives utilisation and spending, a naive model to capture the effect of community-level degree centralisation $C_D(c)$ for patient p's assigned network community c would read as follows:

$$Y_{p,c,t} = \alpha + \beta_1 C_D(c) + \mathbf{X}'_{p,t}\mathbf{\Phi} + \delta_{k,t} + \varepsilon_{p,c,t}, \tag{7}$$

where $Y_{p,c,t}$ represents annual healthcare outcomes for patient p assigned (through her main PCP i) to community c and in a given year t of triennium. The vector $\mathbf{X}_{p,t}$ captures time-varying patient clinical severity (Swiss DRG severity, prior hospital length of stay, age and sex), insurance contract parameters (deductible level, type of managed care model), while $\delta_{k,t}$ captures canton-by-year fixed effects.

Role of provider connectedness While $C_D(c)$ captures overall community topology, it does not reveal whether patient outcomes are driven by the individual connectedness of their specific primary care physician. Shift from community-level topology to individual provider scale, a naive model to estimate the effect of p's main PCP's i log degree $\text{ldeg(i)}_{p,t}$ on utilisation and spending is:

$$Y_{p,c,t} = \alpha + \beta_2 \text{ldeg(i)}_{p,t} + \mathbf{X}'_{p,t}\mathbf{\Phi} + \delta_{k,t} + \varepsilon_{p,c,t}, \tag{8}$$

where $\text{ldeg(i)}_{p,t}$ measures the log number of unique provider-sharing ties maintained by patient p's assigned PCP i in year t. Note that in (8), a patient can change main PCP from year to year. Coefficient β_2 tests whether consulting a highly connected or high-volume referral hub PCP directly alters healthcare utilization.

Role of locally dense collaboration clusters Finally, high degree centrality may simply reflect high patient volume or broad referral reach rather than tight local care coordination. To isolate true local coordination from overall provider scale, a naive model would add the PCP's local clustering coefficient $C(i)_{p,t}$ alongside log degree in a joint specification, as follows:

$$Y_{p,c,t} = \alpha + \beta_2 \text{ldeg(i)}_{p,t} + \beta_3 C(i)_{p,t} + \mathbf{X}'_{p,t}\mathbf{\Phi} + \delta_{k,t} + \varepsilon_{p,c,t}, \tag{9}$$

where $C(i)_{p,t} \in [0,1]$ measures the probability that two providers who share patients with PCP i also share patients with each other (triadic closure). Holding total provider ties constant ($\text{ldeg(i)}_{p,t}$), β_3 isolates the incremental effect of living in a locally tight provider neighbourhood where specialists and generalists intensely collaborate among themselves.

4.2 Identification strategy

The main challenge for causal identifications in the baseline specifications outlined above is represented by endogenous sorting. On the one hand, healthier or more health-conscious patients may systematically choose primary care providers (PCPs) who belong to tightly coordinated, centralised networks (or the opposite). On the other hand, patients may switch provider or community of providers strategically, anticipating a specific healthcare need. I address these threats through a combination of predetermined network construction and a movers design. In the remainder of this section I outline my identification strategy.

Reverse causality in community detection: A primary concern when estimating the effects of micro-level provider metrics is reverse causality driven by patient health status. A patient who experiences an acute illness will see multiple specialists, mechanically creating new patient-sharing ties for their main PCP. This would artificially inflate both the PCP's degree and local clustering coefficient, generating a spurious positive correlation between PCP network density and patient healthcare utilization. To eliminate this source of reverse causality, I calculate all network metrics ($C_D(c)$, $d(i)$, and $C(i)$) using provider-sharing ties established during a predetermined baseline window (years t_1 and t_2 of each triennium τ). By fixing network topology prior to the outcome measurement period t_3, individual health shocks in year t_3 cannot retroactively alter a provider's network metrics.

Endogenous patient sorting: Even with predetermined network metrics, two forms of unobserved patient sorting threaten causal identification:

1. **Macro-level sorting:** Healthier or more health-literate patients may systematically choose to reside in cantons or municipalities served by primary care communities with high or low centralisation ($C_D(c)$).
2. **Micro-level sorting:** Sicker or more demanding patients may intentionally seek out high-volume "hub" GPs with expansive referral networks (i.e. high $d(i)$) or tightly integrated provider cliques (i.e. $C(i) \to 1$).

While standard panel fixed-effects models control for time-invariant individual heterogeneity, healthcare expenditures exhibit strong mean-reversion ($\text{Corr}(Y_{p,t}, Y_{p,t-1}) < 1$) and zero inflation. In these settings, using a pre-post change (first differences) can introduce dynamic bias and reduce statistical efficiency. Following McKenzie (2012), I control for baseline health status using an Analysis of Covariance (ANCOVA) framework, which estimates the baseline outcome correlation freely and maximises power.

Furthermore, to isolate exogenous variation in exposure to integration and connectedness, my primary identification strategy adapts the within-movers design of Finkelstein et al. (2021). Rather than using

non-switchers (stayers) as a counterfactual, which risks confounding switching decisions with unobserved provider stability preferences, the sample is restricted exclusively to movers. I observe patients across triennia ($\tau \in \{2015-2017, 2017-2019, 2022-2024\}$), tracking pre-move health status in year t_2 and post-move outcomes in year t_3. I refer to a patient who changes their main primary care provider ($i_{p,2} \neq i_{p,3}$) between years t_2 and t_3 as a *PCP switcher*. Since patients are assigned to communities c based on the membership of their main PCP i, a switch may or may not carry the patient across community boundaries. In my sample, about 15% of patients are PCP switchers. Table 2 highlights how not all of these changes translate into moves across communities, as patients might move to a different main PCP within a same community. Throughout the analysis a patient is classified as a *mover*, and enters the estimation sample, only if their assigned community differs between t_2 and t_3. This restriction applies to the provider-level specifications as well as to the community-level ones, so that the origin and destination community fixed effects are well defined throughout. The community-level specifications impose one further restriction that the provider-level ones do not: both the origin and the destination community must have at least three providers, since $C_D(c)$ is otherwise undefined. To ensure that moves reflect genuine transitions rather than pre-existing instability, I impose a stability condition: only patients whose dominant community is identical in t_1 and t_2 are retained in the analysis sample. Patients who switched communities in the first two years of a triennium are excluded from the analysis.

	2015-2017	2017-2019	2022-2024	Total
Did not change main PCP i	285,187	308,871	423,031	1,017,089
	(23.82)	(25.79)	(35.33)	(84.94)
Changed main PCP i	49,639	60,150	70,557	180,346
	(4.15)	(5.02)	(5.89)	(15.06)
$\rightarrow$ of which, also changed community c	11,634	12,438	19,092	43,164
	(23.44)	(20.68)	(27.06)	(23.93)
Total	334,826	369,021	493,588	1,197,435

Note: Columns indicate unique patients within each triennium with a switch in main PCP (at t_2 or t_3) and a change in community (at t_3). Triennia 2015-2017 and 2017-2019, and some unique patients are the same across triennia. The overall number of unique patients is of about 672,000.

Table 2: Main PCP switchers and community movers across triennia

To model the effect of moves to communities with different network topology, I use the change in degree centralisation between the origin and destination community, $\Delta C_D(c)_p = C_D(c^d) - C_D(c^o)$, standardised by its sample standard deviation so that coefficients are read per standard deviation of the change. In my main analysis, to increase power, I stack the three triennia and account for triennium fixed effects. Appendix B includes models estimated separately by triennium. My ANCOVA specification uses one observation for individual p in triennium τ, switching from origin PCP community c^o to destination community c^d between t_2 and t_3, observed at t_3.

My causally identified equivalent of (3) reads as follows:

$$Y_{p,3,\tau} = \beta\, \Delta\widetilde{C}_D(c)_p + \psi \widetilde{Y}_{p,2,\tau} + \mathbf{X}'_{p,3}\mathbf{\Phi} + \zeta \ln H_{c^d} + \alpha(c^o) + \delta_{k,3} + \tau + \varepsilon_{p,\tau}, \tag{10}$$

where $\Delta\widetilde{C}_D(c)_p = [C_D(c^d) - C_D(c^o)]/\sigma$, $\widetilde{Y}_{p,2,\tau}$ is the lagged healthcare outcome (not log-transformed in the case of analyses on the intensive margin), and $\alpha(c^o)$ is a fixed effect for the community of origin. By comparing movers exclusively amongst themselves, and within the same origin community, this approach avoids assuming that stayers are a valid counterfactual for movers. Origin-community fixed effects absorb any time-invariant characteristics of the departure community (average integration, provider composition, local care norms) that might otherwise confound the estimated effect. Equation (10) mimics directly the

movers-only approach conditional on origin community used in Finkelstein et al. (2021). $\mathbf{X}_{p,3}$ are individual risk adjusters and insurance contract information defined above, $\ln H_{c^d}$ is log size of destination community, $\delta_{k,3}$ is a canton fixed effect and τ is a triennium fixed effect.

Unlike the provider-level equation below, (10) absorbs origin-community fixed effects only: the treatment is a function of the destination community, so destination fixed effects would absorb it entirely. The specification is therefore more exposed to unobserved destination characteristics, which is why the size of the destination community enters directly. That control is itself not innocuous. Degree centralisation is strongly increasing in community size (appendix A.4), so $\ln H_{c^d}$ absorbs part of the treatment rather than only confounding variation, and results are reported with and without it, as bounds on the effect of moving to a more centralised community. A directional variant of (10), replacing $\Delta\widetilde{C}_D(c)_p$ with indicators for moves up and down the national quintile distribution of $C_D(c)$ relative to lateral movers within the same quintile,

$$Y_{p,3,\tau} = \beta_1(\text{Move Up}_p) + \beta_2(\text{Move Down}_p) + \psi\widetilde{Y}_{p,2,\tau} + \mathbf{X}'_{p,3}\mathbf{\Phi} + \zeta \ln H_{c^d} + \alpha(c^o) + \delta_{k,3} + \tau + \varepsilon_{p,\tau}, \quad (11)$$

is reported in Appendix C as a check on functional form rather than as the primary specification, because binning a continuous measure into quintiles discards the magnitude of each move.

When measuring exposure to continuous provider-level attributes, I model changes in provider degree (Δldeg(i) = ldeg(i)$_{p,3}$ − ldeg(i)$_{p,2}$) and local clustering ($\Delta C(i) = C(i)_{p,3} - C(i)_{p,2}$). Here, the equivalent of (9) reads as follows:

$$Y_{p,3,\tau} = \beta_1\Delta\text{ldeg}(i) + \beta_2\Delta C(i) + \psi\widetilde{Y}_{p,2,\tau} + \mathbf{X}'_{i,y_3}\mathbf{\Phi} + \alpha(c^o) + \alpha(c^d) + \delta_{k,3} + \tau + \varepsilon_{p,\tau}, \quad (12)$$

where $\widetilde{Y}_{p,2,\tau}$ represents the baseline level of the outcome under the origin PCP in year 2. The models absorb origin $\alpha(c^o)$ and destination $\alpha(c^d)$ community fixed effects, canton $\delta_{k,3}$, and triennium fixed effects τ.

Identifying assumptions: The key identifying assumption of Equations (10) and (12) is that conditional on baseline healthcare use $\widetilde{Y}_{p,2,\tau}$, demographic and clinical severity $\mathbf{X}_{p,3}$, and origin and destination fixed effects $\alpha(c^o)$ and $\alpha(c^d)$, the direction or magnitude of a network relocation is uncorrelated with unobserved health shocks at 3.

4.3 Estimation and inference

All models are estimated on a pooled sample that stacks the three triennia, and also separately for each triennium (results in appendix B). In the pooled specification, patient and community identifiers are redefined to be triennium-specific (so that the same individual appearing in two triennia contributes two distinct observations). For community-level analyses, standard errors are clustered three-way: at the origin community, the destination community, and the patient level. For analyses focused on provider connectedness metrics, standard errors are clustered two-way: at the provider and the patient level. All regression analyses were conducted using the statistical software Stata, version 19.

5 Results

5.1 Descriptive statistics

Table 3 reports characteristics of the estimation sample, measured in the pre-move year t_2 and grouped by the direction of the change in community degree centralisation. The community-level estimation sample comprises 37,357 patient-triennia in which the patient's assigned community changes between t_2 and t_3 and $C_D(c)$ is defined at both endpoints; of the 36,968 for which a pre-move year is also observed, 8,901 (24.1%) move to a community in a lower national quintile of $C_D(c)$, 12,331 (33.4%) to a higher quintile, and 15,736 (42.6%) between communities in the same quintile. Movers are on average 46 years old and 69% female, and the three groups are similar in age, sex, deductible choice and prior hospital days.

Table 3: Descriptive statistics by direction of centralisation change

	Direction of centralisation change (valid communities)			
	Mover-down	Mover-lateral	Mover-up	Total
N	8,901 (24.1%)	15,736 (42.6%)	12,331 (33.4%)	36,968 (100.0%)
Triennium				
1517	2,563 (28.8%)	4,146 (26.3%)	3,133 (25.4%)	9,842 (26.6%)
1719	2,874 (32.3%)	3,782 (24.0%)	3,947 (32.0%)	10,603 (28.7%)
2224	3,464 (38.9%)	7,808 (49.6%)	5,251 (42.6%)	16,523 (44.7%)
Age	47.846 (17.843)	45.554 (17.147)	46.274 (17.424)	46.346 (17.432)
Female	0.660 (0.474)	0.696 (0.460)	0.701 (0.458)	0.689 (0.463)
Lowest deductible (300 CHF)	0.563 (0.496)	0.524 (0.499)	0.545 (0.498)	0.541 (0.498)
Highest deductible (2'500 CHF)	0.195 (0.396)	0.237 (0.425)	0.216 (0.412)	0.220 (0.414)
Standard insurance model	0.356 (0.479)	0.313 (0.464)	0.311 (0.463)	0.322 (0.467)
Days in hospital in t-1	0.740 (3.908)	0.569 (3.201)	0.583 (2.986)	0.615 (3.320)
Degree centralisation in origin community	0.763 (0.120)	0.721 (0.115)	0.510 (0.099)	0.661 (0.155)
Degree centralisation in destination community	0.496 (0.121)	0.750 (0.124)	0.788 (0.115)	0.702 (0.168)
Size of origin community	679.724 (482.270)	476.330 (382.729)	181.724 (104.188)	427.034 (398.056)
Size of destination community	164.440 (106.339)	581.335 (430.240)	743.970 (469.494)	535.205 (451.195)
Local clustering around main PCP	0.473 (0.268)	0.517 (0.244)	0.588 (0.163)	0.530 (0.231)
Degree of main PCP	802.717 (1,826.966)	534.760 (1,463.512)	97.819 (90.650)	453.532 (1,338.803)
Any inpatient spending				
0	7,562 (85.0%)	13,612 (86.5%)	10,640 (86.3%)	31,814 (86.1%)
1	1,339 (15.0%)	2,124 (13.5%)	1,691 (13.7%)	5,154 (13.9%)
Any spending on drugs				
0	921 (10.3%)	1,575 (10.0%)	937 (7.6%)	3,433 (9.3%)
1	7,980 (89.7%)	14,161 (90.0%)	11,394 (92.4%)	33,535 (90.7%)
Any spending on lab tests				
0	1,675 (18.8%)	2,952 (18.8%)	2,258 (18.3%)	6,885 (18.6%)
1	7,226 (81.2%)	12,784 (81.2%)	10,073 (81.7%)	30,083 (81.4%)
Any physiotherapy spending				
0	6,742 (75.7%)	12,011 (76.3%)	9,362 (75.9%)	28,115 (76.1%)
1	2,159 (24.3%)	3,725 (23.7%)	2,969 (24.1%)	8,853 (23.9%)
Any emergency care spending				
0	6,303 (70.8%)	11,505 (73.1%)	8,951 (72.6%)	26,759 (72.4%)
1	2,598 (29.2%)	4,231 (26.9%)	3,380 (27.4%)	10,209 (27.6%)
Outpatient spending	1,991.293 (2,313.546)	1,886.115 (2,219.072)	1,837.017 (2,179.634)	1,894.806 (2,229.713)
Spending on drugs	1,020.391 (2,051.645)	883.617 (1,994.187)	811.457 (1,777.318)	891.602 (1,939.061)
Inpatient spending	5,779.773 (4,250.430)	5,545.305 (4,062.713)	5,513.933 (4,126.985)	5,595.228 (4,133.327)
Spending on lab tests	392.609 (411.757)	382.827 (406.463)	361.862 (379.309)	378.138 (399.037)
Physiotherapy spending	832.501 (665.432)	818.602 (661.358)	810.608 (633.680)	819.304 (653.200)

The communities involved are generally large. The estimation sample is therefore not driven by the small communities for which $C_D(c)$ is weakly identified. Those excluded communities are numerous but nearly empty: the retained set covers 98.9–99.3% of patient–community assignments in each triennium.
The last column of Table 3 reports the pre-move prevalence of each extensive-margin outcome in the estimation sample: 90.7% for prescription drugs, 81.4% for laboratory tests, 27.6% for emergency care, 23.9% for physiotherapy and 13.9% for inpatient care. Outpatient care is close to universal in this population and is not tabulated. These are the denominators against which the percentage-point estimates below should be read.

Two features of the treatment deserve comment. First, movers on average shift towards more centralised and larger communities: mean $C_D(c)$ rises from 0.660 at origin to 0.702 at destination, and mean size from 426 to 535 providers. Origin-community and triennium fixed effects absorb the level of this drift, but it is the reason the specification is reported both with and without a control for destination size. Second, $C_D(c)$ is itself strongly increasing in community size: mean centralisation rises from 0.23–0.25 among communities with three or four providers to 0.60–0.63 among those with twenty-five or more (appendix A.4). A change in centralisation is therefore in part a change in scale, and the two cannot be fully separated in these data.

To strengthen the validity of the approach, Figure 3 shows the distribution of residualised (log) expenditure for movers in year 2 (after removing canton, year and origin community fixed effects), stacking all triennia together. The distributions of movers in different directions overlap substantially prior to the community switch, providing support for the identifying assumption that, conditional on fixed effects, they follow comparable expenditure trajectories before the move.

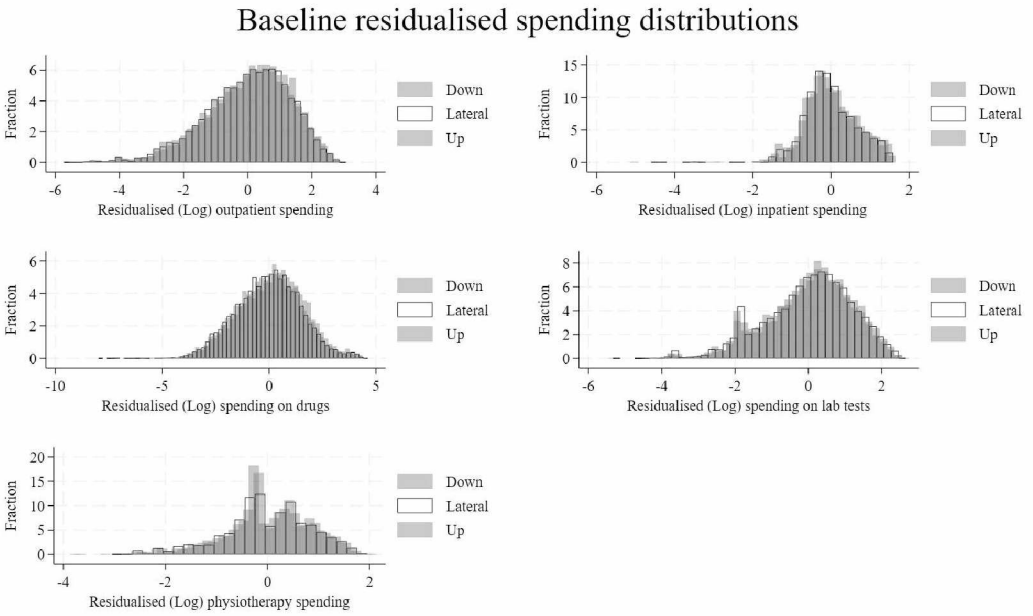

Figure 3: Distribution of residualised variation in outcomes

Furthermore, Figure 4 reports trends on the same residualised log-outcomes for movers split by direction of move. The pre-move trajectories are flat and statistically indistinguishable across groups, suggesting that even the stricter parallel-trends assumption (which would be invoked for identification in a difference-in-differences specification) could plausibly hold. A formal test of this condition is reported in Section 5.3.

To support the empirical plausibility of the communities detected by the algorithm of choice, Figure 5 shows a map of an exemplar canton (Ticino) which overlays a sketch of the provider communities across two triennia. Due to data restrictions, only distances between the providers and the MedStat region where they are legally registered are known, but not their exact geolocation. For representation purposes, the map shows providers randomly displaced around the centroid of their MedStat regions, coloured based on the communities they belong to. The map shows that the algorithm identifies geographically contained collaboration, with few providers in rural areas collaborating with colleagues located in urban centres. Finally, Figure 6 shows the distribution of the three network metrics of interest for the last triennium.

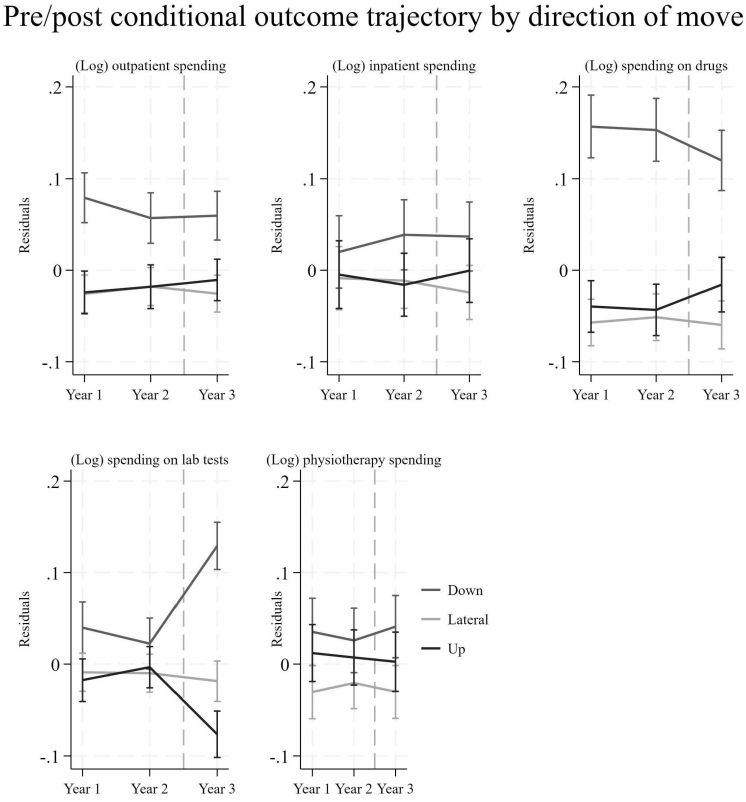

Figure 4: Conditional pre-trends by direction of centralisation change

5.2 Main results

Panels A and B of Tables 4 and 5 report estimates of equation (10). The treatment is the change in Freeman degree centralisation between the origin and destination community, standardised so that coefficients are read per standard deviation of $\Delta C_D(c)$. That standard deviation is 0.250 on a measure bounded between zero and one, so a one-standard-deviation move corresponds to roughly a quarter of the theoretical range of the index. Models are estimated on a consistent analytical sample which only includes patients whose assigned community is comprised of 3 or more providers.

Moves towards more centralised communities reduce healthcare utilisation on every outcome examined. On the extensive margin (Table 4, Panel A), a one-standard-deviation increase in $C_D(c)$ lowers the probability of incurring any expenditure by 3.54 percentage points for prescription drugs, 3.60 for emergency care, 3.10 for laboratory tests, 2.81 for physiotherapy, 2.30 for inpatient care and 0.50 for outpatient care. Relative to the pre-move prevalences in Table 3, these correspond to reductions of between 4% and 17% of baseline prevalence for the five tabulated outcomes, the largest in relative terms being inpatient care (17%), emergency care (13%) and physiotherapy (12%). On the intensive margin (Table 5, Panel A) the effects concentrate in two categories. Conditional on positive spending, a one-standard-deviation increase in $C_D(c)$ reduces laboratory expenditure by 14.3% and drug expenditure by 9.7%. Proportional changes on the intensive margin are reported throughout in log points, that is as $100\hat{\beta}$ rather than $100(e^{\hat{\beta}} - 1)$; the two differ by less than one percentage point at the magnitudes estimated here. The estimates for outpatient care (-1.9%) and physiotherapy (-4.1%) are negative but imprecise.

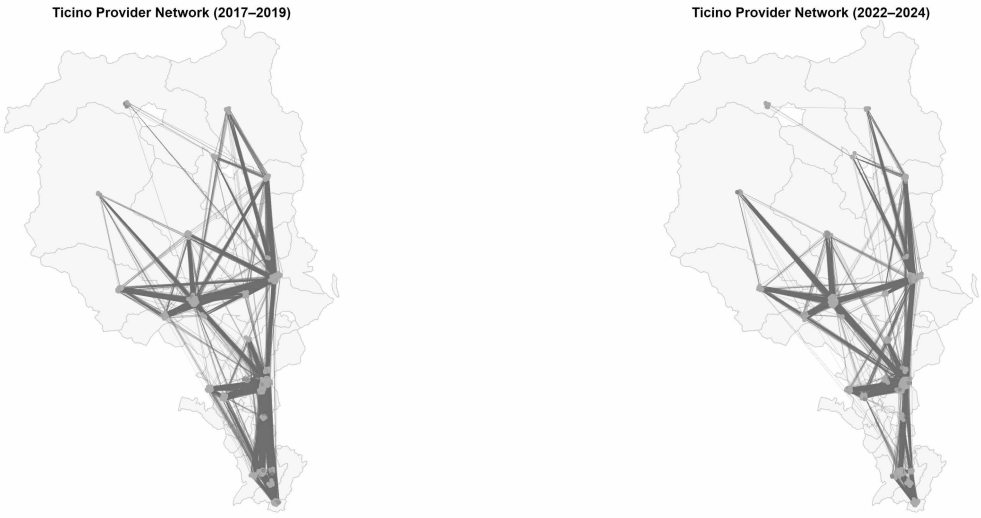

Figure 5: Communities of provision identified in the Canton of Ticino

Notes: 1) Communities are identified by applying the Leiden algorithm (Traag et al., 2019) with the Constant Potts Model (CPM) objective (Traag et al., 2011) to a spatially-weighted patient-sharing network among primary care providers. Each node represents a provider; colours identify communities. 2) To avoid overfitting the figure, grey edges representing patient sharing are plotted only for providers that are part of a same community but registered in two separate MedStat regions, that is: (a) two providers within a same MedStat region in a same network appear simply as two dots of a same colour, despite having shared patients that led them to be in the same community, and (ii) edges connecting providers grouped in different communities are muted.

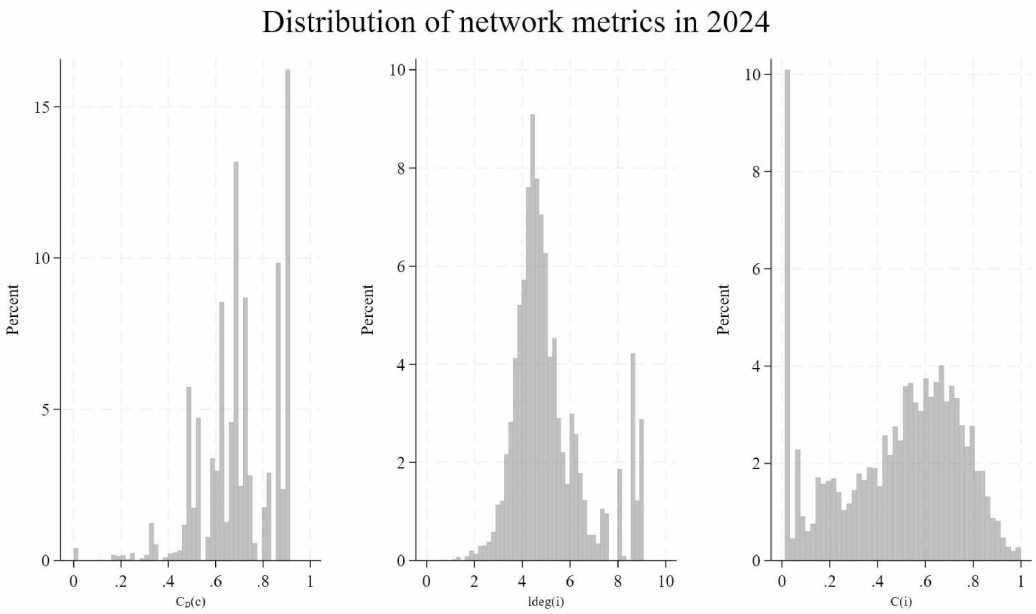

Figure 6: Distribution of network metrics

Table 4: Main results for the extensive margin

Panel A) $\Delta C_D(c)$, community degree centralisation (with $\ln H_{cd}$)

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00501*	-0.0230*	-0.0354***	-0.0310**	-0.0281**	-0.0360**
	(0.00196)	(0.00917)	(0.0103)	(0.0111)	(0.00887)	(0.0113)
N	36704	36704	36704	36704	36704	36704
N_clust	312	312	312	312	312	312
r2	0.0233	0.0608	0.103	0.0803	0.189	0.0783

Panel B) $\Delta C_D(c)$, community degree centralisation (without $\ln H_{cd}$)

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00593***	-0.0221***	-0.0492***	-0.0478***	-0.0445***	-0.0525***
	(0.00102)	(0.00453)	(0.00672)	(0.00741)	(0.00495)	(0.00638)
N	36704	36704	36704	36704	36704	36704
N_clust	312	312	312	312	312	312
r2	0.0232	0.0608	0.103	0.0800	0.189	0.0781

Panel C) Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	-0.00247*	-0.00352*	-0.0160***	-0.0155***	-0.0136***	-0.0104***
	(0.000985)	(0.00158)	(0.00235)	(0.00352)	(0.00150)	(0.00232)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0356	0.0734	0.131	0.101	0.206	0.1000

Panel D) $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	-0.0215***	-0.0276	-0.0569**	0.0369	-0.00460	-0.0549*
	(0.00551)	(0.0165)	(0.0214)	(0.0341)	(0.0184)	(0.0234)
Δldeg(i)	-0.00587***	-0.00787*	-0.0250***	-0.00972	-0.0144***	-0.0190***
	(0.00175)	(0.00352)	(0.00523)	(0.00806)	(0.00329)	(0.00458)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0367	0.0736	0.132	0.101	0.206	0.100

Notes: Control variables included in all models. All outcome variables are binary indicators of any spending; coefficients are therefore percentage-point changes. In Panels A and B the treatment is standardised, so coefficients are per standard deviation of $\Delta C_D(c)$ ($\sigma = 0.250$); in Panels C and D they are semi-elasticities with respect to the provider metric. Panels A and B absorb origin community, canton and triennium fixed effects; Panel A additionally conditions on $\ln H_{cd}$. Panels C and D absorb origin and destination community, canton and triennium fixed effects. Models estimated separately by triennium (Appendix B), the findings are fully consistent with the pooled results above, though individually imprecise. Standard errors (in parentheses) are clustered three-way at the origin community, destination community and patient level in Panels A and B, and two-way at the provider and patient level in Panels C and D; stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel A estimates condition on the log size of the destination community. Because centralisation and size are positively related, this is the conservative specification: it isolates variation in topology holding scale fixed. Removing the destination-size control (Panel B of each table) raises the coefficient in absolute value for eleven of the twelve estimates, by between 16% and 37% on the extensive margin and between 2% and 62% on the intensive margin; the exception is inpatient care, which is marginally larger with the control than without it. The two specifications bracket the effect of moving to a more centralised community: the first attributes to topology only what cannot be explained by scale, the second attributes the combined change to the move. Both are reported; the discussion below relies on the conservative version.

Table 5: Main results for the intensive margin

Panel A) $\Delta C_D(c)$, community degree centralisation (with $\ln H_{c^d}$)

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	-0.0189	-0.0970**	-0.143***	-0.0413
	(0.0323)	(0.0306)	(0.0368)	(0.0277)
N	35972	32495	29117	8284
N_clust	306	306	299	241
r2	0.292	0.454	0.157	0.138

Panel B) $\Delta C_D(c)$, community degree centralisation (without $\ln H_{c^d}$)

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	-0.0493*	-0.132***	-0.174***	-0.0422*
	(0.0224)	(0.0196)	(0.0284)	(0.0195)
N	35972	32495	29117	8284
N_clust	306	306	299	241
r2	0.291	0.454	0.157	0.138

Panel C) Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-0.0351**	-0.0435***	-0.0727***	-0.0206***
	(0.0129)	(0.00771)	(0.00831)	(0.00567)
N	35943	32461	29089	8246
N_clust	5878	5769	5552	3292
r2	0.310	0.466	0.189	0.170

Panel D) $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	-0.173	-0.178**	-0.179*	-0.0826
	(0.107)	(0.0669)	(0.0864)	(0.0689)
Δldeg(i)	-0.0624*	-0.0718***	-0.101***	-0.0341**
	(0.0275)	(0.0149)	(0.0188)	(0.0121)
N	35943	32461	29089	8246
N_clust	5878	5769	5552	3292
r2	0.310	0.466	0.190	0.170

Notes: Control variables included in all models. All outcome variables are log-transformed spending variables. so coefficients approximate proportional changes. In Panels A and B the treatment is standardised, so coefficients are per standard deviation of $\Delta C_D(c)$ ($\sigma = 0.250$); in Panels C and D they are elasticities with respect to the provider metric. All specifications condition on the untransformed lagged outcome and on an indicator for zero spending at baseline. Models estimated separately by triennium (Appendix B), the findings are fully consistent with the pooled results above, though individually imprecise. Fixed effects and clustering are as in Table 4; stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel C reports the effect of a change in the log degree centrality of the patient's main primary care provider. Every coefficient is negative and statistically significant on both margins. Because the regressor is a change in log degree, the coefficients are semi-elasticities on the extensive margin and elasticities on the intensive margin. A doubling of the main PCP's degree reduces the probability of incurring laboratory expenditure by 1.07 percentage points and laboratory spending by 5.0%, with reductions of 1.11 and 0.94 percentage points for prescription drugs and physiotherapy respectively.

Panel D adds the local clustering coefficient around the main PCP, holding degree constant. Three extensive-margin coefficients are individually significant: outpatient care (-2.1 p.p.), prescription drugs (-5.7 p.p.) and emergency care (-5.5 p.p.), as are two intensive-margin coefficients, prescription drugs (-17.8%) and laboratory tests (-17.9%). Triadic closure around a patient's provider lowers the propensity to start drug and emergency care, and the amount spent on drugs and laboratory tests, but has no detectable effect on inpatient care, physiotherapy or the extensive margin of laboratory testing once the provider's overall connectedness is held constant. Given the weight the patient-sharing literature places on local density as a mechanism for informal coordination, the limited reach of this channel relative to overall provider degree is interesting.

5.3 Robustness Checks

To evaluate the validity of my primary approach, I perform a battery of complementary robustness checks (see appendix C for further details).

First I regress the pre-move value of each outcome on the future treatment, using the same covariates as the estimating equation. The sole exception is inpatient care, where lagged length of stay and lagged severity are near-deterministic functions of having had an inpatient stay in the baseline year; for that outcome both are excluded. For the majority of outcomes, the results suggest either no selection or a mild negative selection pattern that - if anything - would attenuate my estimates. The resulting interpretable set is slightly narrower due to mild residual selection on some outcomes. For community centralisation, five of six extensive-margin pre-move associations carry the *opposite* sign to the estimated effect and none is significant, so residual selection works against the finding rather than generating it. The exception is emergency care, where the pre-move association is significant, of the same sign, and should therefore be read as suggestive only. For provider degree the same pattern holds on five of six extensive-margin outcomes; inpatient care is the exception and is not interpretable. On the intensive margin, laboratory tests and physiotherapy pass cleanly my test for selection while outpatient care and prescription drugs do not. For local clustering, the drug and emergency-care results survive on both margins, whereas the outpatient extensive-margin estimate does not.

Second, following McKenzie (2012), I re-estimate the main empirical specifications using a First Differences (FD) approach. ANCOVA and FD - which agree closely in my results - bracket the true effect under opposing assumptions about outcome dynamics; in panel settings with baseline controls the two are complementary rather than competing (Angrist and Pischke, 2009).

Third, I re-estimate my ANCOVA specification dropping all baseline patient demographics, health severity and insurance plan controls (while retaining the lagged outcome). This check leaves the sign of every community-level estimate unchanged and raises the magnitudes, as expected when covariates absorb part of the variation; the two intensive-margin estimates that are imprecise in the main specification, outpatient care and physiotherapy, become significant. The laboratory elasticity for provider degree is essentially unaffected.

Finally, two further specifications corroborate the main result. First, replacing the continuous treatment with indicators for moves up and down the national quintile distribution of $C_D(c)$ as in Equation (11) yields coefficients of the same sign where significant, -2.28 percentage points for laboratory tests and -3.19 for emergency care among upward movers, but is otherwise imprecise and null throughout on the intensive margin. This is expected: binning a continuous measure into quintiles discards the magnitude of each move, and the tied mass of communities at $C_D(c) = 0$ leaves one quintile boundary poorly identified. Second, an indicator for the transition between solo and networked practice (a treatment defined also for the patients whom the $H_c \geq 3$ restriction excludes) shows that moving from a networked practice to a solo one raises conditional spending on laboratory tests by 11.8%, on prescription drugs by 11.5% and on outpatient care by 5.3%. Movement in the opposite direction is imprecisely estimated, since few patients make it: every extensive-margin coefficient is insignificant, as are three of the four intensive-margin ones, the exception being laboratory spending (-26.7%), which falls when a patient moves from a solo to a networked practice. Both point the same way as the main estimates: less connectedness leading to more spending.

6 Discussion

Integrated care and provider collaboration are widely promoted as instruments for containing healthcare expenditure and improving coordination for patients with complex needs. The empirical evidence remains mixed, in large part because informal collaboration is difficult to observe and harder to measure. This paper constructs patient-sharing networks from a decade of Swiss claims, characterises primary care communities using a spatially weighted Leiden partition, and asks whether connectedness in these communities through the patients' main primary care providers affects patient-level utilisation and spending. To isolate this effect, the paper uses a within-movers design that compares patients who change provider community against each other rather than against stayers.

The central finding is that greater connectedness is associated with lower utilisation, and that this holds at both scales examined. A one-standard-deviation move towards a more centralised community lowers the propensity to consume care in all six categories, by between 0.5 and 3.6 percentage points, and reduces conditional spending on laboratory tests by 14.3% and on prescription drugs by 9.7%. A doubling of the main provider's degree lowers the propensity to consume in all six categories and conditional spending in all four, by 5.0% for laboratory tests, 3.0% for drugs, 2.4% for outpatient care and 1.4% for physiotherapy. Local clustering around the provider, holding degree fixed, lowers the propensity to incur drug and emergency spending and conditional spending on drugs and laboratory tests. The three sets of estimates point the same way, which is what makes them worth reading together.

The effects are precisely estimated but economically modest. Among patients who use them, mean annual spending in the pre-move year is 378 CHF for laboratory tests and 892 CHF for prescription drugs (Table 3). A one-standard-deviation increase in community centralisation therefore corresponds to roughly 54 CHF less laboratory spending and 86 CHF less drug spending per patient-year; a doubling of the main provider's degree, to roughly 19 CHF and 27 CHF. Set against mean annual outpatient spending of about 1,900 CHF for the same patients, these are second-order quantities. What the estimates establish is a direction and a mechanism, not a lever large enough to bend an expenditure curve on its own.

Whilst the set of results that survives the selection test is narrower than the coefficient tables suggest, the valid patterns are consistent with avoided duplication. Reductions concentrate in laboratory testing and prescription drugs, the two categories in which repetition is both clinically plausible and cheap to trigger, and in which a coordinating provider with visibility over a patient's recent results can decline to re-order. Inpatient care, which is largely outside primary care gatekeeping, moves least in relative terms. This connects directly to Shin (2023), who find that centrally positioned physicians are better placed to navigate the system and route patients towards more efficient specialised care. I emphasise that this reading is an interpretation of a coefficient pattern rather than a tested mechanism. The design does not identify duplication directly, and the data contain no test-level detail that would allow a repeated-test measure to be constructed. Two implications would be testable with richer data and would discriminate this channel from alternatives: the effect should be larger for patients seeing more distinct providers, where the scope for duplication is greater, and it should concentrate in repeat rather than first-time tests within a category.

Brekke et al. (2024) set out two opposing forces that structural integration sets in motion: efficiency gains through information sharing, clinical coordination and the elimination of duplicated care, against softened competition and restricted patient choice, which can raise service volume. The estimates here are consistent with the first channel operating at both the market and the clinician level, but they are silent on the second. A within-movers design with origin-community fixed effects incorporates no price variation and no market-level counterfactual; it cannot detect softened competition even if it were present. Distinguishing the two would require variation in market structure (e.g. provider entry and exit, or practice consolidation) of a kind this design deliberately conditions away.

This work has several limitations, four of which directly limit the interpretation of my estimates. First, centralisation and community size cannot be fully separated: mean centralisation more than doubles between the smallest and largest communities in the sample, so a change in topology is partly a change in scale. Conditioning on destination size attenuates the two significant intensive-margin estimates by 18% and 27%, and the specifications with and without that control should be read as bounding the effect rather than as competing estimates of it. Second, the design accounts for time-invariant patient heterogeneity but not for time-varying shocks coincident with a move; a patient's decision to change community may itself coincide with an unmeasured health shock. In this work I only provide weak evidence that this does not drive the findings. Third, networks are constructed entirely from billing records, and so capture neither informal clinical discussion, multi-disciplinary team meetings, nor any coordination that is not billed; this biases towards understating collaboration, and one that no claims-based network measure escapes. Fourth, with the available data the community detection algorithm cannot distinguish collaboration mandated by insurance arrangements such as HMOs from collaboration emerging endogenously between providers, and the two have different policy implications.
Two further constraints limit external validity. The setting is a mandatory private insurance system with cantonal governance and near-universal registration with a single primary care provider, which limits extrapolation to tax-funded or single-payer systems. And the data contain no clinical parameters, health outcomes or measures of patient satisfaction, so reductions in expenditure cannot be distinguished from under-provision. That distinction is not a technicality: it determines whether these findings describe efficiency gains or rationing, and nothing in the analysis can settle it.

Discounting these limitations, the results support a few actionable conclusion. Community centralisation is not a policy instrument: no authority sets the degree distribution of a patient-sharing network. However, provider connectedness is different on both counts. It survives every sample restriction and partition tested here, and it is amenable to policy through practice consolidation, shared referral infrastructure, and the incentives that determine whether a primary care physician maintains a broad set of working relationships or a narrow one. This actionable finding is also the most robust. Yet, this should not be read as an argument for promoting connectedness as such. As Brekke et al. (2024) emphasise, organisational integration does not automatically deliver efficiency or welfare gains. The estimates here are equally consistent with better-connected providers suppressing appropriate care alongside inappropriate care, and without health outcomes the two cannot be told apart. What the results do establish is narrower and, I think, still worth stating: the concern that hub-and-spoke structures mechanically inflate service volume finds no support amongst the communities where it is defined, and the connectedness of a patient's own provider is consistently associated with lower spending in exactly the categories where duplication is easiest to avoid.

References

Albertini, M., Bakx, P., and Mazzonna, F. (2025). Health and Labor Market Consequences of Low-Value Care: The Role of Practice Style. Working Paper 17771, IZA Discussion Papers.

Angrist, J. D. and Pischke, J.-S. (2009). Mostly harmless econometrics: an empiricist's companion. Princeton University Press, Princeton, USA.

Badinski, I., Finkelstein, A., Gentzkow, M., and Hull, P. (2023). Geographic Variation in Healthcare Utilization: The Role of Physicians.

Bautista, M. A. C., Nurjono, M., Lim, Y. W., Dessers, E., and Vrijhoef, H. J. (2016). Instruments Measuring Integrated Care: A Systematic Review of Measurement Properties. The Milbank Quarterly, 94(4):862–917. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0009.12233.

Baxter, S., Johnson, M., Chambers, D., Sutton, A., Goyder, E., and Booth, A. (2018). The effects of integrated care: a systematic review of UK and international evidence. BMC Health Services Research, 18(1):350.

Bischof, T. and Kaiser, B. (2021). Who cares when you close down? The effects of primary care practice closures on patients. Health Economics, 30(9):2004–2025. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hec.4287.

Bouton, C., Journeaux, M., Jourdain, M., Angibaud, M., Huon, J.-F., and Rat, C. (2023). Interprofessional collaboration in primary care: what effect on patient health? A systematic literature review. BMC Primary Care, 24(1):253.

Brekke, K. R., Siciliani, L., and Straume, O. R. (2024). Competition, quality and integrated health care. Journal of Health Economics, 95:102880.

Burns, L. R., Nembhard, I. M., and Shortell, S. M. (2022). Integrating network theory into the study of integrated healthcare. Social Science & Medicine, 296:114664.

Damery, S., Flanagan, S., and Combes, G. (2016). Does integrated care reduce hospital activity for patients with chronic diseases? An umbrella review of systematic reviews. Publisher: British Medical Journal Publishing Group Section: Health services research.

de Graaf-Ruizendaal, W. A., van der Hoek, L., and de Bakker, D. H. (2018). Are low and high utilization related to the way GPs manage their practices? An observational study. BMC Family Practice, 19(1):46.

Ehlert, A. and Oberschachtsiek, D. (2014). Does managed care reduce health care expenditure? Evidence from spatial panel data. International Journal of Health Care Finance and Economics, 14(3):207–227.

Finkelstein, A., Gentzkow, M., and Williams, H. (2021). Place-Based Drivers of Mortality: Evidence from Migration. American Economic Review, 111(8):2697–2735.

Foy, R., Hempel, S., Rubenstein, L., Suttorp, M., Seelig, M., Shanman, R., and Shekelle, P. G. (2010). Meta-analysis: Effect of Interactive Communication Between Collaborating Primary Care Physicians and Specialists. Annals of Internal Medicine, 152(4):247–258. Publisher: American College of Physicians.

Goldzahl, L., Stokes, J., and Sutton, M. (2022). The effects of multi-disciplinary integrated care on healthcare utilization: Evidence from a natural experiment in the UK. Health Economics, 31(10):2142–2169. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hec.4561.

Gyger, P. (2020). Regionenreport 2020. Technical report, Helsana Group, Zurich.

Hietapakka, L., Sinervo, T., Väisänen, V., Niemi, R., Gutvilig, M., Linnaranta, O., Suvisaari, J., Hakulinen, C., and Elovainio, M. (2025). Patient-sharing networks among Finnish primary healthcare professionals taking care of patients with mental health or substance use problems: a register study. Publisher: British Medical Journal Publishing Group Section: Health services research.

Hjalmarsson, L., Kaiser, B., and Bischof, T. (2023). The impact of physician exits in primary care: A study of practice handovers. Health Policy, 135:104867.

Jackson, M. O., Rogers, B. W., and Zenou, Y. (2017). The Economic Consequences of Social-Network Structure. Journal of Economic Literature, 55(1):49–95.

Katon, W., Russo, J., Von Korff, M., Lin, E., Simon, G., Bush, T., Ludman, E., and Walker, E. (2002). Long-term effects of a collaborative care intervention in persistently depressed primary care patients. Journal of General Internal Medicine, 17(10):741–748.

Kee, K., Nies, H., van Wieringen, M., and Beersma, B. (2023). From Integrated Care to Integrating Care: A Conceptual Framework of Behavioural Processes Underlying Effective Collaboration in Care. International Journal of Integrated Care, 23(4):4.

Kesternich, E. and Rank, O. (2022). Beyond patient-sharing: Comparing physician- and patient-induced networks. Health Care Management Science, 25(3):498–514.

Landon, B. E., Onnela, J.-P., Keating, N. L., Barnett, M. L., Paul, S., O'Malley, A. J., Keegan, T., and Christakis, N. A. (2013). Using Administrative Data to Identify Naturally Occurring Networks of Physicians. Medical Care, 51(8):715.

Matulis, J., Liu, S., Mecchella, J., North, F., and Holmes, A. (2017). Choosing Wisely: A Quality Improvement Initiative to Decrease Unnecessary Preoperative Testing. BMJ Quality Improvement Reports, 6(1). Publisher: British Medical Journal Publishing Group.

McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. Journal of Development Economics, 99(2):210–221.

Merçay, C. (2016). Expérience de la population âgée de 18 ans et plus avec le système de santé – Situation en Suisse et comparaison internationale. Analyse de l'International Health Policy Survey 2016 du Commonwealth Fund sur mandat de l'Office fédéral de la santé publique (OFSP). Technical Report 56, Observatoire suisse de la santé, Bern.

Morley, L. and Cashell, A. (2017). Collaboration in Health Care. Journal of Medical Imaging and Radiation Sciences, 48(2):207–216.

Peikes, D., Chen, A., Schore, J., and Brown, R. (2009). Effects of Care Coordination on Hospitalization, Quality of Care, and Health Care Expenditures Among Medicare Beneficiaries: 15 Randomized Trials. JAMA, 301(6):603–618.

Rocks, S., Berntson, D., Gil-Salmerón, A., Kadu, M., Ehrenberg, N., Stein, V., and Tsiachristas, A. (2020). Cost and effects of integrated care: a systematic literature review and meta-analysis. The European Journal of Health Economics, 21(8):1211–1221.

Saint-Pierre, C., Herskovic, V., and Sepúlveda, M. (2018). Multidisciplinary collaboration in primary care: a systematic review. Family Practice, 35(2):132–141.

Schmid, C. P. R., Beck, K., and Kauer, L. (2018). Chapter 16 - Health Plan Payment in Switzerland. In McGuire, T. G. and van Kleef, R. C., editors, Risk Adjustment, Risk Sharing and Premium Regulation in Health Insurance Markets, pages 453–489. Academic Press.

Shin, E. (2023). Physician Connectedness and Referral Choice*. Oxford Bulletin of Economics and Statistics, 85(6):1238–1261.

Stokes, J., Shah, V., Goldzahl, L., Kristensen, S. R., and Sutton, M. (2021). Does prevention-focused integration lead to the triple aim? An evaluation of two new care models in England. Journal of Health Services Research & Policy, 26(2):125–132. Publisher: SAGE Publications.

Struckmann, V., Leijten, F. R. M., van Ginneken, E., Kraus, M., Reiss, M., Spranger, A., Boland, M. R. S., Czypionka, T., Busse, R., and Rutten-van Mölken, M. (2018). Relevant models and elements of integrated care for multi-morbidity: Results of a scoping review. Health Policy, 122(1):23–35.

Suter, E., Oelke, N. D., Lima, M. A. D. d. S., Stiphout, M., Janke, R., Witt, R. R., Vliet-Brown, C. V., Schill, K., Rostami, M., Hepp, S., Birney, A., Al-Roubaiai, F., and Marques, G. Q. (2017). Indicators and Measurement Tools for Health Systems Integration: A Knowledge Synthesis. International Journal of Integrated Care, 17(6).

Traag, V. A., Van Dooren, P., and Nesterov, Y. (2011). Narrow scope for resolution-limit-free community detection. Physical Review E, 84(1):016114.

Traag, V. A., Waltman, L., and van Eck, N. J. (2019). From Louvain to Leiden: guaranteeing well-connected communities. Scientific Reports, 9(1):5233.

Valentijn, P. P., Schepman, S. M., Opheij, W., and Bruijnzeels, M. A. (2013). Understanding integrated care: a comprehensive conceptual framework based on the integrative functions of primary care. International Journal of Integrated Care, 13:e010.

Wang, C., , and Wang, F. (2022). GIS-automated delineation of hospital service areas in Florida: from Dartmouth method to network community detection methods. Annals of GIS, 28(2):93–109. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/19475683.2022.2026470.

Wang, C., Wang, F., and Onega, T. (2021). Network optimization approach to delineating health care service areas: Spatially constrained Louvain and Leiden algorithms. Transactions in GIS, 25(2):1065–1081. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/tgis.12722.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. Nature, 393(6684):440–442.

Wulsin, L. R., Söllner, W., and Pincus, H. A. (2006). Models of Integrated Care. Medical Clinics, 90(4):647–677. Publisher: Elsevier.

Yeager, S. (2005). Interdisciplinary Collaboration: The Heart and Soul of Health Care. Critical Care Nursing Clinics, 17(2):143–148. Publisher: Elsevier.

Zhang, Y., Stokes, J., Anselmi, L., Bower, P., and Xu, J. (2025). Can integrated care interventions strengthen primary care and improve outcomes for patients with chronic diseases? A systematic review and meta-analysis. Health Research Policy and Systems, 23(1). Publisher: Springer Science and Business Media LLC.

Zonneveld, N., Driessen, N., Stüssgen, R. A. J., and Minkman, M. M. N. (2018). Values of Integrated Care: A Systematic Review. International Journal of Integrated Care, 18(4).

Appendix

A Implementation of community detection algorithms

A.1 Sensitivity to threshold for minimal relevant patient sharing

The main analysis uses a threshold of at least 3 shared patients to establish an edge between two providers i and j; that is $e_{i,j} = 0$ for all providers who only share 1 or 2 patients. To assess the sensitivity of my provider network partitions against weak or spurious patient-sharing links, I relax this threshold. Specifically, I rerun my community detection algorithm for all triennia imposing a stricter relative thresholding filter to only the top 20% edges for each provider. That is, an edge between two providers is retained if and only if its distance-decayed weight falls within the top 20% of all incident edges (i.e., $\geq 80^{\text{th}}$ percentile) for *both* the sender and receiver nodes. This approach is consistent with Landon et al. (2013). With the resulting partition, I compare the number of communities identified in each canton between the baseline and stricter thresholds, and the results for my main empirical specifications.

As expected, the number of communities identified across cantons is substantially smaller compared to my main parametrisation.

Number of communities identified across cantons over time restricting edges to top 20% patient sharing relationships

Canton	1517	1719	2224
AG	18	11	19
AI	2	2	1
AR	6	5	6
BE	127	122	81
BL	8	12	13
BS	35	34	25
FL	4	7	12
FR	10	10	8
GE	44	35	45
GL	2	2	1
GR	28	25	20
JU	4	3	3
LU	13	13	10
NE	13	8	12
NW	1	1	2
OW	2	2	2
SG	14	16	15
SH	13	9	10
SO	11	13	15
SZ	3	5	3
TG	16	15	11
TI	18	13	25
UR	2	3	2
VD	67	57	60
VS	19	17	20
ZG	5	4	1
ZH	84	81	67

Under this stricter calibration for edge strength, *ex ante* one would expect a clearer or stronger effect for individual provider scale $\Delta ldeg(i)$ due to the reduction of non-systematic measurement noise in shared billing ties. Conversely, the influence on both community-level centralisation $C_D(c)$ and local clustering $\Delta C(i)$ is *ex ante* ambiguous. Filtering for mutual top 20% ties alters the internal geometry of detected communities: it can flatten relative degree hierarchies within clusters (potentially dampening $C_D(c)$) and mechanically shift $C(i)$ from a measure of open patient-sharing networks to one of highly integrated, closed co-management silos. Provider degree survives essentially intact: the result is not an artefact of counting weak or incidental ties, since the strongest-tie subgraph gives the same answer. Community centralisation collapses to insignificance, so the filter removes the variation the estimate is identified from rather than testing it. Local clustering reverses sign, from negative on the full network to positive and significant on the filtered one. This too is interpretable, and it is the sharpest evidence in the paper for the distinction it draws: $C(i)$ computed on all ties measures an open, loosely coupled neighbourhood, whereas $C(i)$ computed on strong ties measures a closed co-management silo. They are different constructs, and they carry opposite associations with spending. Taken together with the fact that provider degree is the metric that survives every sample and partition tested, the evidence favours a reading in which what matters is a provider's reach across the system rather than the density of their immediate peer group.

Main results for the extensive margin

Panel A) $\Delta C_D(c)$, community degree centralisation (with $\ln H_{c^d}$)

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00137	-0.00833*	-0.00191	-0.00459	-0.00213	-0.00779
	(0.00115)	(0.00423)	(0.00699)	(0.00593)	(0.00554)	(0.00703)
N	38537	38537	38537	38537	38537	38537
N_clust	488	488	488	488	488	488
r2	0.0274	0.0678	0.107	0.0838	0.196	0.0849

Panel B) Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	-0.00233*	-0.00130	-0.0123***	-0.0137***	-0.0116***	-0.00647**
	(0.000923)	(0.00149)	(0.00251)	(0.00297)	(0.00152)	(0.00243)
N	38485	38485	38485	38485	38485	38485
N_clust	5955	5955	5955	5955	5955	5955
r2	0.0408	0.0838	0.134	0.106	0.214	0.108

Panel D) $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	0.00232	0.00808	0.0118**	0.0226**	0.0206**	0.0151
	(0.00156)	(0.00590)	(0.00420)	(0.00805)	(0.00679)	(0.00777)
Δldeg(i)	-0.00206*	-0.000364	-0.0109***	-0.0111***	-0.00922***	-0.00471*
	(0.000851)	(0.00161)	(0.00267)	(0.00278)	(0.00168)	(0.00236)
N	38485	38485	38485	38485	38485	38485
N_clust	5955	5955	5955	5955	5955	5955
r2	0.0409	0.0839	0.134	0.106	0.214	0.109

Notes: Control variables included in all models. All outcome variables are binary indicators of any spending; coefficients are therefore percentage-point changes. In Panel A the treatment is standardised, so coefficients are per standard deviation of $\Delta C_D(c)$ ($\sigma = 0.250$); in Panel B it is a semi-elasticity with respect to the provider metric. Panel A absorbs origin community, canton and triennium fixed effects and condition on $\ln H_{c^o}$; Panel A additionally conditions on $\ln H_{c^d}$. Panels B and C absorb origin and destination community, canton and triennium fixed effects. Standard errors (in parentheses) are clustered three-way at the origin community, destination community and patient level in Panel, and two-way at the provider and patient level in Panels B and C; stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Main results for the intensive margin

Panel A) $\Delta C_D(c)$, community degree centralisation (with $\ln H_{c^d}$)

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	0.0587**	-0.00571	-0.0127	0.000635
	(0.0213)	(0.0215)	(0.0285)	(0.0167)
N	37768	34222	30578	8825
N_clust	480	484	469	357
r2	0.294	0.460	0.159	0.142

Panel B) Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-0.0291*	-0.0402***	-0.0618***	-0.0199***
	(0.0120)	(0.00711)	(0.00887)	(0.00577)
N	37717	34166	30521	8741
N_clust	5920	5828	5609	3410
r2	0.314	0.475	0.193	0.181

Panel D) $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	0.0165	0.0425*	0.0530*	0.0399
	(0.0222)	(0.0213)	(0.0216)	(0.0273)
Δldeg(i)	-0.0272*	-0.0353***	-0.0557***	-0.0156*
	(0.0114)	(0.00709)	(0.00896)	(0.00640)
N	37717	34166	30521	8741
N_clust	5920	5828	5609	3410
r2	0.314	0.475	0.193	0.181

Notes: Control variables included in all models. All outcome variables are log-transformed spending variables, so coefficients approximate proportional changes. In Panel A the treatment is standardised, so coefficients are per standard deviation of $\Delta C_D(c)$ ($\sigma = 0.250$); in Panels B and C they are elasticities with respect to the provider metric. All specifications condition on the untransformed lagged outcome and on an indicator for zero spending at baseline. Fixed effects and clustering are as in main specification; stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

A.2 Calibration of γ

In this section I set out the procedure used to calibrate the density threshold γ which governs the Constant Potts Model objective function. To objectively define healthcare communities without imposing arbitrary geographic or administrative boundaries, I partition the patient-sharing network using the Leiden community detection algorithm under the Constant Potts Model (CPM) objective function (Traag et al., 2019). A key feature of the CPM framework is the resolution parameter γ, which dictates the structural granularity of the identified clusters. The number of communities detected varies directly with γ: lower values ($\gamma \to 0$) yield coarser partitions that capture macro-level network structures, whereas higher values of γ isolate finer, more granular clusters of provider interactions.

To select an optimal γ that reflects the underlying scale of provider interactions while avoiding stochastic instability, I implement a data-driven calibration routine on the first year (t_1) of each triennium. Specifically, I evaluate a logarithmic grid of 30 candidate values ranging from $\gamma = 10^{-6}$ to $\gamma = 10^{0}$. For each candidate γ, I evaluate the stability and reproducibility of the resulting network partitions by running the Leiden algorithm across $S = 5$ distinct random seeds.

I quantify partition stability across random seeds using pairwise Normalised Mutual Information (NMI). Formally, for two membership partitions A and B, $\text{NMI}(A, B)$ measures the information shared between the two groupings, normalised to yield 1 for identical partitions and 0 for independent assignments. For each resolution parameter γ, I compute the average pairwise NMI across all seed combinations:

$$\overline{\text{NMI}}(\gamma) = \frac{2}{S(S-1)} \sum_{i=1}^{S-1} \sum_{j=i+1}^{S} \text{NMI}\left(\mathbf{c}_i(\gamma), \mathbf{c}_j(\gamma)\right) \tag{A1}$$

where $\mathbf{c}_i(\gamma)$ denotes the vector of community assignments generated under seed i at resolution γ. Following this stability evaluation, I select the optimal resolution parameter γ^* according to a two-part decision rule:

$$\gamma^* = \arg\max_{\gamma} \left\{ \gamma \;\middle|\; \overline{\text{NMI}}(\gamma) \geq 0.995 \quad \text{and} \quad \mathcal{C}(\gamma) > 1 \right\} \tag{A2}$$

where $\mathcal{C}(\gamma)$ is the number of detected non-trivial communities. This rule selects the finest-grained structural partition (highest possible γ) that preserves near-perfect partition determinism ($\overline{\text{NMI}} \geq 0.995$) across algorithm runs.[5] By calibrating γ^* on baseline data, I ensure that the community definitions hold network resolution constant while allowing patient movement and provider degree centrality to evolve naturally over time.

[5] In the rare event that no candidate achieves $\overline{\text{NMI}} \geq 0.995$, the routine selects the non-trivial parameter γ that maximises $\overline{\text{NMI}}(\gamma)$.

A.3 Validation of community partition against institutional geography

A potential concern with data-driven community detection algorithms operating on health insurance claims is that detected partitions might yield geographically erratic clusters that bear little resemblance to established political and administrative healthcare boundaries. In Switzerland, healthcare governance, hospital planning, and tariff structures are organised predominantly at the cantonal level. To evaluate whether my unconstrained, spatially-weighted Leiden partitions respect institutional geography, I define a measure of cantonal spatial coherence.

For any detected multi-provider community $c \in \mathcal{C}$ (where size $H_c > 1$), let $H_{c,k}$ denote the number of healthcare providers in community c residing in canton $k \in \mathcal{K} = \{1, \ldots, 26\}$. The individual community coherence, Coherence_c, is defined as the proportion of providers belonging to the community's dominant canton:

$$\text{Coherence}_c = \frac{\max_{k \in \mathcal{K}} H_{c,k}}{N_c}, \quad \text{where } H_c = \sum_{k \in \mathcal{K}} H_{c,k}. \tag{A3}$$

To summarise network-wide adherence to cantonal borders across the entire physician population while excluding trivial singleton partitions ($H_c = 1$, for which $\text{Coherence}_c = 1.0$ by definition), I compute the overall size-weighted cantonal coherence $\Omega_{\mathcal{C}}$:

$$\Omega_{\mathcal{C}} = \frac{\sum_{c \in \mathcal{C}, H_c > 1} \max_{k \in \mathcal{K}} H_{c,k}}{\sum_{c \in \mathcal{C}, H_c > 1} H_c}. \tag{A4}$$

An overall weighted purity $\Omega_{\mathcal{C}} \to 1$ indicates that fine-grained referral clusters are near-perfectly nested within cantonal borders. Conversely, substantial departures from unity highlight cross-cantonal functional care markets.

Summary of community coherence across triennia

Triennium	Number of communities	Number of providers	Coherence $\Omega_{\mathcal{C}}$ (%)
2015-2017	299	15,770	91.57%
2017-2019	249	16,751	89.73%
2022-2024	448	16,080	92.33%

Across all evaluated triennia, I find high weighted cantonal coherence (e.g., $\Omega_{\mathcal{C}} \approx 90\%$), confirming that patient-sharing relationships naturally aggregate within institutional cantonal boundaries without requiring hard geographical constraints. The minority of low-coherence communities ($\text{Coherence}_c < 0.80$) concentrate systematically in areas with known cross-cantonal healthcare sharing patterns (Gyger, 2020). This is reassuring in relation to the algorithm's ability to accurately recover functional health markets where patient care flows naturally reflect political lines.

A.4 Community size, singletons, and the range over which $C_D(c)$ is defined

The size of communities in my preferred partitions for the three triennia analysed is characterised by wide variation and a few mass points, most notably at singleton communities. Singleton communities are defined as isolated clusters containing only a single provider ($H_c = 1$), which the algorithm identifies as not having tight collaborations with other generalists or specialists. In the context of Swiss ambulatory healthcare data, a singleton community represents an isolated, self-contained, or specialised solo practice that operates without sufficient patient-sharing ties to be clustered into a broader multi-provider network. Real-world singleton communities across Switzerland likely fall into three distinct clinical and geographic archetypes:

1. Isolated peripheral and alpine solo practices (*Landärzte*): Located in topographically isolated mountain valleys or rural peripheries (e.g., Val Müstair in Graubünden, the Bernese Oberland, or the upper Valais), these primary care physicians serve small, remote populations. Because geographic barriers restrict local co-utilization and patient referrals go directly to regional acute hospitals, these practices share very few routine ambulatory billing links with other local doctors, resulting in isolated single-node clusters.

2. Urban niche specialists: Operating in major metropolitan centres (e.g., Zurich, Geneva, or Lausanne), these highly specialised solo practitioners (e.g., paediatric endocrinology, highly specialised surgical subspecialties, private reproductive medicine) operate largely outside standard regional physician networks. Because their patient base is thinly scattered across many cantons rather than concentrated within a local primary care referral loop, their network ties fail to reach the threshold required for integration into a local functional market.

3. Unaffiliated out-of-hours and emergency-focused solo practices: Solo physicians providing high-volume primary care or urgent care outside organised regional call networks or integrated emergency centres. These episodic visits lack longitudinal follow-up or shared billing linkages with local GP clusters, producing isolated claim streams that do not cluster into broader ambulatory physician networks.

The table below presents the network summary statistics and singleton community diagnostics across the three triennia analysed (2015-2017, 2017-2019, and 2022-2024). Across all periods, singleton communities represent a relatively small share of overall healthcare providers (12.3% to 18.2%). Filtering out these single-provider units consolidates the sample into 249 to 448 non-singleton functional healthcare delivery communities ($H_c > 1$), preserving the vast majority of active healthcare providers (81.8% to 87.7%) within meaningful multi-provider communities.

Diagnostics on singleton communities across triennia

Triennium	Total Providers (H)	Total Communities (C)	Non-singleton communities ($H_c > 1$)	% Singleton communities	% Singleton providers
2015-2017	18,212	2,741	299	89.1%	13.4%
2017-2019	19,106	2,604	249	90.4%	12.3%
2022-2024	19,669	4,037	448	88.9%	18.2%

To provide further support for the validity of my partitioning algorithm in terms of singletons, the figure below projects the spatial distribution of singleton practices across Swiss MedStat regions. The geographic pattern strongly aligns with the taxonomy above: singleton rates (as a percentage of local providers) peak sharply in peripheral and high-altitude alpine MedStat regions (e.g., Valais, Graubünden, and the Bernese Oberland), confirming the spatial isolation of rural *Landärzte*. Conversely, urban MedStat centres (e.g. Zurich, Geneva) display a high absolute share but low relative share of singletons, predominantly driven by highly specialised niche practitioners operating outside local primary care referral clusters.

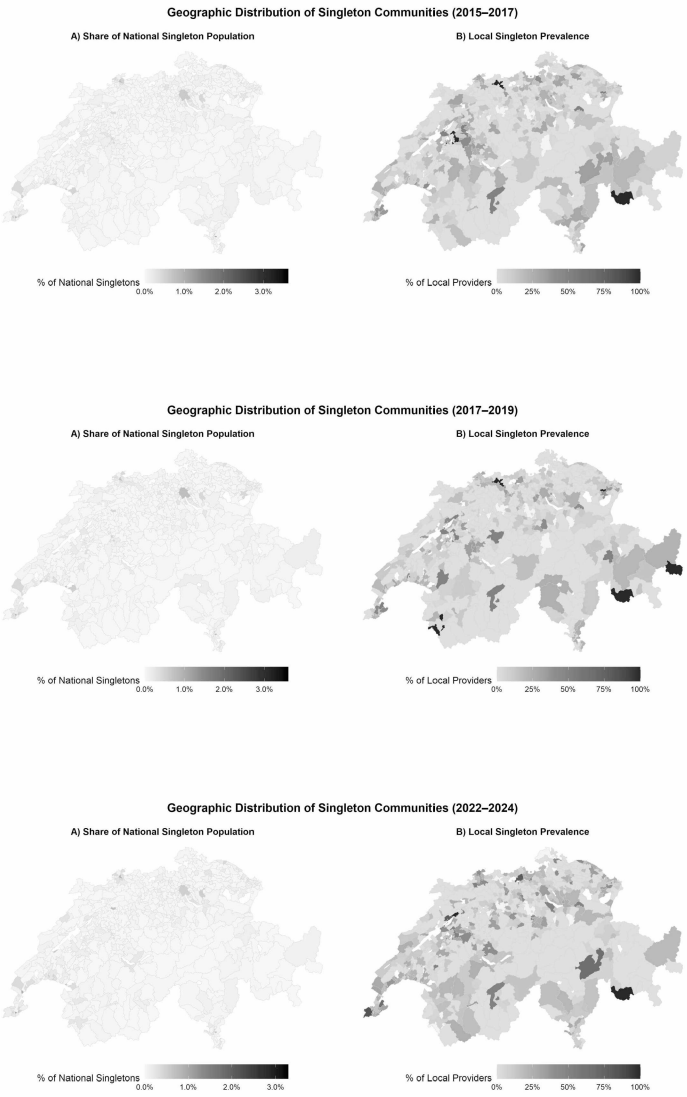

Note: Maps represent the distribution of listed outcomes by MedStat region. White regions represent areas with no providers (i.e. a null numerator and denominator).

Consequences for degree centralisation. The size distribution matters directly for the community-level treatment. Freeman degree centralisation has normaliser $(H_c - 1)(H_c - 2)$, so it is undefined for $H_c \leq 2$ and returned as missing by the community-detection routine. Imposing $H_c \geq 3$ leaves 148 of 2,741 communities in 2015-2017 (5.4%), 109 of 2,604 in 2017-2019 (4.2%) and 194 of 4,037 in 2022-2024 (4.8%). Because these are precisely the communities in which patients are concentrated, the restriction retains 98.9%, 99.3% and 98.9% of patient-community assignments respectively. It costs almost no patients and removes a tied mass that no discretisation of the index can handle.

Among the communities where it is defined, $C_D(c)$ is strongly increasing in size, as Table A1 shows. Mean centralisation roughly doubles between the smallest and largest size bands in every triennium. Two implications follow. First, a change in $C_D(c)$ is partly a change in scale, which is why the main specification is reported with and without a control for destination community size. Second, restricting to $H_c \geq 3$ does not fully remove the degenerate region: between 10.1% and 12.8% of valid communities still have $C_D(c)$ exactly equal to zero, almost all of them in the 3–4 provider band, which is why one quintile boundary is poorly identified in the directional specification of equation (11).

Table A1: Mean degree centralisation by community size band, communities with $H_c \geq 3$

	2015–2017		2017–2019		2022–2024	
Size band	N_c	mean C_D	N_c	mean C_D	N_c	mean C_D
$H_c = 3$–4	54	0.228	36	0.231	79	0.249
$H_c = 5$–9	17	0.274	14	0.289	34	0.305
$H_c = 10$–24	10	0.402	7	0.383	14	0.363
$H_c \geq 25$	67	0.596	52	0.624	67	0.630
All valid	148	0.412	109	0.436	194	0.399

Note: N_c is the number of communities in the band.

B Main results by triennium

The tables in this section re-estimate the specifications of Tables 4 and 5 separately for each triennium. Panel A uses the continuous, standardised change in community degree centralisation of equation (10), conditioning on the size of the destination community. Estimates are individually imprecise because each triennium supplies only 63-130 origin communities. Signs are nonetheless consistent with the pooled results.

Main results for the extensive margin 2015-2017

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00167	-0.0201*	-0.0335*	0.000670	-0.0517**	-0.0247
	(0.00260)	(0.00847)	(0.0165)	(0.0149)	(0.0160)	(0.0172)
N	9774	9774	9774	9774	9774	9774
N_clust	104	104	104	104	104	104
r2	0.0240	0.0641	0.113	0.0819	0.186	0.0842

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	-0.00213	-0.00156	-0.0170***	-0.0150***	-0.0119***	-0.00989**
	(0.00115)	(0.00230)	(0.00299)	(0.00441)	(0.00255)	(0.00353)
N	9763	9763	9763	9763	9763	9763
N_clust	2945	2945	2945	2945	2945	2945
r2	0.0352	0.0764	0.140	0.0995	0.203	0.106

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	-0.0190*	-0.0463	-0.0386	0.0798	0.0210	-0.0546
	(0.00835)	(0.0259)	(0.0278)	(0.0420)	(0.0319)	(0.0376)
Δldeg(i)	-0.00508*	-0.00873	-0.0230***	-0.00269	-0.00863	-0.0183**
	(0.00222)	(0.00479)	(0.00636)	(0.00885)	(0.00546)	(0.00664)
N	9763	9763	9763	9763	9763	9763
N_clust	2945	2945	2945	2945	2945	2945
r2	0.0362	0.0767	0.140	0.100	0.203	0.106

Notes: Control variables included in all models. All outcome variables are binary indicators of any spending. In Panel A the treatment is standardised, so coefficients are per standard deviation of $\Delta C_D(c)$; in Panels B and C they are semi-elasticities. Standard errors (in parentheses) are clustered three-way at the origin community, destination community and patient level in Panel A, and two-way at the provider and patient level in Panels B and C; stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Main results for the intensive margin 2015-2017

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	0.0432	-0.0646	-0.0714	-0.0120
	(0.0559)	(0.0448)	(0.0618)	(0.0386)
N	9568	8633	7794	2129
N_clust	102	103	101	79
r2	0.287	0.443	0.157	0.120

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-0.0411**	-0.0379**	-0.0593***	-0.0212
	(0.0142)	(0.0118)	(0.0133)	(0.0137)
N	9557	8622	7787	2119
N_clust	2910	2828	2666	1257
r2	0.308	0.454	0.189	0.151

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	-0.0999	-0.0779	-0.261*	-0.176
	(0.110)	(0.0998)	(0.129)	(0.123)
Δldeg(i)	-0.0566*	-0.0500*	-0.0997***	-0.0496*
	(0.0253)	(0.0211)	(0.0267)	(0.0239)
N	9557	8622	7787	2119
N_clust	2910	2828	2666	1257
r2	0.308	0.454	0.190	0.152

Notes: Control variables included in all models. All outcome variables are log-transformed spending variables. In Panel A the treatment is standardised, so coefficients are per standard deviation of $\Delta C_D(c)$; in Panels B and C they are elasticities. Standard errors (in parentheses) are clustered three-way at the origin community, destination community and patient level in Panel A, and two-way at the provider and patient level in Panels B and C; stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Main results for the extensive margin 2017-2019

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00412	-0.0420*	-0.0172*	-0.0232*	-0.0286	-0.0323
	(0.00241)	(0.0174)	(0.00690)	(0.0114)	(0.0151)	(.)
N	10537	10537	10537	10537	10537	10537
N_clust	78	78	78	78	78	78
r2	0.0192	0.0431	0.0883	0.0828	0.180	0.0714

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	-0.00159	-0.00789**	-0.0118***	-0.0125*	-0.0153***	-0.00906*
	(0.000925)	(0.00244)	(0.00224)	(0.00492)	(0.00239)	(0.00374)
N	10534	10534	10534	10534	10534	10534
N_clust	2914	2914	2914	2914	2914	2914
r2	0.0262	0.0551	0.103	0.103	0.195	0.0906

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	-0.0149*	-0.0364	-0.0271	0.0405	0.0245	-0.0802*
	(0.00656)	(0.0254)	(0.0281)	(0.0391)	(0.0299)	(0.0371)
Δldeg(i)	-0.00382*	-0.0133**	-0.0158**	-0.00648	-0.0116*	-0.0210**
	(0.00173)	(0.00499)	(0.00575)	(0.00956)	(0.00500)	(0.00703)
N	10534	10534	10534	10534	10534	10534
N_clust	2914	2914	2914	2914	2914	2914
r2	0.0269	0.0553	0.103	0.103	0.195	0.0912

Notes: As Table above for 2015-2017. The missing standard error in column (6) of Panel A is a failure of the three-way cluster-robust variance estimator with 78 origin communities, not a zero.

Main results for the intensive margin 2017-2019

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	-0.00273	-0.138**	-0.145**	-0.0793
	(0.0361)	(0.0408)	(0.0527)	(0.0405)
N	10350	9463	8423	2444
N_clust	77	76	73	63
r2	0.283	0.434	0.144	0.137

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-0.0291*	-0.0404***	-0.0786***	-0.0180
	(0.0135)	(0.00946)	(0.00961)	(0.00938)
N	10347	9458	8419	2437
N_clust	2891	2782	2643	1319
r2	0.298	0.444	0.176	0.163

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	-0.164	-0.219*	0.00563	-0.0240
	(0.139)	(0.108)	(0.0956)	(0.122)
Δldeg(i)	-0.0536	-0.0731***	-0.0778***	-0.0217
	(0.0311)	(0.0204)	(0.0188)	(0.0202)
N	10347	9458	8419	2437
N_clust	2891	2782	2643	1319
r2	0.298	0.444	0.176	0.163

Notes: As Table above for 2015-2017.

Main results for the extensive margin 2022-2024

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00924**	-0.0105	-0.0535**	-0.0663**	-0.00805	-0.0498*
	(0.00303)	(0.0119)	(0.0196)	(0.0197)	(0.0138)	(0.0198)
N	16393	16393	16393	16393	16393	16393
N_clust	130	130	130	130	130	130
r2	0.0310	0.0774	0.111	0.0845	0.203	0.0806

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	-0.00329**	-0.00164	-0.0185***	-0.0176***	-0.0132***	-0.0115***
	(0.00108)	(0.00190)	(0.00286)	(0.00319)	(0.00207)	(0.00238)
N	16377	16377	16377	16377	16377	16377
N_clust	3368	3368	3368	3368	3368	3368
r2	0.0451	0.0903	0.145	0.107	0.220	0.103

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	-0.0299***	-0.00531	-0.103***	0.00153	-0.0420	-0.0347
	(0.00721)	(0.0231)	(0.0249)	(0.0403)	(0.0254)	(0.0292)
Δldeg(i)	-0.00824***	-0.00252	-0.0355***	-0.0174*	-0.0201***	-0.0172**
	(0.00200)	(0.00461)	(0.00608)	(0.00867)	(0.00460)	(0.00546)
N	16377	16377	16377	16377	16377	16377
N_clust	3368	3368	3368	3368	3368	3368
r2	0.0466	0.0903	0.146	0.107	0.220	0.104

Notes: As Table above for 2015-2017.

Main results for the intensive margin 2022-2024

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	-0.0906	-0.102	-0.206**	-0.0426
	(0.0653)	(0.0608)	(0.0648)	(0.0542)
N	16054	14399	12900	3709
N_clust	127	127	125	99
r2	0.304	0.482	0.171	0.163

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-0.0371**	-0.0496***	-0.0777***	-0.0244**
	(0.0140)	(0.00863)	(0.00893)	(0.00935)
N	16039	14381	12883	3687
N_clust	3340	3273	3115	1760
r2	0.324	0.495	0.204	0.198

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	-0.253	-0.257**	-0.283*	-0.0499
	(0.132)	(0.0891)	(0.118)	(0.110)
Δldeg(i)	-0.0791*	-0.0926***	-0.125***	-0.0330
	(0.0325)	(0.0184)	(0.0242)	(0.0210)
N	16039	14381	12883	3687
N_clust	3340	3273	3115	1760
r2	0.325	0.495	0.205	0.198

Notes: As Table above for 2015-2017.

C Robustness checks

C.1 Baseline balance test: association between t_3 metrics and pre-move outcomes

Each cell regresses the pre-move (t_2) value of the outcome on the future treatment, using the covariates and fixed effects of the corresponding estimating equation. A coefficient of the same sign as the estimated effect indicates selection that could generate the result; a coefficient of the opposite sign indicates selection that works against it. Statistical significance should be interpreted alongside the size of the pre-move association relative to the estimated effect.

Balance test on the extensive margin

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	0.00317	0.0120	0.00423	0.00979	0.00218	-0.0200*
	(0.00185)	(0.00629)	(0.00413)	(0.00690)	(0.00763)	(0.00870)
N	36704	36704	36704	36704	36704	36704
N_clust	312	312	312	312	312	312
r2	0.0212	0.0349	0.0804	0.0576	0.0788	0.0658

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	0.00297***	-0.00353**	0.0134***	0.00344**	0.00405*	0.00335
	(0.000374)	(0.00120)	(0.00143)	(0.00117)	(0.00179)	(0.00172)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0315	0.0473	0.0934	0.0677	0.0894	0.0767

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	0.0255***	-0.00452	0.0181	-0.0541**	-0.0172	0.0142
	(0.00446)	(0.0135)	(0.0151)	(0.0166)	(0.0179)	(0.0192)
Δldeg(i)	0.00698***	-0.00425	0.0162***	-0.00507	0.00135	0.00559
	(0.000992)	(0.00247)	(0.00342)	(0.00280)	(0.00369)	(0.00363)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0332	0.0473	0.0935	0.0680	0.0895	0.0767

Notes: Dependent variables are the t_2 values of the extensive-margin indicators. Controls and fixed effects as in the corresponding panels of Table 4, except that lagged length of stay and lagged severity are excluded from the inpatient column. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Balance test on the intensive margin

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	-1.695	-36.79	-5.492	2.135
	(42.26)	(41.92)	(7.231)	(10.18)
N	36704	36704	36704	36704
N_clust	312	312	312	312
r2	0.200	0.123	0.0992	0.102

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-21.98*	-19.50*	0.871	1.768
	(9.349)	(7.821)	(1.750)	(2.106)
N	36674	36674	36674	36674
N_clust	5916	5916	5916	5916
r2	0.212	0.136	0.112	0.115

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	-21.21	-68.66	4.110	-47.72*
	(104.7)	(84.20)	(22.41)	(22.83)
Δldeg(i)	-25.32	-30.31*	1.518	-5.743
	(20.35)	(15.23)	(4.362)	(4.383)
N	36674	36674	36674	36674
N_clust	5916	5916	5916	5916
r2	0.212	0.136	0.112	0.115

Notes: Dependent variables are the t_2 values of spending in Swiss francs, not log-transformed, so coefficients are in CHF and are not comparable in magnitude to the log-point estimates in Table 5; only their sign and significance should be read. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.2 Models estimated in First Differences (FD) instead of ANCOVA

First-difference estimates for the extensive margin

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00799**	-0.0276**	-0.0391***	-0.0395**	-0.0296***	-0.0188
	(0.00286)	(0.0106)	(0.0110)	(0.0123)	(0.00790)	(0.0111)
N	36704	36704	36704	36704	36704	36704
N_clust	312	312	312	312	312	312
r2	0.0171	0.290	0.0332	0.0207	0.0189	0.0212

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	-0.00525***	-0.000736	-0.0275***	-0.0185***	-0.0163***	-0.0133***
	(0.00100)	(0.00161)	(0.00189)	(0.00339)	(0.00154)	(0.00209)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0319	0.299	0.0624	0.0383	0.0350	0.0366

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	-0.0454***	-0.0270	-0.0725***	0.0837**	0.00675	-0.0671**
	(0.00642)	(0.0174)	(0.0210)	(0.0317)	(0.0181)	(0.0255)
Δldeg(i)	-0.0124***	-0.00499	-0.0389***	-0.00534	-0.0153***	-0.0238***
	(0.00179)	(0.00362)	(0.00451)	(0.00740)	(0.00303)	(0.00447)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0346	0.299	0.0630	0.0388	0.0350	0.0369

Notes: The dependent variable is the change in the outcome between t_2 and t_3; the lagged outcome is dropped from the right-hand side. Other controls and fixed effects as in Table 4. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

First-difference estimates for the intensive margin

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	-0.00344	-0.0587*	-0.130***	-0.0601
	(0.0289)	(0.0237)	(0.0358)	(0.0534)
N	35524	29873	24280	4184
N_clust	305	301	294	209
r2	0.0265	0.0202	0.0335	0.0657

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-0.0373***	-0.0520***	-0.102***	-0.0415***
	(0.0107)	(0.00571)	(0.00566)	(0.0107)
N	35493	29845	24252	4126
N_clust	5860	5595	5202	2042
r2	0.0419	0.0368	0.0668	0.119

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	-0.163	-0.236***	-0.236***	-0.0370
	(0.0868)	(0.0505)	(0.0708)	(0.134)
Δldeg(i)	-0.0631**	-0.0897***	-0.139***	-0.0477
	(0.0218)	(0.00989)	(0.0121)	(0.0254)
N	35493	29845	24252	4126
N_clust	5860	5595	5202	2042
r2	0.0422	0.0375	0.0673	0.119

Notes: As above. The samples are smaller than in Table 5 because the first difference of a log outcome requires positive spending in both years. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.3 ANCOVA without controls

Estimates without demographic, severity and insurance controls, extensive margin

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C\_D(c)$, SD units	-0.00727***	-0.0300***	-0.0607***	-0.0517***	-0.0549***	-0.0571***
	(0.00132)	(0.00414)	(0.00868)	(0.00759)	(0.00541)	(0.00688)
N	36704	36704	36704	36704	36704	36704
N_clust	312	312	312	312	312	312
r2	0.0174	0.0435	0.0639	0.0534	0.167	0.0630

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Δldeg(i)	-0.00235*	-0.00302*	-0.0155***	-0.0155***	-0.0127***	-0.0115***
	(0.00104)	(0.00138)	(0.00279)	(0.00383)	(0.00166)	(0.00245)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0307	0.0580	0.0970	0.0752	0.187	0.0866

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
$\Delta C(i)$	-0.0222***	-0.0315*	-0.0621*	0.0391	-0.00945	-0.0558*
	(0.00609)	(0.0156)	(0.0266)	(0.0355)	(0.0208)	(0.0251)
Δldeg(i)	-0.00585**	-0.00798**	-0.0253***	-0.00929	-0.0141***	-0.0203***
	(0.00191)	(0.00307)	(0.00656)	(0.00865)	(0.00395)	(0.00496)
N	36674	36674	36674	36674	36674	36674
N_clust	5916	5916	5916	5916	5916	5916
r2	0.0318	0.0582	0.0977	0.0753	0.187	0.0869

Notes: All patient demographics, severity measures and insurance plan controls are dropped; the lagged outcome, the community size terms and the fixed effects are retained, so the comparison with Table 4 isolates the contribution of the covariate adjustment. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Estimates without demographic, severity and insurance controls, intensive margin

Panel A) Results for $\Delta C_D(c)$, community degree centralisation

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C\_D(c)$, SD units	-0.101***	-0.210***	-0.206***	-0.0660***
	(0.0284)	(0.0240)	(0.0316)	(0.0186)
N	35972	32495	29117	8284
N_clust	306	306	299	241
r2	0.222	0.349	0.134	0.112

Panel B) Results for Δldeg(i), log degree of main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Δldeg(i)	-0.0274	-0.0313**	-0.0727***	-0.0170**
	(0.0155)	(0.0108)	(0.00976)	(0.00603)
N	35943	32461	29089	8246
N_clust	5878	5769	5552	3292
r2	0.244	0.366	0.169	0.147

Panel C) Results for $\Delta C(i)$, local clustering around main PCP

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
$\Delta C(i)$	-0.198	-0.221*	-0.193	-0.0790
	(0.131)	(0.101)	(0.0992)	(0.0704)
Δldeg(i)	-0.0586	-0.0664**	-0.103***	-0.0300*
	(0.0341)	(0.0237)	(0.0225)	(0.0126)
N	35943	32461	29089	8246
N_clust	5878	5769	5552	3292
r2	0.244	0.366	0.170	0.147

Notes: As above. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.4 Directional specification

Equation (11) replaces the continuous treatment with indicators for moves up and down the national quintile distribution of $C_D(c)$, with lateral movers as the reference group. Where significant, the coefficients agree in sign with the continuous specification; elsewhere they are imprecise, and on the intensive margin they are null throughout. Two features of the data explain the loss of precision. Binning a continuous measure into quintiles discards the magnitude of each move, and 10-13% of valid communities have $C_D(c)$ exactly equal to zero (appendix A.4), which leaves the lowest quintile boundary poorly identified.

Panel A) Extensive margin

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Mover-down	0.00162	0.00494	0.00334	-0.00112	0.0101	0.00379
	(0.00199)	(0.00707)	(0.00834)	(0.00969)	(0.00882)	(0.00866)
Mover-up	-0.00241	0.00102	-0.0145	-0.0228*	-0.00737	-0.0319**
	(0.00253)	(0.00997)	(0.0100)	(0.0115)	(0.0123)	(0.0119)
N	36704	36704	36704	36704	36704	36704
N_clust	312	312	312	312	312	312
r2	0.0230	0.0604	0.102	0.0798	0.189	0.0779

Panel B) Intensive margin

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Mover-down	-0.0481	-0.0389	-0.0180	0.0240
	(0.0291)	(0.0283)	(0.0352)	(0.0259)
Mover-up	-0.0136	-0.0329	-0.0534	0.0382
	(0.0340)	(0.0433)	(0.0455)	(0.0380)
N	35972	32495	29117	8284
N_clust	306	306	299	241
r2	0.292	0.454	0.156	0.138

Notes: Coefficients are relative to lateral movers, i.e. patients moving between communities in the same national quintile of $C_D(c)$. Controls, fixed effects and clustering as in Panel A of Tables 4 and 5. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.5 Practice setting: transitions between solo and networked practice

The $H_c \geq 3$ restriction excludes patients whose main provider practises in a community too small for $C_D(c)$ to be defined. Those patients are not uninformative: the transition between solo and networked practice is a coarse but well-defined change in connectedness, and it is available for exactly the observations the main treatment drops. The estimates below use an indicator for the direction of that transition, with patients who remain in a networked community as the reference group.

Moving from a networked practice to a solo one raises conditional spending on prescription drugs by 11.5% and on outpatient care by 5.3%. The corresponding laboratory estimate (+11.8%) is significant. Transitions in the opposite direction, from solo to networked practice, are null throughout and very imprecisely estimated, since few patients make that move. The interpretable estimates point the same way as the main results: less connectedness, more spending.

Panel A) Extensive margin

	(1) Outpat.	(2) Inpat.	(3) Drugs	(4) Lab	(5) Physio	(6) Emerg. care
Solo to networked	0.00124	0.0127	0.00343	-0.00589	0.0563	-0.00167
	(0.000958)	(0.0497)	(0.0273)	(0.0466)	(0.0670)	(0.0712)
Networked to solo	-0.000963	-0.00590	0.0146*	0.0166	0.0253**	-0.0172
	(0.00143)	(0.00829)	(0.00701)	(0.0113)	(0.00946)	(0.0113)
N	41717	41717	41717	41717	41717	41717
N_clust	701	701	701	701	701	701
r2	0.0254	0.0691	0.100	0.0824	0.202	0.0838

Panel B) Intensive margin

	(1) Outpat.	(2) Drugs	(3) Lab	(4) Physio
Solo to networked	0.0288	0.132	-0.267*	0.0594
	(0.179)	(0.142)	(0.134)	(0.246)
Networked to solo	0.0529*	0.115***	0.118***	0.0248
	(0.0260)	(0.0310)	(0.0315)	(0.0284)
N	40859	37152	33277	9689
N_clust	682	668	616	362
r2	0.303	0.468	0.161	0.149

Notes: The reference group is patients who move between two networked communities. Controls, fixed effects and clustering as in Panel A of Tables 4 and 5. Dependent variables in Panels C and D are t_2 values, with spending in Swiss francs rather than logs. Stars indicate significance as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$