Rationed Choice: Capacity Constraints and Patient Demand for Primary Care

Temulun Borjigen*

June 2026

Abstract

This paper studies how capacity constraints distort observed patient choices and bias the estimation of primary care demand. In many primary care systems, providers can close their lists when demand exceeds regulated capacity, thus the set of providers a patient can join differs from the set operating in the market. I develop a discrete-choice framework that separates patient preferences from feasible access, using variation from forced switches after GP retirements and practice closures, together with changes in providers' open-list status over time. I complement the model with machine-learning methods to capture heterogeneity in welfare loss due to patients' switch motivations and characteristics. Using Danish administrative data, I compare estimates from this constrained-choice model to a standard specification that treats choices as unconstrained. Ignoring capacity constraints confounds preferences with rationing, understating the role of distance and provider characteristics, and distorting substitution patterns used in counterfactual analysis. I then use the corrected estimates to study the welfare and distributional effects of changes in list-size regulation and provider entry in underserved areas.

*Department Firms and Markets, DIW Berlin; Department of Economics, University of Copenhagen; and Berlin School of Economics. tborjigen@diw.de

1 Introduction

Patient choice is a central feature of many healthcare systems. Allowing patients to select their provider is intended to strengthen provider incentives by making demand responsive to differences in quality, access, and convenience (Gaynor et al. 2016; Santos et al. 2017). In primary care, these mechanisms are especially important because patients are usually attached to a single provider over time, and the chosen practice acts as the first point of contact with the healthcare system. Yet the value of patient choice depends not only on whether patients are formally allowed to choose, but also on whether the providers they prefer are actually available.

This distinction matters in list-based primary care systems. Even when care is free at the point of use and patients have a formal right to choose their general practitioner, providers may close their lists once they reach capacity. In that case, patients can choose only among practices that are accepting new patients. Observed choices therefore reflect both patient preferences and supply-side rationing. A patient who registers with a more distant or less preferred practice may not be revealing weak preferences for proximity or provider quality. She may simply be unable to access the practice she would otherwise choose. As a result, capacity constraints can create a wedge between patients' preferred and realized providers, reducing the surplus they derive from provider choice and potentially generating substantial welfare losses.

Estimating this welfare loss requires recovering patient preferences separately from the constraints that determine the feasible choice set. Standard revealed-preference models are not sufficient for this task when they treat all geographically relevant providers as available alternatives. If nearby or high-utility practices are closed to new patients, observed choices are second-best choices from a restricted set rather than first-best expressions of demand. Treating these constrained choices as unconstrained revealed preferences confounds tastes with rationing, which can bias estimated preference parameters, implied substitution patterns, and the welfare effects of policies that expand access.

The conceptual issue mirrors a broader challenge in discrete choice analysis under constrained or unobservable choice sets. A central concern in this literature is that observed choices may not reveal preferences over the full set of alternatives when the econometrician does not observe which options were feasible or considered. For instance, in school-choice settings, submitted rankings may combine true preferences with strategic responses to assignment mechanisms and capacity constraints (Agarwal and Somaini 2018; Fack et al. 2019). More generally, discrete-choice models with heterogeneous choice sets show that ignoring variation in the alternatives available to decision makers can bias preference estimates and substitution patterns (Barseghyan et al. 2021; Crawford et al. 2021). In healthcare, similar issues arise when provider acceptance decisions or other supply-side restrictions are latent, requiring instruments or structural assumptions to recover the feasible choice set (Agarwal and Somaini 2025).

This paper estimates the welfare effect of capacity constraints in the Danish primary-care market. The setting studied here differs from the constrained choice literature because the key constraint is directly observed. Primary care practices in Denmark are permitted to close their lists to new enrollments when list size reaches a regulated threshold (Pedersen et al. 2012; Nørøxe et al. 2018; Haastrup et al. 2025), and Danish administrative records disclose whether each GP practice is open or closed to new patients in each period. This allows me to construct feasible choice sets directly and to compare demand estimates and welfare measures under constrained and unconstrained choice-set definitions. The main welfare object is the loss in patient surplus from being unable to choose from the full set of relevant providers because some clinics have closed lists. I then examine how ignoring these constraints changes estimated preferences and the predicted welfare gains from counterfactual capacity expansions.

Besides the observable list status, the Danish primary care market provides additional features that support the empirical analysis. First, patients face only nominal monetary switching costs and Danish administrative registers link each individual to her registered GP, residential address, and health history. These data allow both the chosen provider and the relevant set of potential providers to be measured at the population level. Second, the institutional environment generates variation that is central to the identification strategy. When a practice closes, displaced patients must select a new provider within a specified window or are administratively reassigned to a replacement practice by the regional health authority. A subset of patients from closed clinics actively choose a new clinic within that window at the time of closure. These active choice occasions identify preferences over providers because the timing of the choice is induced by the closure of the previous practice, not by the patient's endogenous decision to switch providers. This reduces selection on unobserved dissatisfaction with the incumbent provider or on patient-specific search motives. Conditional on the observed feasible choice set, the clinic selected during the closure window can therefore be interpreted as revealing preferences over available alternatives.

Building on this variation, I estimate a discrete-choice model of primary care clinic choice to recover patient preferences separately from the capacity constraints that shape feasible access. I estimate the model for switcher groups that differ in the extent to which the timing of choice is induced by institutional events rather than by endogenous switching motives. For each group, I compare demand estimates under two choice-set definitions: one that treats all geographically relevant clinics as available, and one that restricts the choice set to clinics that are open to new patients at the time of choice. This comparison quantifies how much preference estimates and welfare calculations are distorted when constrained choices are treated as unconstrained revealed preferences.

The analysis delivers three main results. First, capacity constraints impose a quantitatively important welfare burden. Closed lists frequently remove high-utility alternatives from patients' feasible choice sets. In my sample, 65 percent of switch events occur when the patient's highest-utility clinic has a closed list. These constraints are not evenly distributed across local markets, implying that the welfare burden of rationing depends on where patients live and which clinics are accepting new patients at the time of choice. Second, ignoring capacity constraints biases revealed-preference estimates. A conditional logit model that treats all clinics as available understates the magnitude of patients' distance aversion. In the preferred active-at-closure specification, the unconstrained model understates distance aversion by roughly one-third relative to the full-clinic benchmark and by approximately 46 percent relative to the specification that uses the observed feasible choice set. This means that models ignoring closed lists make patients appear more willing to travel than they are when choices are evaluated against the providers they could actually access. The bias also changes implied substitution patterns, because unavailable high-utility clinics are incorrectly treated as feasible competitors. Third, these demand biases matter for welfare and policy counterfactuals. Using the corrected demand estimates, I quantify the welfare loss from capacity constraints and evaluate counterfactual expansions in provider availability. The counterfactuals show that models ignoring capacity constraints understate both the average welfare cost of rationing and the heterogeneity of this burden across patient groups and market conditions. The welfare gains from expanding capacity are therefore not captured correctly by models that treat observed allocations as if they were generated by unconstrained choice.

These results contribute most directly to work on barriers to patient access in healthcare markets. Besides decisions to stop accepting new patients, constraints on patient choice may arise from other supply-side behavior by providers. For example, in hospital markets, referral behavior and physician advice can determine the alternatives that patients consider or can access (Beckert 2018). In insurance markets, provider networks restrict the set of hospitals or physicians that patients can use on favorable terms, which affects provider demand, market power, and patient welfare (Ho 2006; Gowrisankaran et al. 2015). In contrast, other work studies policy reforms that aim to expand patient choices through mandating a minimum number of choices and lowering provider entry barriers(Gaynor et al. 2016; Dietrichson et al. 2020). This paper studies a related barrier to access in primary care, capacity constraints that arise when GPs close their patient lists to new registrations, and contributes to the welfare discussion by quantifying how these constraints affect patient surplus.

The paper also adds to the literature estimating patient preferences over healthcare providers. A large body of work shows that patients respond to provider attributes such as distance, quality, waiting time, and physician characteristics when choosing healthcare providers. Geographic proximity is one of the strongest predictors of provider choice in both hospital and primary-care settings (Sivey 2012; Gowrisankaran et al. 2015; Raval et al. 2022). Distance is also central to welfare analysis because, in healthcare markets where prices are limited or absent, the disutility of distance provides a metric for valuing access and for comparing proximity against other provider attributes such as quality (Romley and Goldman 2011; Chandra et al. 2016). Patients also respond to quality and physician characteristics. In primary care, patients are more likely to choose higher-quality family doctors (Santos et al. 2017; Brown et al. 2023), while in hospital markets mortality, readmission rates, and waiting times affect patient demand (Beckert et al. 2012; Gutacker et al. 2016). Physician characteristics matter as well, including gender and age similarity between patients and doctors (Godager 2012). This paper builds on this revealed-preference approach, but shows that the interpretation of observed choices depends on whether preferred providers are actually available. When high-utility providers are closed to new patients, observed distance choices and estimated trade-offs between provider attributes reflect both preferences and rationing.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and the registry data. Section 3 documents stylized facts on the incidence and distribution of capacity constraints. Section 4 develops the discrete-choice framework, presents conditional logit estimates across choice sets and switcher groups, and quantifies the bias from ignoring constraints. Section 5 translates the preferred estimates into welfare measures, examines distributional heterogeneity across switcher types and market conditions, and supplements the structural analysis with OLS regressions and causal machine learning. Section 6 evaluates two counterfactual capacity expansions. Section 7 concludes.

2 Institutional Setting and Data

2.1 Danish Primary Care: Institutional Background

General practice (GP) occupies a central role in Danish primary care. The system is organized around patient lists, with nearly all residents registered with a specific general practice. The GP serves as the patient's first point of contact and acts as gatekeeper for most specialist and hospital care (Simonsen et al. 2021; Haastrup et al. 2025). Patients are generally listed with a practice located within a regulated distance from their residence, typically within 15 kilometers, and within 5 kilometers in densely populated areas such as central Copenhagen. GP services are tax-financed and free at the point of use for patients in the main public insurance group. As a result, patients' provider choices are not mediated by consultation-level prices.

General practice is governed through collective agreements between the Danish Regions and the Organization of General Practitioners (*Praktiserende Lægers Organisation*, PLO), which is part of the Danish Medical Association. These agreements regulate remuneration, opening hours, service obligations, and rules concerning practice capacity (Haastrup et al. 2025). GPs are self-employed contractors and are reimbursed through a mixed payment system combining capitation and fee-for-service payments. Capitation provides a payment for each listed patient, while fee-for-service payments remunerate specific consultations and procedures.

A key institutional feature for this paper is that patient lists are capacity constrained. Danish general practice operates with a list norm of approximately 1,600 patients per full-time GP, although the exact capacity depends on practice type, staffing, and the relevant collective agreement. Practices may close their lists to new patients when their list size reaches the applicable capacity threshold. Conversely, practices with open lists can accept new registrants. List status is recorded administratively and varies across practices and over time. In the empirical analysis, I use this recorded open-list status to construct the set of providers that is feasible for each patient at the time of choice.

Patients can change their registered GP, subject to the requirement that the new practice is accepting new patients. A standard voluntary switch requires a small administrative fee, DKK 150 (about USD 20)[1], while switching is free in institutional cases such as practice closure, prac-

[1] The fee may be lower in some regions. For example, in Copenhagen and Aarhus, the fee is DKK 40 (roughly USD 6). See https://international.kk.dk/live/healthcare/going-to-a-doctor/changing-your-doctor and https://international.aarhus.dk/live/health/change-of-general-practitioner-family-doctor#what-does-it-cost-to-change-to-a-new-general-practitioner-family-doctor--91

tice restructuring, or residential moves. These rules imply that the main constraint on switching is not the consultation price, but whether the desired practice is open to new registrants. Publicly available information about practices is also limited, consisting mainly of the physician's name and the practice address. This makes selective switching based on detailed provider characteristics less likely than in settings where quality, practice style, or patient-experience measures are widely available (Huang and Ullrich 2024). A patient who wants to move to a closed-list practice cannot generally do so unless an exception applies. She must instead choose among open-list practices or remain with her incumbent provider.

When a GP practice closes, listed patients receive a letter from the regional health authority specifying a period during which they can choose a new GP without paying the usual administrative fee. Patients who make a choice during this window are recorded as active re-registrations after closure. If a patient does not choose a new GP within the stated period, the municipality assigns the patient to a new doctor and the patient receives a new health insurance card. A related but distinct case arises when a GP transfers the practice to a successor. In that case, patients are automatically assigned to the physician who takes over the practice. Patients who do not wish to remain with the successor may request a free change of GP within the period specified by the regional health authority. These rules create three empirically distinct closure-related transitions: automatic transfer to a successor, administrative reassignment after closure, and active choice during the closure window.

This institutional setting creates a direct link between capacity constraints and observed patient allocations. In an unconstrained revealed-preference setting, the observed GP registration would be interpreted as the patient's preferred alternative among the relevant set of providers. In the Danish list system, however, the observed allocation reflects both patient preferences and the supply-side availability of practice lists. The empirical implication is that the relevant choice set is the set of practices accepting new patients, not the full set of geographically proximate practices. Because list status is observed in the administrative data, I can distinguish providers that are physically available in the market from providers that are feasible alternatives for new registrants.

2.2 Data Sources and Sample Construction

Clinics. I identify primary care clinics using authorization numbers from the Danish Clinic Registry (*Yderregisteret*) and link these identifiers to claims records from the Danish National Health Insurance Service Registry (*Sygesikringsregistret*). The Clinic Registry contains clinic-level information, including opening and closing dates, postcode, and the unique personal identifiers (*personnummer*, PNR) of physicians and other healthcare workers affiliated with each clinic. [2] The Claims Registry records the clinic authorization number, claim date, patient identifier, and type of medical service provided for each reimbursed service. Services are classified using six-digit codes, which allow me to identify claims corresponding to services provided in the primary care sector.

[2] The PNR is the unique pseudonymized personal identifier assigned to residents of Denmark.

Table 2.1 reports summary statistics for the clinic-quarter panel. The data contain 38,481 clinic-quarter observations over 2014 - 2019, with an average of 1,749 clinics observed per quarter. Only 27 percent of clinic-quarter observations have an open list, indicating that list closures are common during the sample period. This feature is central for the empirical design, since it implies that the set of feasible providers differs substantially from the full set of active clinics. The average clinic has two physicians and 3,122 listed patients, corresponding to 1,627 listed patients per physician. The latter figure is close to the institutional list norm of 1,600 patients per full-time GP, which is consistent with capacity constraints being relevant in practice. Clinics are heterogeneous in both size and activity: the standard deviation of total listed patients is 1,956, and the standard deviation of newly listed patients is 239 relative to a mean of 159. On average, clinics record 8,026 claims and 2,112 visited patients per quarter. Physician composition also varies across clinics. The mean share of female physicians is 48 percent, the average physician age is 54 years, and the mean share of physicians with an immigration background is 8 percent. These characteristics enter the choice model as provider attributes and allow the analysis to distinguish capacity constraints from observable differences in clinic composition.

Table 2.1: Clinic summary statistics

	Mean	SD
Number of clinics	1749.14	72.07
Open status	0.27	0.44
Clinic size	2.00	1.29
Share of female physicians	0.48	0.42
Physician age	54.19	8.35
Physicians with immigration background	0.08	0.24
Number of claims	8026.00	5556.06
Number of total listed patients	3122.31	1956.11
Number of listed patients per physician	1626.94	461.44
Number of newly listed patients	159.08	239.03
Number of visited patients	2111.84	1433.42
Observations	38481	

Note: This table summarizes the mean and the standard deviation of the listed variables from an unbalanced quarterly panel of active primary care clinics from 2014 to 2019. Open status (0/1) indicates whether a clinic is open to accepting new patients. Clinic size refers to the number of physicians in a clinic, excluding other supportive roles such as assistants, administrative staff, and interns. Physicians with immigration background include physicians who are first- or second-generation immigrants. Number of visited patients refers to the number of unique patients who visited the clinic, regardless of whether they were listed at that clinic.

Patients. I construct a patient-quarter panel for the years 2014 through 2019 by linking each resident to her registered GP in each quarter, following the patient-list registration records. They record the GP affiliation used for health insurance administration and capitation payments, and can be linked to the clinic authorization numbers in the provider and claims registers. The resulting panel contains one observation per patient per quarter and yields 134.6 million patient-quarter observations. A switch event is recorded whenever the patient's registered GP changes between two consecutive quarters. I classify each switch according to the institutional reason for the change: residential relocation, practice closure, or a voluntary switch unrelated to a closure. Closure-related switches include practice takeovers, administrative reassignments, and active re-registrations made during the closure window.

Table 2.2 summarizes the distribution of switch events by type. The sample contains 4.0 million switch events over 2014–2019. Residential moves account for 28 percent of switches, while clinic closures account for 32 percent. The remaining switches are changes that are not directly linked to either a residential move or a clinic closure. Closure-related switches are divided into three subgroups. Among all switch events, 15 percent occur through takeover after closure, 13 percent through administrative reassignment, and 3 percent through active choice during the closure window. These closure subgroups are central to the empirical design because they differ in the extent to which the timing and nature of the switch are driven by institutional events rather than by patient-initiated search. In particular, active choices after closure provide the cleanest switcher sample for estimating preferences: these patients are forced to make a new GP choice by the closure of their previous clinic, but they actively select their new provider from the available set.

Table 2.2: Switch events by type

	Share	Number of events
Moving address	0.28	1131393
Clinic closure	0.32	1278095
Takeover after closure	0.15	605032
Assigned after closure	0.13	539658
Active choice after closure	0.03	133391
Other reasons	0.40	1606408
Observations	4015896	

Note: The table reports the share and number of switch events by institutional reason. Closure-related subcategories are nested within the clinic-closure category and therefore do not represent additional mutually exclusive categories.

Table 2.3 compares patients who do not switch GP during the sample period with patients who switch at least once. The table is based on 5.7 million patient-level observations, of which 3.3 million are non-switchers and 2.4 million are switchers. Switchers are younger than non-switchers at first observation, with a mean age of 35.0 compared with 43.5 among non-switchers. The gender composition is similar across the two groups, with women accounting for about half of patients in both samples. Patients with an immigration background are more represented among switchers: they account for 15 percent of switchers and 9 percent of non-switchers. Switchers also have slightly fewer claims over the sample period, with an average of 57.4 claims compared with 61.7 among non-switchers. Among patients who switch, the average number of switches is 1.56, indicating that most switchers change GP once, while a smaller share switch multiple times. The average age at switching is 37.6 years. These differences indicate that switching is not randomly distributed in the population. Switchers are younger and more likely to have an immigration background than patients who remain with the same GP. The choice-model analysis therefore conditions on patient characteristics and compares switcher groups separately, rather than treating all observed GP registrations as arising from the same choice process.

Table 2.3: Patient summary statistics

	(1) Non-switchers	(2) Switchers
Age first observed	43.45	35.02
	(23.76)	(22.27)
Share of female patients	0.51	0.50
	(0.50)	(0.50)
Share of patients with immigration background	0.09	0.15
	(0.29)	(0.36)
Number of claims	61.73	57.42
	(62.91)	(59.35)
Number of switches		1.56
		(0.95)
Age when switching		37.64
		(22.52)
Observations	3312982	2415763

Note: This table summarizes the mean and the standard deviation of patient characteristics for 5,728,745 patients from 2014 to 2019. Non-switchers are patients who are listed with only one GP during the sample period. Switchers are patients who change their listed GP at least once. Patients with immigration background include patients who are first- or second-generation immigrants.

Choices. For each switch event, I define a choice occasion at the time of the change in registration. The patient's feasible choice set consists of all general practices within the relevant distance radius whose lists are open in the quarter of the switch. The radius is 15 kilometers in general and 5 kilometers in Copenhagen. Patient-to-GP distances are computed as distances between the patient's residential address coordinates and the GP's clinic address coordinates.[3] The chosen alternative is the practice with which the patient is registered after the switch. I keep choice occasions where at least one feasible alternative is observed. Cases in which patients are administratively reassigned after a practice closure are not treated as standard active choices. Instead, I retain them as a distinct switcher group. This distinction is important because administrative reassignment reflects the regional allocation process, whereas active re-registration during the closure window reflects a patient's choice among feasible providers.

Table 2.4 describes the choice environments faced by switchers. Panel A compares the number and distance of all clinics to the corresponding measures for clinics with open lists. Within 5 kilometers, patients have on average 60.6 clinics nearby, but only 9.1 of these clinics are open to new patients. Within 15 kilometers, the average number of nearby clinics rises

[3] Patient residences and clinic locations are geocoded in the WGS84 coordinate reference system. For each patient–clinic pair, I compute the great-circle distance between the two coordinate points using the Haversine formula.

to 107.3, while the average number of open clinics is only 17.8. Thus, the formal geographic market contains many more clinics than the feasible choice set available to a patient who must register with a new GP. The distance measures show the same pattern. Average distance to all clinics within the relevant radius is 2.5 kilometers within 5 kilometers and 6.2 kilometers within 15 kilometers. The corresponding average distance to open clinics is slightly larger, at 2.6 and 6.5 kilometers. This indicates that excluding closed-list clinics not only reduces the number of feasible alternatives, but also shifts the available set toward more distant providers. Panel B compares the distance to the chosen clinic with the distance to the nearest clinic and the nearest open clinic. Patients choose clinics that are on average 3.5 kilometers away, compared with 1.5 kilometers for the nearest clinic and 2.8 kilometers for the nearest open clinic. The gap between the nearest clinic and the nearest open clinic shows the role of capacity constraints in limiting access to geographically proximate providers. The remaining gap between the nearest open clinic and the chosen clinic indicates that patients do not choose only on distance, but may also respond to other clinic attributes, prior attachment, or institutional features of the switch occasion.

Table 2.4: Choice set composition at the time of switch events

Panel A: Patient choices	Within 5km	Within 10km	Within 15km
Number of all clinics	60.6	102.7	107.3
	(86.1)	(131.6)	(128.8)
Distance to all clinics	2.5	4.6	6.2
	(1.1)	(2.1)	(2.7)
Number of open clinics	9.1	16.5	17.8
	(14.5)	(23.7)	(23.2)
Distance to open clinics	2.6	4.7	6.5
	(1.2)	(2.3)	(3.0)
Panel B: Distance to chosen and nearest	Chosen	Nearest	Nearest open
Distance	3.5	1.5	2.8
	(3.4)	(2.0)	(3.1)

Note: The table reports the choice set characteristics for the 4,015,896 switch events described in Table 2.2. The distance unit is kilometers. Standard deviations are in parentheses.

2.3 Descriptive evidence on capacity constraints

Figure 2.1 provides descriptive evidence that list capacity is a binding constraint in the Danish primary care market. Panel (a) plots the number of active clinics and the share of clinics with closed lists over time. The number of clinics declines steadily during the sample period, from more than 1,800 clinics in the middle of the period to about 1,600 by the end of 2019. At the same time, the share of closed-list clinics remains high throughout the period, generally between

65 and 75 percent. This pattern shows that patients face a market in which many geographically proximate clinics are not available for new registration. Capacity constraints are therefore not confined to a small subset of clinics or periods, but are a persistent feature of the market.

Panel (b) relates open-list status to the number of newly listed patients. Clinics that are open to new patients receive substantially more new registrations than clinics that are less likely to be open. This positive correlation between open status and new registrants confirms that the administrative open-list indicator is meaningful for patient flows. Open-list status therefore captures an effective access margin rather than only a formal classification. Panel (c) shows that open-list status is systematically related to existing capacity pressure. The probability that a clinic is open declines with the lagged number of listed patients per physician. The estimated slope is negative, with a coefficient of -0.375 relative to an outcome mean of 0.278. Clinics with larger patient lists per physician are therefore less likely to accept new patients in the following period. This is consistent with the institutional rule that practices may close their lists when their patient load approaches the regulated capacity threshold.

Taken together, the three panels establish the empirical relevance of capacity constraints for patient choice. The evidence shows that limited capacity is both pervasive and consequential. During the sample period from 2014 to 2019, many clinics are unavailable to new patients, whereas open-list status strongly predicts patient inflows, and list availability responds to underlying patient load. These patterns support the central premise of the empirical analysis that patients choose among the subset of clinics that are actually accepting new patients. Ignoring list status would therefore treat unavailable clinics as feasible alternatives and confound patient preferences with supply-side rationing.

3 Model, Identification, and Estimation

3.1 Preferences and Utility

Building on the descriptive evidence in the previous section, I model GP switching as a discrete-choice problem in which patients select among feasible practices. The framework follows the standard random-utility approach widely used in the healthcare choice and differentiated-products literatures, but differs from most applications by explicitly incorporating administratively observed capacity constraints through the construction of the feasible choice set. This allows preferences to be estimated under alternative assumptions about provider availability and forms the basis for the empirical comparisons that follow. For each switch event, patient i chooses among a set of feasible general practices. The indirect utility from registering with practice j in period t is

$$U_{ijt} = \beta_d^i d_{ij} + \beta_{d2} d_{ij}^2 + \boldsymbol{\gamma}' \mathbf{X}_{ijt} + \varepsilon_{ijt}, \tag{1}$$

where d_{ij} is the Euclidean distance in kilometers between patient i's residential address and practice j's address. The vector $\mathbf{X}_{ijt}$ contains observed practice and patient-practice characteristics, including clinic size, physician age, the share of female physicians, the share of physicians with an immigration background, an indicator for gender match, and an indicator for immigration-

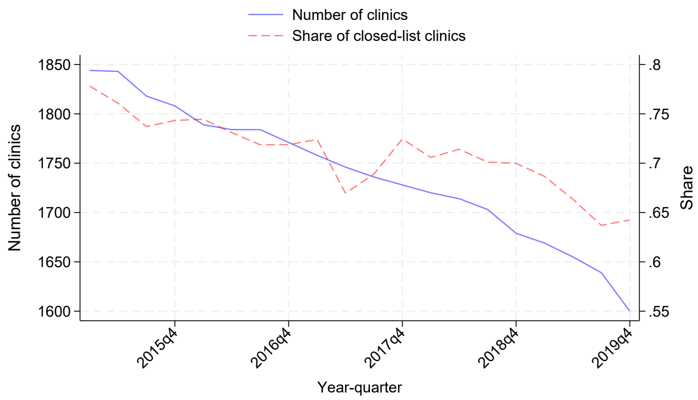

(a) Number of clinics and share of closed-list clinics

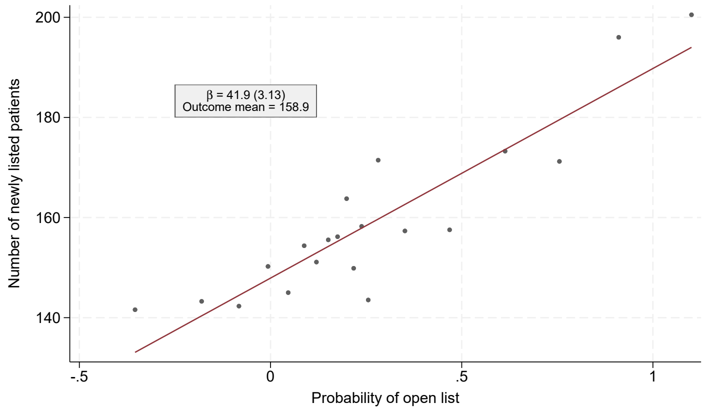

(b) Newly listed patients and open-list status

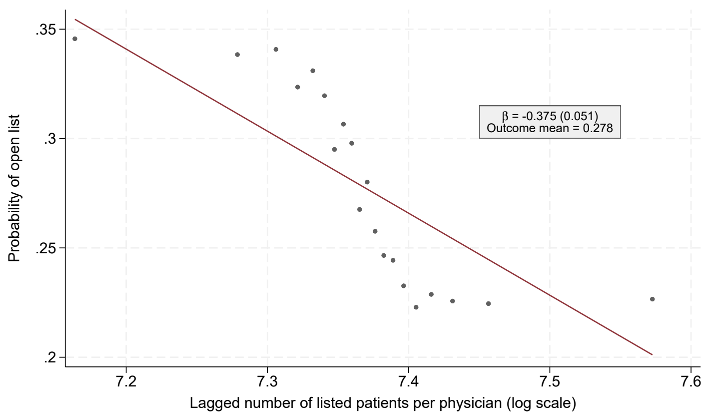

(c) Open-list status and lagged patient load

Figure 2.1: Descriptive evidence on capacity constraints in Danish primary care

Notes: Panel (a) plots the number of clinics and the share of closed-list clinics over time. Panels (b) and (c) present binned scatter plots based on clinic-quarter observations. Variables are residualized with respect to clinic and year-quarter fixed effects before binning. Each dot represents the mean of the outcome within a bin of the explanatory variable. The solid line shows the fitted linear relationship based on the underlying data.

background match. The idiosyncratic component ε_{ijt} is assumed to be independently distributed Type I Extreme Value across alternatives.

The linear distance coefficient is therefore allowed to vary with observed patient characteristics:

$$\beta_d^i = \beta_d + \beta_{\text{age}}\text{age}_i + \beta_{\text{female}}\text{female}_i + \beta_{\text{imm}}\text{immigrant}_i, \tag{2}$$

The model does not include practice fixed effects in the baseline specification. This choice reflects both computational constraints and the fact that the main object of interest is the comparison between estimates obtained under different definitions of the feasible choice set. The estimated distance coefficient should therefore be interpreted as the association between distance and choice conditional on observed practice attributes and the constructed choice set. The identifying assumption is that, within a choice occasion and conditional on observed attributes, unobserved determinants of practice attractiveness are not systematically correlated with distance in a way that differs across the constrained and unconstrained choice-set specifications.

3.2 Choice Set Formation

The main departure from a standard revealed-preference choice model is the definition of the feasible choice set. In an unconstrained specification, the patient is assumed to be able to choose any practice within the relevant geographic radius:

$$\mathcal{C}_{it}^{\text{all}} = \{j \in \mathcal{J}_t : d_{ij} \leq \bar{r}_i\}, \tag{3}$$

where $\mathcal{J}_t$ is the set of active practices in period t and $\bar{r}_i$ is the relevant distance radius. In the baseline construction, $\bar{r}_i$ is 15 kilometers outside Copenhagen and 5 kilometers in Copenhagen. In the constrained specification, a practice is included in the feasible choice set only if its list is open to new patients in the period of the switch:

$$\mathcal{C}_{it}^{\text{open}} = \left\{j \in \mathcal{J}_t : d_{ij} \leq \bar{r}_i,\ \text{open}_{jt} = 1\right\}, \tag{4}$$

where open_{jt} is an indicator for whether practice j is accepting new registrants in period t. Because list status is recorded administratively, the rationing rule is observed directly rather than inferred from observed choices or modeled as a latent acceptance rule.

Given a feasible choice set $\mathcal{C}_{it}$, the conditional logit choice probability is

$$P_{ijt}(\mathcal{C}_{it}) = \frac{\exp\left(\beta_d^i d_{ij} + \beta_{d2} d_{ij}^2 + \boldsymbol{\gamma}'\mathbf{X}_{ijt}\right)}{\sum_{k \in \mathcal{C}_{it}} \exp\left(\beta_d^i d_{ik} + \beta_{d2} d_{ik}^2 + \boldsymbol{\gamma}'\mathbf{X}_{ikt}\right)}. \tag{5}$$

The comparison between $\mathcal{C}_{it}^{\text{all}}$ and $\mathcal{C}_{it}^{\text{open}}$ measures how estimated preferences change when closed-list practices are removed from the set of feasible alternatives.

3.3 Identification and Switcher Groups

The comparison across choice-set definitions identifies the role of capacity constraints in revealed-preference estimation. If closed-list practices are treated as feasible alternatives, patients who would have preferred a closed practice but choose another provider are interpreted as having voluntarily selected the observed alternative from the full set of nearby practices. This misclassification can attenuate distance aversion: a patient who chooses a more distant open practice because nearby practices are closed appears, in the unconstrained specification, to have a relatively weak preference for proximity. Conditioning on observed open-list status removes unavailable practices from the denominator of the choice probability and therefore separates preferences over providers from supply-side restrictions on access.

I estimate the model separately for switcher groups that differ in the source of the choice occasion. The main estimation sample consists of patients who actively choose a provider after their original practices have closed (*active closure switchers*). This group combines two features that are particularly attractive for preference estimation. First, patients make an active choice among available providers, so the observed registration decision directly reveals preferences. Second, the timing of the choice is induced by the closure of the previous practice rather than by a patient-initiated decision to search. Because the need to switch is generated by an institutional shock, this sample is less susceptible to selection on the timing of switching than voluntary switchers, whose choices may reflect unobserved dissatisfaction, health shocks, or other motives for changing provider.

The analysis focuses on switchers more generally because these individuals reveal preferences through an observed choice among alternatives. Non-switchers are not included in the baseline estimation sample. Remaining with the current GP is difficult to interpret in the presence of capacity constraints. For example, some patients may stay because they are satisfied with their existing provider, while others may prefer to switch but be unable to find an acceptable alternative with an open patient list. Observing no switch therefore does not necessarily imply that the current provider is the preferred option. Restricting attention to switchers avoids mixing genuine satisfaction with constrained inaction and ensures that each observation corresponds to a realized choice among feasible providers.

To assess the importance of selection into switching, I also estimate the model for broader switcher samples. Voluntary switchers make active changes that are not directly linked to practice closure. Closure switchers include all patients displaced by the exit of their previous practice, regardless of whether they actively choose a new provider or are subsequently assigned by administrative procedures. Comparing estimates across these groups helps evaluate the extent to which preference estimates depend on the source and timing of the switching decision.

Patients who are administratively assigned after closure are treated separately from active choices. Their post-closure provider reflects the regional assignment process rather than a standard patient choice, so they are not used as active choices in the baseline conditional logit estimation. They are retained for descriptive and welfare analyses, where they provide evidence on the consequences of assignment under constrained availability.

The two main choice-set specifications are:

- *Unconstrained specification:* $\mathcal{C}_{it} = \mathcal{C}_{it}^{\text{all}}$. All active practices within the relevant radius are treated as feasible alternatives.
- *Constrained specification:* $\mathcal{C}_{it} = \mathcal{C}_{it}^{\text{open}}$. Only practices with open lists within the relevant radius are treated as feasible alternatives.

Both specifications are estimated across switcher groups, including all switchers, voluntary switchers, closure switchers, and active-at-closure switchers.

3.4 Estimation Results

Table 3.1: Preference estimates: all switchers and active-at-closure switchers

	All switchers		Active-at-closure	
	(1) All clinics	(2) Open clinics	(3) All clinics	(4) Open clinics
Distance	-0.689	-0.733	-0.872	-1.011
	(0.002)	(0.002)	(0.005)	(0.008)
Distance squared	0.031	0.032	0.037	0.040
	(0.000)	(0.000)	(0.000)	(0.000)
Distance x Patient age	-0.002	-0.002	-0.001	-0.002
	(0.000)	(0.000)	(0.000)	(0.000)
Distance x Female patient	-0.005	-0.007	-0.001	-0.005
	(0.001)	(0.001)	(0.003)	(0.005)
Distance x Immigrant patient	0.009	0.004	-0.014	-0.015
	(0.001)	(0.002)	(0.005)	(0.007)
Clinic size	0.094	0.048	0.279	0.263
	(0.001)	(0.001)	(0.003)	(0.004)
Physician age	-0.078	-0.083	-0.025	-0.037
	(0.000)	(0.000)	(0.000)	(0.001)
Share of female physicians	-0.167	0.348	-0.399	0.163
	(0.003)	(0.004)	(0.011)	(0.014)
Physicians with immigration background	-0.039	-0.753	0.357	-0.592
	(0.005)	(0.005)	(0.013)	(0.017)
Gender match	0.222	0.384	0.277	0.443
	(0.004)	(0.006)	(0.013)	(0.018)
Immigration background match	0.422	0.424	0.318	0.312
	(0.007)	(0.009)	(0.021)	(0.028)
Cases	1003707	882922	101563	89054
Observations	103680207	18761618	12124897	1769357

Notes: The table reports conditional logit estimates of primary care provider choice. Cases are patient-quarter choice occasions. Observations are patient-clinic alternatives. Columns (1) and (3) define the choice set as all active practices within the relevant distance radius. Columns (2) and (4) restrict the choice set to practices with open lists. The all-switchers specifications are estimated using a 50 percent stratified random sample due to computational constraints. Sampling is stratified by age group, patient municipality, and year-quarter. Age groups are defined as below 18, 18-39, 40-64, and 65 and above. Standard errors are reported in parentheses.

Table 3.1 reports conditional logit estimates for GP choice under two definitions of the choice set. Columns (1) and (3) treat all active clinics within the relevant distance radius as feasible alternatives. Columns (2) and (4) restrict the choice set to clinics with open lists. The comparison is shown both for all switchers and for active-at-closure switchers, the preferred sample in which the timing of the choice is induced by the closure of the previous practice but the new provider is actively selected by the patient.

Across all specifications, patients are less likely to choose more distant clinics. The distance coefficient is negative and precisely estimated. For all switchers, the coefficient changes from -0.689 when all nearby clinics are treated as feasible to -0.733 when the choice set is restricted to open-list clinics. For active-at-closure switchers, the corresponding change is larger, from -0.872 to -1.011. Thus, imposing the observed capacity constraint increases the estimated magnitude of distance aversion. This pattern is consistent with the central mechanism of the paper: when closed-list clinics are incorrectly treated as available, patients who choose more distant open clinics are interpreted as having a weaker preference for proximity, even though nearby clinics may not have been feasible options. The difference across switcher groups points in the same direction. Distance aversion is substantially stronger among active-at-closure switchers than among all switchers. These patients are required to choose a new GP because their previous practice exits, rather than switching at a time when their preferred clinic is open for new intakes. Therefore, the stronger distance aversion estimated for this group suggests that selection into the timing of switching may bias the estimates for broader switcher groups.

The remaining coefficients show that patients also respond to observable practice characteristics. Larger practices are more likely to be chosen, and gender and immigration-background matches are positively associated with choice. The coefficients on physician composition differ across choice-set definitions, which is consistent with the idea that closed lists are not randomly distributed across practices. Removing closed-list clinics changes not only the number of alternatives, but also the composition of the alternatives against which each chosen clinic is compared.

Figure 3.1 summarizes the estimated distance coefficients across switcher groups and choice-set definitions. For each group, the figure compares the coefficient obtained when all clinics within the relevant radius are treated as feasible alternatives with the coefficient obtained when the choice set is restricted to open-list clinics. The estimates are negative in all groups, indicating that patients are less likely to choose more distant providers. Appendix Tables A.1 and A.2 report the full conditional logit estimates for the main four groups.

The main pattern is that ignoring capacity constraints leads to a less negative estimated distance coefficient. In other words, failing to account for clinic availability consistently understates the magnitude of distance aversion across almost all switcher groups. The difference is largest for active closure switchers, where the unconstrained model underestimates the distance coefficient by 13.7 percent. The corresponding differences are 7.7 percent for closure switchers, 6.0 percent for all switchers, and 4.4 percent for voluntary switchers.

An additional group shown in the figure, not part of the four main sample definitions, is switchers who are assigned to a clinic following a clinic closure and later switch away from that

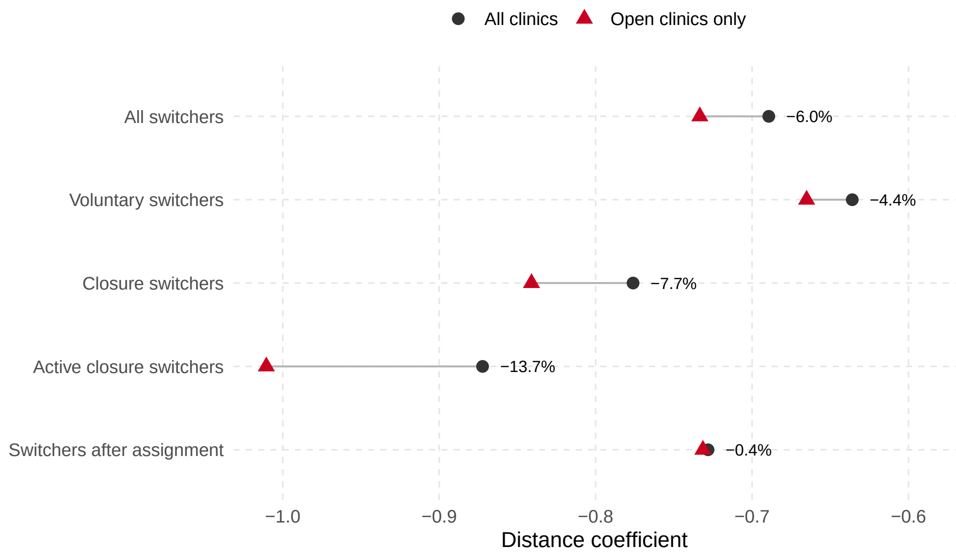

Figure 3.1: Distance coefficient estimates across switcher groups and choice sets

Notes: All clinics define the choice set as all active practices within the relevant distance radius. Open clinics restrict the choice set to practices with open lists. Percentage differences are with respect to the estimates from the model using open clinics as choice sets.

assignment. For this group, the estimates are nearly unchanged, with a difference of 0.4 percent. The observed choice of these patients is therefore not the immediate closure-induced choice, but a later correction to an assigned match. The small difference between the all-clinic and open-clinic specifications suggests that, for this group, the main selection margin is the decision to leave the assigned clinic, rather than the exclusion of closed clinics from the feasible choice set. This distinguishes them from active-at-closure switchers, for whom the choice occasion is directly induced by closure, and the open-list restriction substantially increases the estimated magnitude of distance aversion.

The negative difference supports the interpretation that unconstrained choice sets attenuate estimated distance aversion. When closed-list clinics are included as if they were feasible alternatives, patients who choose a more distant open clinic are treated as having passed over nearby clinics voluntarily. Restricting the choice set to open clinics removes unavailable alternatives from the denominator of the choice probability. The resulting estimates attribute the observed choice to preferences over feasible providers rather than to a mixture of preferences and capacity-driven exclusion. The discrepancy driven by feasible alternatives being the largest among active closure switchers also supports the identification strategy because this group makes an active provider choice at a time when the need to switch is induced by the closure of the previous practice.

4 Welfare Implications

4.1 Welfare Measures

I translate the preference estimates into two complementary measures of the burden imposed by capacity constraints. The first measure is the *logsum welfare loss*:

$$\Delta V_{it} = \log \sum_{j \in \mathcal{C}_{it}^{\text{all}}} \exp\Big(\hat{V}_{ijt}\Big) - \log \sum_{j \in \mathcal{C}_{it}^{\text{open}}} \exp\Big(\hat{V}_{ijt}\Big) \geq 0, \tag{6}$$

where $\hat{V}_{ijt}$ denotes the estimated deterministic component of utility. This logsum difference is the standard inclusive-value measure of the welfare change associated with restricting the choice set. It captures the reduction in expected utility when closed-list clinics are removed from the feasible choice set. I report the logsum loss in utility units and compute both welfare measures using the constrained active-at-closure estimates as the preferred parameters.

The second measure is the *expected distance gap*:

$$\Delta D_{it} = \mathrm{E}[d_{ij} \mid j \in \mathcal{C}_{it}^{\text{open}}] - \mathrm{E}\Big[d_{ij} \mid j \in \mathcal{C}_{it}^{\text{all}}\Big], \tag{7}$$

where expectations are taken with respect to the logit choice probabilities implied by equation (5). The constrained set $\mathcal{C}_{it}^{\text{open}}$ contains open-list clinics within the relevant distance radius, while $\mathcal{C}_{it}^{\text{all}}$ contains all active clinics within the same radius. The expected distance gap measures the additional expected travel distance generated by list closures. It compares the predicted distance under the observed capacity constraint with the predicted distance that would arise if all clinics in the geographic choice set were open. Unlike the logsum measure, it is expressed directly in kilometers and therefore provides an intuitive measure of the access burden imposed by capacity constraints.

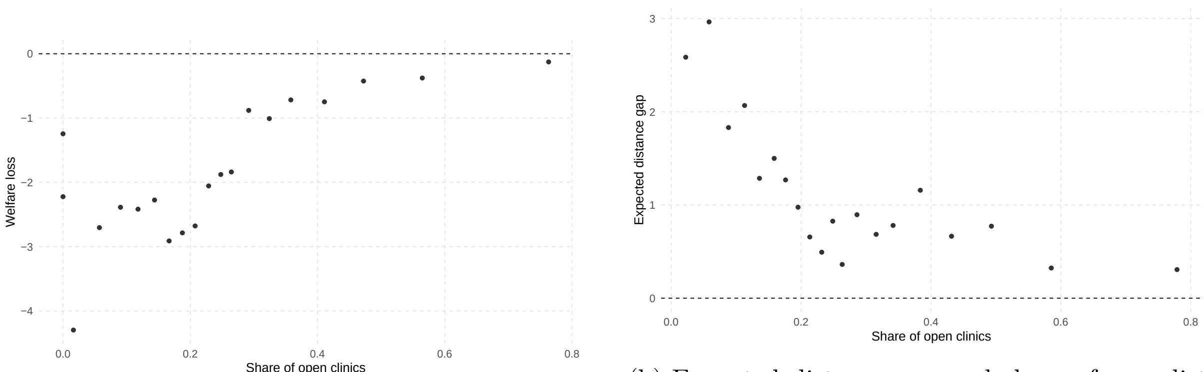

(a) Welfare loss and share of open-list clinics

(b) Expected distance gap and share of open-list clinics

Figure 4.1: Relationship between welfare measures and share of open-list clinics in the choice set

Notes: Panels (a) and (b) present binned scatter plots based on choice case (patient-quarter) observations. Alternatives in each choice set are defined as clinics within 15 km (5 km for Copenhagen) of a patient's residential address. The solid line shows the fitted linear relationship based on the underlying data.

Figure 4.1 shows how the welfare burden of capacity constraints varies with local provider availability for the sample of all switchers. Panel (a) plots the logsum welfare loss against the share of open clinics in the patient's choice set, while Panel (b) plots the expected distance gap against the same measure of availability. Both panels show a negative relationship, where switchers facing a lower share of open clinics experience larger welfare losses. In Panel (a), lower open-list availability is associated with a larger reduction in expected utility, consistent with the idea that closed lists remove attractive alternatives from the feasible set. In Panel (b), lower open-list availability is associated with a larger expected distance gap, indicating that patients are predicted to register farther away when nearby clinics are closed. The two panels therefore provide complementary evidence on the same mechanism, suggesting that when the local market offers fewer open clinics, patients face both a larger utility loss and a greater expected travel burden.

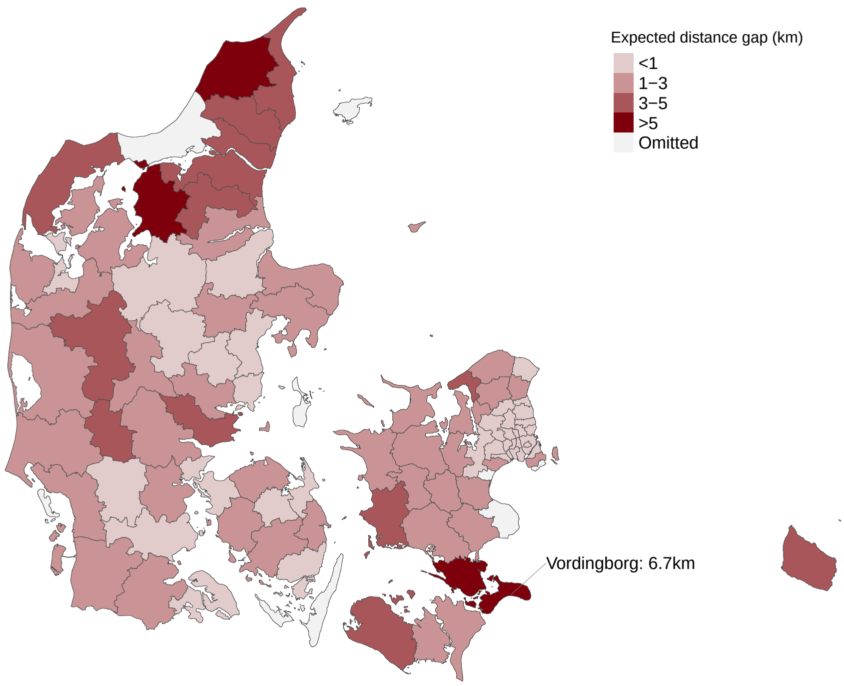

Figure 4.2: Geographic distribution of expected distance gaps

Notes: The figure maps the expected distance gap averaged across clinics and year-quarters within each municipality over 2014-2019. The expected distance gap is the difference between predicted distance under the observed open-list choice set and predicted distance under the counterfactual choice set in which all clinics are open. Darker shading indicates larger average expected distance gaps. Municipalities with fewer than 5 clinics are omitted to ensure anonymity.

Figure 4.2 maps the geographic distribution of the expected distance gap. For each municipality, the measure is averaged across clinics and year-quarters over the period 2014-2019. The map therefore summarizes the average additional expected travel distance associated with capacity constraints in each local market over the sample period. Darker areas correspond to municipalities where closed lists generate larger predicted increases in travel distance. The figure shows substantial geographic variation, with many municipalities having average gaps between one and three kilometers and several municipalities exceeding five kilometers. Vordingborg is highlighted with an average expected distance gap of 6.7 kilometers. This spatial pattern indicates that the welfare burden of capacity constraints is concentrated in local markets where closed lists substantially restrict access to nearby providers.

4.2 Aggregate Welfare Cost

Table 4.1 reports the welfare measures by switcher group. Across all switchers, the average logsum welfare loss is 1.80 utility units, while capacity constraints increase expected distance by 1.12 kilometers on average. The burden is not evenly distributed across switcher groups. Active closure switchers experience the largest welfare losses under both measures, with a logsum loss of 2.43 utility units and an expected distance gap of 1.52 kilometers. Takeover switchers and closure switchers also face relatively large losses, while assigned switchers experience a smaller expected distance gap of 0.89 kilometers despite still incurring a substantial utility loss. These differences show that the welfare cost of capacity constraints depends on the institutional source of the switch and on the set of alternatives available at the time of re-registration.

Table 4.2 reports complementary evidence on whether patients' predicted highest-utility clinics are available. Among all switchers, 65 percent have a highest-utility clinic that is closed. The share is higher among moving switchers and active closure switchers, both at 72 percent, and among assigned switchers, at 70 percent. Only 19 percent of all switchers are observed in their predicted highest-utility clinic, and 17 percent are observed in their predicted highest-utility open clinic. These patterns indicate that capacity constraints are not only present in the aggregate, but often bind for the specific clinic that the model predicts the patient would prefer most.

Table 4.1: Expected distance gap and logsum welfare loss, by switcher group

Population	Cases	Welfare constrained	Welfare loss	Expected distance constrained	Distance gap
All switchers	2007618	0.34 (0.57)	-1.80 (2.53)	3.71 (3.36)	1.12 (2.80)
Voluntary switchers	1032198	0.36 (0.58)	-1.98 (2.67)	3.56 (3.23)	1.02 (2.68)
Moving switchers	562139	0.43 (0.64)	-2.20 (2.75)	3.37 (3.13)	1.01 (2.63)
Closure switchers	975420	0.32 (0.55)	-1.61 (2.35)	3.86 (3.47)	1.22 (2.92)
Takeover switchers	486590	0.33 (0.56)	-1.47 (2.28)	4.11 (3.76)	1.44 (3.31)
Assigned switchers	387245	0.33 (0.56)	-1.70 (2.35)	3.47 (3.03)	0.89 (2.31)
Active closure switchers	101559	0.24 (0.43)	-1.96 (2.61)	4.13 (3.49)	1.52 (2.93)

Table 4.2: Share of switch events, by switcher group

Population	Cases	Share top clinic closed	Share in top clinic	Share in top open clinic
All switchers	2007618	0.65	0.19	0.17
Voluntary switchers	1032198	0.63	0.18	0.15
Moving switchers	562139	0.72	0.16	0.17
Closure switchers	975420	0.67	0.20	0.19
Takeover switchers	486590	0.64	0.24	0.17
Assigned switchers	387245	0.70	0.16	0.21
Active closure switchers	101559	0.72	0.18	0.13

Notes: Indirect utilities are computed using conditional logit estimates from active switchers at clinic closure. Top clinic refers to the patient's predicted highest-utility clinic. Cases refer to patient-quarter level switch events.

Additionally, Figure 4.3 shows the distribution of utility gap between the predicted highest-utility clinic in an unconstrained scenario and the actual choice. Many patients experience moderate losses, but there is a long upper tail of patients with losses above four utility units, indicating substantial heterogeneity. Capacity constraints impose limited costs for patients whose preferred open alternatives are close substitutes, but much higher costs for patients whose preferred clinics are closed and whose remaining feasible alternatives are less attractive.

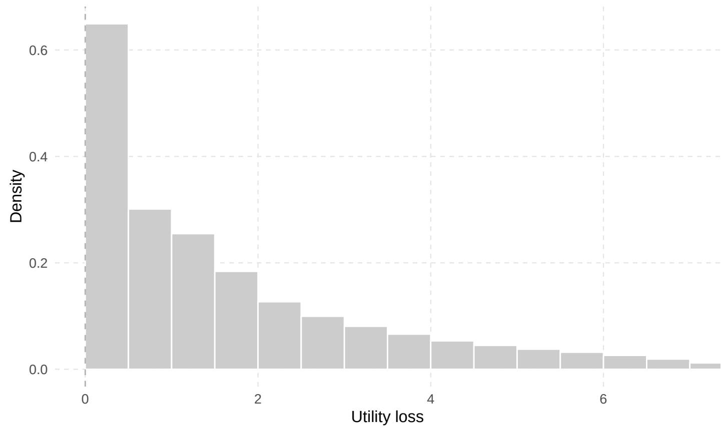

Figure 4.3: Distribution of utility gap between highest-utility and chosen clinic

4.3 Causal Machine Learning: Heterogeneous Effects of Constraint Exposure

The previous section shows that welfare losses differ across switcher groups, highlighting one important dimension of heterogeneity linked to the institutional circumstances under which patients switch providers. However, switcher type alone does not capture the full range of variation in welfare losses, which may also depend on patient characteristics and the availability of alternative clinics in the local market. I therefore examine which switch events are most affected by capacity constraints using a causal machine learning approach. The treatment is an indicator that the patient's predicted highest-utility clinic has a closed list. The main outcome is the expected distance gap, and I also examine chosen distance as a complementary outcome. The causal forest estimator of Wager and Athey (2018) allows the treatment effect to vary flexibly with patient characteristics, local market supply, and switch type.

Figure 4.4 shows the distribution of estimated conditional average treatment effects for the expected distance gap. The distribution is right-skewed, indicating that the effect of a closed top-choice clinic is modest for many switch events but large for a smaller group of high-burden cases. This pattern is consistent with the welfare results above: capacity constraints do not impose the same cost on all patients, because some patients have close open substitutes while others face a much smaller feasible choice set.

Figure 4.5 illustrates the corresponding distribution for chosen distance. This outcome captures realized registration distances rather than the model-based expected distance gap. The distribution is centered near zero but exhibits substantial heterogeneity, with both positive and negative estimated effects. The right tail indicates that, for some patients, having a closed top-

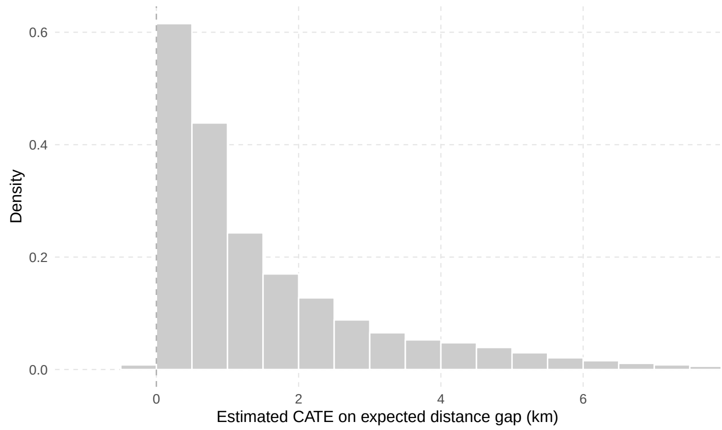

Figure 4.4: Distribution of estimated CATE for expected distance gap

Notes: The treatment is an indicator that the patient's predicted highest-utility clinic has a closed list. The outcome is the expected distance gap (km). The sample consists of all switch events (all switchers) used in the welfare analysis.

choice clinic is associated with a considerably longer chosen distance. These are cases in which the capacity constraint appears to push patients toward more distant providers. The mass near zero suggests that many patients have nearby open substitutes, so the closure of the predicted top-choice clinic does not substantially change the distance to the chosen clinic.

The presence of negative estimated effects reflects that chosen distance is a realized outcome shaped by the full choice environment, including idiosyncratic preferences, the location of open substitutes, and the institutional circumstances of the switch. For example, a patient whose predicted top-choice clinic is closed may choose a different open clinic that is closer but less preferred along other dimensions. Administrative reassignment after clinic closures may also send patients to the remaining nearby clinic, even if it is not their preferred option in terms of overall utility. These negative estimated effects of the top-utility clinic being closed indicate that patients often trade off distance against other attributes of alternative clinics.

Consistent with the negative correlation between open-list clinics and overall welfare loss, local supply is a prominent source of heterogeneity. Figure 4.6 plots mean estimated treatment effects for the expected distance gap by the number of open-list clinics within the choice set. The treatment effect is largest when patients have few open alternatives and declines as the number of open clinics increases. This pattern shows that the welfare cost of a closed top-choice clinic depends not only on whether the preferred clinic is unavailable, but also on the quality of the remaining feasible choice set. In thick markets, patients are more likely to have close open substitutes. In markets where few clinics are open, the same closure of the top-choice clinic generates a larger expected distance penalty.

Figure 4.7 summarizes heterogeneity by switcher type using chosen distance as the outcome. The estimated effects are largest for active-at-closure and assigned switchers, and smaller for voluntary non-closure switchers. This ordering is consistent with the structural welfare results in Table 4.1. Patients displaced by clinic closure face a choice occasion generated by the exit of

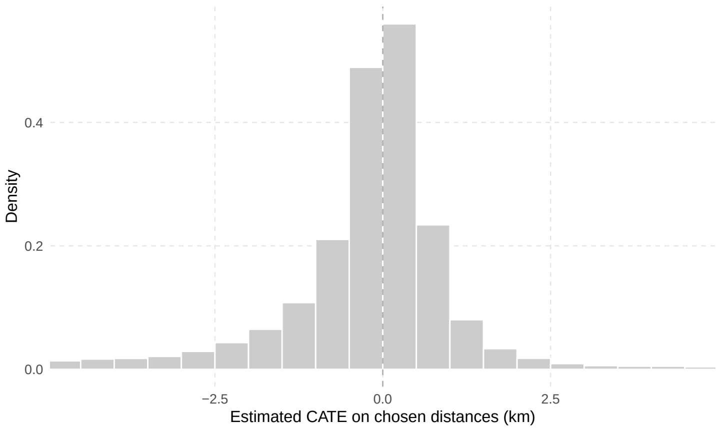

Figure 4.5: Distribution of estimated CATE for chosen distance

Notes: The treatment is an indicator that the patient's predicted highest-utility clinic has a closed list. The outcome is the distance to the chosen clinic (km). The sample consists of all switch events (all switchers) used in the welfare analysis.

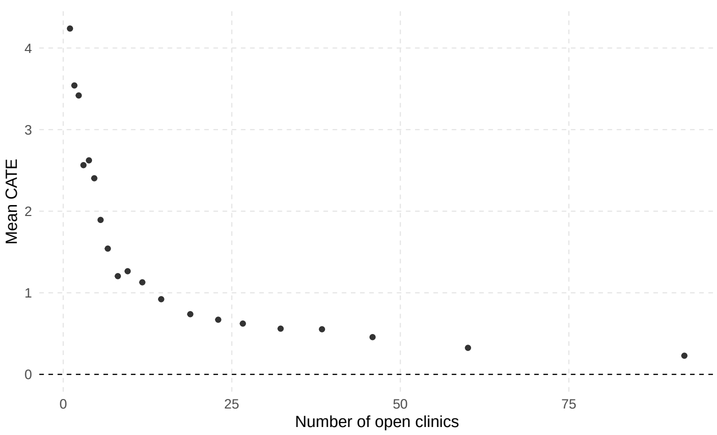

Figure 4.6: Mean estimated CATE for expected distance gap by number of open-list clinics

their previous provider, and the consequences of a closed top-choice clinic are amplified when the local set of open alternatives is limited. Voluntary switchers, by contrast, are less exposed to this institutional displacement margin and show smaller distance effects.

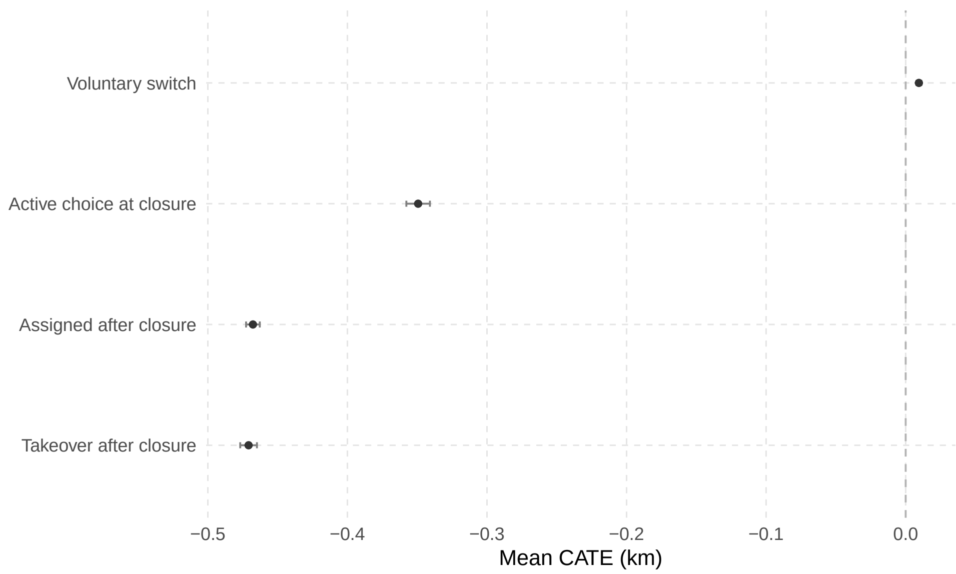

Figure 4.7: Mean estimated CATE for chosen distance by switcher type

Notes: The treatment is an indicator that the patient's predicted highest-utility clinic has a closed list. The outcome is the distance to the chosen clinic. Error bars report 95 percent confidence intervals.

Overall, the causal forest results indicate that the burden associated with a closed top-choice clinic varies systematically with local supply. As a robustness check, Appendix B reports linear OLS regressions of the expected distance gap on patient characteristics, past primary care utilization, local market supply, and switch type. These regressions support the same qualitative conclusion: the number of open-list clinics within the choice set is the strongest predictor of the expected distance gap, and closure-related switches are associated with larger welfare losses.

5 Counterfactual Policy Analysis

The welfare measures in Section 4.1 describe the status-quo burden of capacity constraints. They show how much expected utility and expected travel distance are affected by the fact that some clinics are closed to new patients. I now use the preferred demand estimates to evaluate two counterfactual capacity expansions. Both counterfactuals are implemented at the municipality level. For each municipality, I identify the clinic with the highest average excess demand over the sample period, but its list is closed under the observed capacity constraint. A closed-list clinic's excess demand is defined as the sum of predicted choice probabilities from all switching patients. I then aggregate this measure across switch events and year-quarters to identify the closed-list clinic in each municipality that has the highest predicted demand. Figure 5.1 maps the geographic distribution of this excess-demand measure.

The two counterfactuals differ in how new capacity is introduced. The first expands the

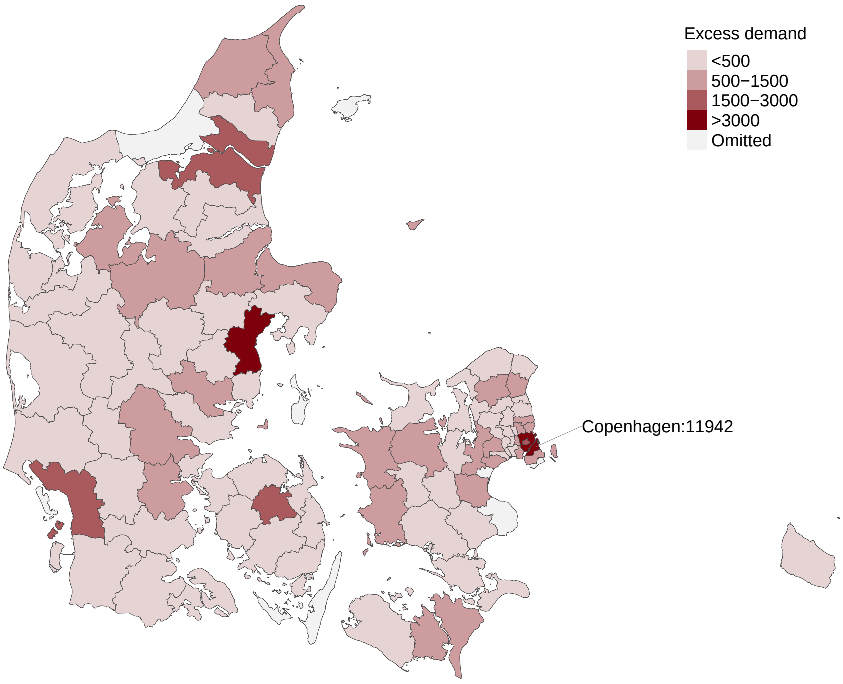

Figure 5.1: Geographic distribution of excess demand

Notes: The figure maps municipality-level excess demand for primary care clinics, measured by the total demand for closed-list clinics in each municipality, averaged across quarters from 2014-2019. Excess demand is measured as the sum of choice probabilities from all switching patients within 15 km (within 5 km for Copenhagen). Municipalities with fewer than 5 clinics are omitted to ensure anonymity.

existing top-demand clinic by one physician, thereby allowing the clinic to re-open its list to new patients. The second adds a new clinic at the same location as the top-demand clinic. Comparing these two scenarios separates the role of location from the role of incumbency. If the two counterfactuals generate similar gains, most of the welfare improvement comes from placing additional capacity at the right location. If expanding the incumbent generates larger gains, this suggests that patients value features of the existing clinic beyond its geographic location, such as observed clinic characteristics or prior attachment.

Counterfactual 1: Expanding the Top-Demand Clinic by One Physician. The first counterfactual adds one physician to the top-demand clinic in each municipality. In the Danish list system, one additional full-time GP corresponds to an increase in list capacity of approximately 1,600 patients. I implement this counterfactual by allowing the selected top-demand clinic to enter the feasible choice set for patients in the municipality. This captures the direct effect of relaxing the list constraint at the clinic where the model predicts unmet demand to be highest. Let $\mathcal{C}_{it}^{\text{open}}$ denote the observed feasible choice set and let $\mathcal{C}_{it}^{\text{expand}}$ denote the counterfactual choice set after the top-demand clinic is made available. The expected-distance gain for patient i in period t is

$$\Delta D_{it}^{\text{expand}} = \mathrm{E}[d_{ij} \mid j \in \mathcal{C}_{it}^{\text{open}}] - \mathrm{E}\Big[d_{ij} \mid j \in \mathcal{C}_{it}^{\text{expand}}\Big] \geq 0, \tag{8}$$

where expectations are computed using the logit choice probabilities from equation (5). A positive value indicates that adding capacity at the top-demand clinic reduces the patient's predicted registration distance. I compute the corresponding logsum gain as the increase in the inclusive value when moving from the observed feasible set to the expanded feasible set.

Counterfactual 2: New Clinic Entry at the Top-Demand Location. The second counterfactual adds a new clinic at the same geographic location as the top-demand clinic in each municipality. The entrant is assigned the capacity of one full-time physician, corresponding to 1,600 listed patients, and is treated as open to new registrants. Because the new clinic is located at the same address as the top-demand clinic, it offers the same distance to patients as the constrained incumbent. The counterfactual therefore isolates the value of adding capacity at a high-demand location without expanding the incumbent clinic itself. For simplicity, I set clinic characteristics, including physician age, gender and immigration background, to the clinic sample average.

Let $\mathcal{C}_{it}^{\text{entry}}$ denote the feasible choice set after adding the new clinic. The expected-distance gain is computed analogously to equation (8), replacing $\mathcal{C}_{it}^{\text{expand}}$ with $\mathcal{C}_{it}^{\text{entry}}$. This new-entry scenario should be interpreted as a location-targeted capacity expansion rather than as a full model of entry. It abstracts from equilibrium responses such as changes in list status at other clinics, physician relocation, and changes in patient congestion after entry. The exercise is nevertheless useful because it asks whether placing one additional clinic at the location of estimated excess demand would materially reduce the welfare burden of capacity constraints.

Figure 5.2 relates the logsum welfare gain from each counterfactual to municipality-level excess demand. The figure indicates that municipalities with higher excess demand experience larger welfare gains from capacity expansion. This supports the targeting rule used in the counterfactual design. Municipalities where the total demand of capacity-constrained clinics is larger are precisely the municipalities where adding capacity generates larger welfare gain. The welfare gains are consistently larger when capacity is added by expanding the top constrained incumbent clinic rather than by opening a new clinic at the same location. Expanding an existing clinic increases capacity at a provider that already has an established patient base and high predicted demand, so the additional slots are immediately valued by many patients who would otherwise be diverted to less preferred alternatives. In contrast, a new clinic at the same location inherits the geographic advantage but not the incumbent-specific utility captured by the model, which limits the increase in expected utility. As a result, incumbent expansion relaxes the binding capacity constraint at a highly valued provider and generates larger welfare gains.

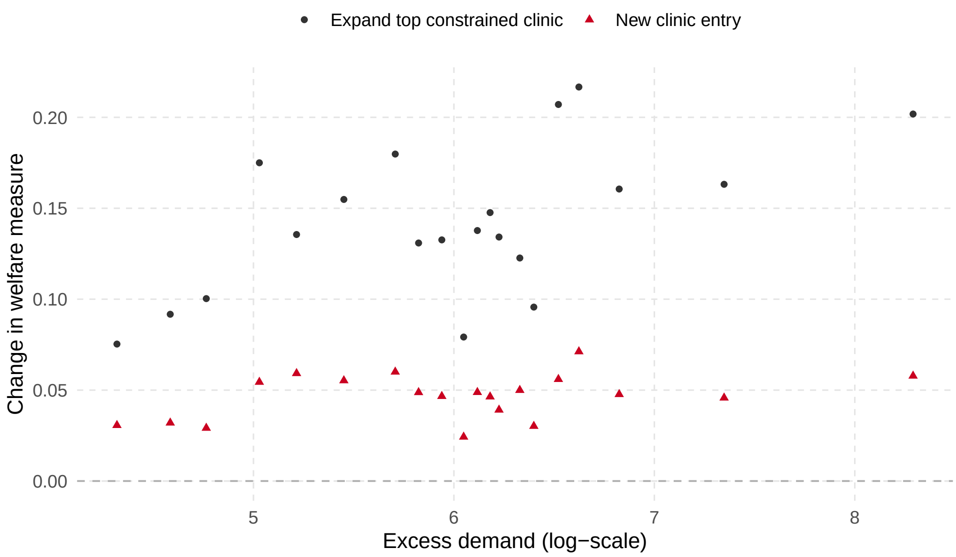

Figure 5.2: Logsum welfare gain from counterfactual capacity expansion and municipality-level excess demand

Notes: The figure plots the change in the logsum welfare measure under each counterfactual against municipality-level excess demand as a binned scatter of 20 bins. Excess demand is measured as the frequency with which a closed clinic is the predicted highest-utility choice under the unconstrained choice set and is shown in log scale. The two series compare expanding the top constrained clinic and establishing a new clinic at the same location.

Figure 5.3 plots the change in the expected distance gap under the two counterfactual scenarios against the baseline expected distance gap. Both counterfactuals generate larger reductions for switch events with larger baseline gaps. This pattern suggests that the interventions are concentrated on high-burden cases. Patients whose predicted distance to the provider is most distorted by capacity constraints are also those who gain most when additional capacity is placed at the top-demand location. The fitted relationships are similar for incumbent expansion and new clinic entry, which is consistent with the two policies adding capacity at the same geographic location.

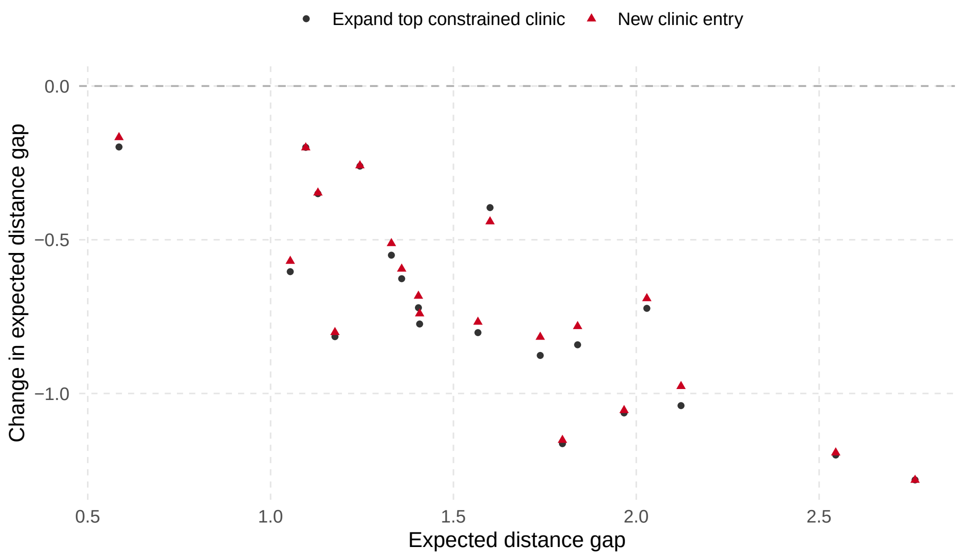

Figure 5.3: Counterfactual change in expected distance and baseline expected distance

Notes: The figure plots the change in the expected distance gap against the baseline expected distance gap as a binned scatter of 20 bins. The baseline expected distance gap is the difference between predicted distance under the observed open-list choice set and predicted distance under the choice set in which all clinics are open. The two series compare expanding the top constrained clinic and adding a new clinic at the same location.

Table 5.1 reports the baseline expected distance gap and the counterfactual gaps under both capacity expansions, by switcher group. Across all switchers, the baseline expected distance gap is 1.12 kilometers. Expanding the top constrained clinic reduces this gap to 0.54 kilometers, while adding a new clinic at the same location reduces it to 0.56 kilometers. The two interventions therefore reduce the average expected distance gap by roughly one half. The gains are larger for switchers whose provider change is connected to clinic closure. Closure switchers have a baseline gap of 1.22 kilometers, which falls to 0.40 kilometers under incumbent expansion and 0.44 kilometers under new entry. Active closure switchers experience the largest baseline burden, with an expected distance gap of 1.52 kilometers. For this group, the gap falls to 0.52 kilometers under incumbent expansion and 0.57 kilometers under new entry. Takeover switchers also experience large reductions, from 1.44 kilometers at baseline to 0.52 and 0.55 kilometers under the two counterfactuals. These patterns indicate that targeted capacity additions are most valuable for patients whose choice occasion is directly shaped by exogenous shock to the existing patient-provider relationship.

Assigned switchers have a lower baseline gap of 0.89 kilometers, but still experience substantial gains. Their expected distance gap falls to 0.22 kilometers under incumbent expansion and 0.27 kilometers under new entry. Across every switcher group, expanding the existing top constrained clinic produces a slightly larger reduction in the expected distance gap than adding a new clinic at the same location. This difference is small, but it is consistent across groups. The result suggests that most of the gain comes from placing capacity in the high-demand location, while a smaller additional gain comes from relaxing the constraint at an incumbent clinic that already has high predicted demand.

Table 5.1: Expected distance gap under baseline and counterfactual capacity expansions, by switcher group

Population	Baseline	Expand top constrained clinic	New clinic entry
All switchers	1.12	0.54	0.56
Voluntary switchers	1.02	0.67	0.68
Moving switchers	1.01	0.43	0.46
Closure switchers	1.22	0.40	0.44
Takeover switchers	1.44	0.52	0.55
Assigned switchers	0.89	0.22	0.27
Active closure switchers	1.52	0.52	0.57

Notes: The table reports expected distance gaps in kilometers. The baseline expected distance gap is the difference between predicted distance under the observed open-list choice set and predicted distance under the choice set in which all clinics are open. The counterfactual columns report the expected distance gap after adding capacity either by expanding the top constrained clinic in each municipality or by opening a new clinic at the same location. Choice probabilities are computed using the conditional logit estimates from active closure switchers. Cases refer to patient-quarter switch events.

Overall, the counterfactual analysis shows that capacity expansions can substantially reduce the welfare burden generated by closed patient lists, but that the gains depend on where new capacity is placed. Adding one physician to the top constrained clinic reduces the average expected distance gap by roughly one half, and opening a new clinic at the same location produces a similar, though slightly smaller, improvement. The largest gains accrue to patients whose switching events are directly linked to clinic closure and to municipalities where excess demand for closed clinics is highest. These findings suggest that capacity constraints matter not only for current access, but also for the evaluation of access policies. They also indicate that accounting for list closures can affect estimates of the benefits from targeted capacity expansions and the distribution of those benefits across locations.

6 Conclusion

This paper examines how capacity constraints affect revealed-preference estimates of patient demand in primary care. Using the Danish GP list system, where clinics can close their lists to new patients and list status is directly observed, I distinguish between geographically nearby providers and providers that are actually feasible choices. I use this feature to estimate GP choice under alternative definitions of the choice set and to quantify how ignoring closed lists changes estimated preferences, welfare measures, and counterfactual predictions.

The results show that capacity constraints are a central feature of the market. Closed lists are common, substantially reducing patients' feasible choice sets, and in 65 percent of switch events the patient's predicted preferred clinic is closed. Ignoring these constraints leads to an understatement of distance aversion, particularly among patients who actively choose a new GP following a practice closure. The welfare analysis shows that capacity constraints impose meaningful costs on patients. Across all switchers, closed lists increase expected travel distance by 1.12 kilometers on average, with larger effects for closure-related switchers. These burdens are concentrated in areas with limited local supply, highlighting the importance of market structure for access to care. Counterfactual simulations suggest that targeted capacity expansions can substantially reduce these access costs. Adding one physician to the most constrained clinic in each municipality reduces the average expected distance gap from 1.12 to 0.54 kilometers, while opening a new clinic at the same location yields a similar reduction.

More broadly, the paper demonstrates the importance of accounting for feasible choice sets when estimating demand. Administrative information on list status provides a clearer picture of access than measures based solely on the number of nearby providers. While the empirical setting is Danish primary care, the underlying issue extends to many markets where institutional constraints shape observed choices. Such information can help evaluate access policies by distinguishing between areas with many clinics and areas with many feasible clinics. This distinction matters for reforms such as increasing list capacity, supporting recruitment in underserved municipalities, or targeting new practice entry. The results suggest that policies based only on observed registrations or on the geographic density of clinics may understate the welfare gains from expanding capacity in high-demand constrained locations.

Overall, the findings show that capacity constraints change both the measurement and the consequences of patient choice. When closed-list providers are treated as available alternatives, observed registrations understate patients' aversion to distance and understate the welfare burden of restricted access. Accounting for list status reveals that patient choice in primary care is shaped not only by preferences over providers, but also by the institutional rules that determine which providers are open to new patients.

References

Agarwal, N. and Somaini, P. (2018). Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism. *Econometrica*, 86(2):391–444.

Agarwal, N. and Somaini, P. (2025). Demand Analysis under Latent Choice Constraints. *The Review of Economic Studies*, page rdaf093.

Barseghyan, L., Coughlin, M., Molinari, F., and Teitelbaum, J. C. (2021). Heterogeneous Choice Sets and Preferences. *Econometrica*, 89(5):2015–2048.

Beckert, W. (2018). Choice in the presence of experts: The role of general practitioners in patients' hospital choice. *Journal of Health Economics*, 60:98–117.

Beckert, W., Christensen, M., and Collyer, K. (2012). Choice of NHS-funded Hospital Services in England. *The Economic Journal*, 122(560):400–417.

Brown, Z. Y., Hansman, C., Keener, J., and Veiga, A. F. (2023). Information and Disparities in Health Care Quality: Evidence from GP Choice in England.

Chandra, A., Finkelstein, A., Sacarny, A., and Syverson, C. (2016). Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector. *American Economic Review*, 106(8):2110–2144.

Crawford, G. S., Griffith, R., and Iaria, A. (2021). A survey of preference estimation with unobserved choice set heterogeneity. *Journal of Econometrics*, 222(1, Part A):4–43.

Dietrichson, J., Ellegård, L. M., and Kjellsson, G. (2020). Patient choice, entry, and the quality of primary care: Evidence from Swedish reforms. *Health Economics*, 29(6):716–730.

Fack, G., Grenet, J., and He, Y. (2019). Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions. *American Economic Review*, 109(4):1486–1529.

Gaynor, M., Propper, C., and Seiler, S. (2016). Free to Choose? Reform, Choice, and Consideration Sets in the English National Health Service. *American Economic Review*, 106(11):3521–3557.

Godager, G. (2012). Birds of a feather flock together: A study of doctor–patient matching. *Journal of Health Economics*, 31(1):296–305.

Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers When Prices Are Negotiated: Evidence from the Hospital Industry. *American Economic Review*, 105(1):172–203.

Gutacker, N., Siciliani, L., Moscelli, G., and Gravelle, H. (2016). Choice of hospital: Which type of quality matters? *Journal of Health Economics*, 50:230–246.

Haastrup, P., Møller, A., Kristensen, J. K., and Huibers, L. (2025). Danish primary care: A focus on general practice in the Danish healthcare system. *Scandinavian Journal of Primary Health Care*, 43(4):839–845.

Ho, K. (2006). The welfare effects of restricted hospital choice in the US medical care market. *Journal of Applied Econometrics*, 21(7):1039–1079.

Huang, S. and Ullrich, H. (2024). Provider effects in antibiotic prescribing: Evidence from physician exits. *Journal of Human Resources*.

Nørøxe, K. B., Pedersen, A. F., Bro, F., and Vedsted, P. (2018). Mental well-being and job satisfaction among general practitioners: A nationwide cross-sectional survey in Denmark. *BMC Family Practice*, 19:130.

Pedersen, L. B., Kjær, T., Kragstrup, J., and Gyrd-Hansen, D. (2012). General practitioners' preferences for the organisation of primary care: A discrete choice experiment. *Health Policy*, 106(3):246–256.

Raval, D., Rosenbaum, T., and Wilson, N. E. (2022). Using disaster-induced closures to evaluate discrete choice models of hospital demand. *The RAND Journal of Economics*, 53(3):561–589.

Romley, J. A. and Goldman, D. P. (2011). How Costly is Hospital Quality? A Revealed-Preference Approach. *The Journal of Industrial Economics*, 59(4):578–608.

Santos, R., Gravelle, H., and Propper, C. (2017). Does Quality Affect Patients' Choice of Doctor? Evidence from England. *Economic Journal*, 127(600):445–494.

Simonsen, M., Skipper, L., Skipper, N., and Thingholm, P. R. (2021). Discontinuity in care: Practice closures among primary care providers and patient health care utilization. *Journal of Health Economics*, 80:102551.

Sivey, P. (2012). The effect of waiting time and distance on hospital choice for English cataract patients. *Health Economics*, 21(4):444–456.

Wager, S. and Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

A Additional Results for Preference Estimates

Table A.1: Preference estimates: all switchers and voluntary switchers

	All switchers		Voluntary switchers	
	(1) All clinics	(2) Open clinics	(3) All clinics	(4) Open clinics
Distance	-0.689	-0.733	-0.636	-0.665
	(0.002)	(0.002)	(0.002)	(0.002)
Distance squared	0.031	0.032	0.030	0.031
	(0.000)	(0.000)	(0.000)	(0.000)
Distance x Patient age	-0.002	-0.002	-0.002	-0.002
	(0.000)	(0.000)	(0.000)	(0.000)
Distance x Female patient	-0.005	-0.007	-0.012	-0.015
	(0.001)	(0.001)	(0.001)	(0.001)
Distance x Immigrant patient	0.009	0.004	0.015	0.013
	(0.001)	(0.002)	(0.002)	(0.002)
Clinic size	0.094	0.048	0.242	0.213
	(0.001)	(0.001)	(0.001)	(0.002)
Physician age	-0.078	-0.083	-0.039	-0.045
	(0.000)	(0.000)	(0.000)	(0.000)
Share of female physicians	-0.167	0.348	-0.360	0.085
	(0.003)	(0.004)	(0.005)	(0.006)
Physicians with immigration background	-0.039	-0.753	0.161	-0.644
	(0.005)	(0.005)	(0.006)	(0.008)
Gender match	0.222	0.384	0.404	0.535
	(0.004)	(0.006)	(0.006)	(0.007)
Immigration background match	0.422	0.424	0.349	0.382
	(0.007)	(0.009)	(0.009)	(0.011)
Cases	1003707	882922	516344	445484
Observations	103680207	18761618	56636159	9853215

Notes: Cases refer to patient × period choice cases. Observations refer to the total number of alternatives available to each choice case. All clinics include all primary care clinics within 15 km (5 km if the patient is in Copenhagen) as available alternatives. Open clinics restrict the choice set to open-list clinics within 15 km. Standard errors in parentheses.

Table A.2: Preference estimates: closure switchers and active-at-closure switchers

	Closure switchers		Active-at-closure	
	(1) All clinics	(2) Open clinics	(3) All clinics	(4) Open clinics
Distance	-0.776	-0.841	-0.872	-1.011
	(0.002)	(0.002)	(0.005)	(0.008)
Distance squared	0.030	0.031	0.037	0.040
	(0.000)	(0.000)	(0.000)	(0.000)
Distance x Patient age	-0.001	-0.001	-0.001	-0.002
	(0.000)	(0.000)	(0.000)	(0.000)
Distance x Female patient	-0.007	-0.010	-0.001	-0.005
	(0.001)	(0.001)	(0.003)	(0.005)
Distance x Immigrant patient	-0.038	-0.071	-0.014	-0.015
	(0.002)	(0.002)	(0.005)	(0.007)
Clinic size	-0.077	-0.182	0.279	0.263
	(0.001)	(0.001)	(0.003)	(0.004)
Physician age	-0.120	-0.128	-0.025	-0.037
	(0.000)	(0.000)	(0.000)	(0.001)
Share of female physicians	0.034	0.672	-0.399	0.163
	(0.003)	(0.005)	(0.011)	(0.014)
Physicians with immigration background	-0.192	-0.838	0.357	-0.592
	(0.004)	(0.006)	(0.013)	(0.017)
Gender match	0.059	0.243	0.277	0.443
	(0.005)	(0.006)	(0.013)	(0.018)
Immigration background match	0.474	0.522	0.318	0.312
	(0.008)	(0.010)	(0.021)	(0.028)
Cases	975412	875464	101563	89054
Observations	94175498	17822845	12124897	1769357

Notes: Cases refer to patient × period choice cases. Observations refer to the total number of alternatives available to each choice case. All clinics include all primary care clinics within 15 km (5 km if the patient is in Copenhagen) as available alternatives. Open clinics restrict the choice set to open-list clinics within 15 km. Standard errors in parentheses.

B Regression Analysis of Welfare Heterogeneity

Table B.1: OLS: expected distance gap on patient and market characteristics

	(1)	(2)	(3)	(4)
Patient age	0.0014	0.0008	0.0004	0.0041
	(0.0001)	(0.0001)	(0.0001)	(0.0132)
Female patient	0.0122	0.0154	0.0136	
	(0.0044)	(0.0044)	(0.0032)	
Patient immigration background	-0.1726	-0.1664	0.0033	
	(0.0051)	(0.0051)	(0.0037)	
Past care utilization	0.0008	-0.0003	-0.0003	-0.0009
	(0.0002)	(0.0002)	(0.0002)	(0.0006)
Number of clinics	0.0012	0.0013	0.0055	0.0040
	(0.0000)	(0.0000)	(0.0000)	(0.0001)
Number of open clinics	-0.0348	-0.0328	-0.0298	-0.0320
	(0.0001)	(0.0001)	(0.0002)	(0.0003)
Closure switch		0.2046	0.1391	0.1609
		(0.0077)	(0.0079)	(0.0160)
Closure x Past care utilization		0.0041	0.0013	0.0011
		(0.0004)	(0.0003)	(0.0007)
Closure x Number of clinics		-0.0001	-0.0006	-0.0003
		(0.0000)	(0.0000)	(0.0001)
Closure x Number of open clinics		-0.0041	0.0001	-0.0012
		(0.0002)	(0.0001)	(0.0003)
Observations	1787024	1787024	1786226	538430
R-squared	0.06	0.07	0.48	0.69
Within R-squared			0.0055	0.0330
Region x time FE	No	No	Yes	No
Patient FE	No	No	No	Yes
Time FE	No	No	No	Yes
Outcome mean	1.12	1.12	1.12	1.04

Notes: Outcome expected distance gap is in kilometers. Closure switch (0/1) refers to patients who switch to a new clinic because their originally listed clinic has closed. Fixed effects included are Danish municipality (*Kommune*) interacted with year-quarter. Standard errors in parentheses are clustered at the patient level.

Table B.2: OLS: expected distance gap, supplementary specifications

	(1)	(2)	(3)
Patient age	0.0006	0.0005	0.0005
	(0.0001)	(0.0001)	(0.0001)
Female patient	0.0122	0.0130	0.0124
	(0.0032)	(0.0032)	(0.0032)
Patient immigration background	0.0019	0.0023	0.0031
	(0.0037)	(0.0037)	(0.0037)
Past care utilization	0.0001	-0.0004	0.0000
	(0.0002)	(0.0002)	(0.0001)
Number of clinics	0.0052	0.0053	0.0053
	(0.0000)	(0.0000)	(0.0000)
Number of open clinics	-0.0302	-0.0290	-0.0296
	(0.0002)	(0.0002)	(0.0002)
Assigned after closure	-0.1203		
	(0.0080)		
Assigned x Past care utilization	-0.0011		
	(0.0003)		
Assigned x Number of clinics	-0.0004		
	(0.0000)		
Assigned x Number of open clinics	0.0057		
	(0.0002)		
Takeover after closure		0.2009	
		(0.0094)	
Takeover x Past care utilization		0.0022	
		(0.0003)	
Takeover x Number of clinics		-0.0005	
		(0.0000)	
Takeover x Number of open clinics		-0.0031	
		(0.0002)	
Active choice at closure			0.2124
			(0.0154)
Active x Past care utilization			0.0016
			(0.0007)
Active x Number of clinics			-0.0004
			(0.0001)
Active x Number of open clinics			-0.0011
			(0.0003)
Observations	1786226	1786226	1786226
R-squared	0.48	0.48	0.48
Within R-squared	0.0054	0.0058	0.0054
Region x time FE	Yes	Yes	Yes
Patient FE	No	No	No
Time FE	No	No	No
Outcome mean	1.12	1.12	1.12

Notes: Outcome expected distance gap is in kilometers. Takeover switch (0/1) refers to patients who stay at new clinics that have taken over their original clinics. Assigned switch (0/1) refers to patients who are administratively reassigned to a clinic when their original clinics have closed. Active switch (0/1) are patients who switch voluntarily at clinic closure. Fixed effects included are Danish municipality (*Kommune*) interacted with year-quarter. Standard errors in parentheses are clustered at the patient level.