



Supply-side variation in the use of emergency departments[☆]

Dan Zeltzer^{a,*}, Liran Einav^{b,c}, Avichai Chasid^d, Ran D. Balicer^{e,f}

^a Berglas School of Economics, Tel Aviv University, Tel Aviv 6997801, Israel

^b Department of Economics, Stanford University, Stanford, CA, United States

^c NBER, Cambridge, MA, United States

^d Clalit Research Institute, Clalit Health Services, Tel Aviv, Israel

^e Clalit Research Institute, Clalit Health Services, Tel Aviv, Israel

^f Department of Epidemiology, Faculty of Health Sciences, Ben Gurion University, Beersheba, Israel



ARTICLE INFO

Article history:

Received 3 May 2020

Received in revised form 11 March 2021

Accepted 11 March 2021

Available online 14 March 2021

JEL classification:

I11

Keywords:

Regional variation

Emergency department

Healthcare utilization

ABSTRACT

We study the role of person- and place-specific factors in explaining geographic variation in emergency department utilization using detailed data on 150,000 patients who moved regions within Israel. We document that about half of the destination-origin differences in the average emergency department utilization rate across districts translates to the change (up or down) in movers' propensity to visit the emergency department. In contrast, we find no change in the probability of having a hospital admission through the emergency department. Similar results are obtained in a complementary event study, which uses hospital entry as a source of variation. The results from both approaches suggest that supply-side variation in emergency department access affects only the less severe cases—for which close substitutes likely exist—and that variation across emergency physicians in their propensity to admit patients is not explained by place-specific factors, such as differences in incentives, capacity, or diagnostic quality.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Admitting a patient to the hospital from the emergency department (ED) is one of the more expensive routine healthcare decisions (Sabbatini et al., 2014). It directly influences the probability of hospital stays, which account for a large share of total healthcare spending.

Both ED utilization and admissions through the ED exhibit large variation over time and across space, and the sources of that variation remain unclear. To better understand this variation, it is important to acknowledge the

interrelated nature of patients' decisions to visit the ED and emergency physicians' decisions as to which ED patients to admit. Either of these margins may exhibit regional variation. The decision to visit the ED for urgent care may be influenced by place-specific factors, such as distance to the nearest ED and the presence of nearby urgent care centers and retail clinics (Alexander et al., 2019; Allen et al., 2019). The decision to admit patients from the ED may be influenced by place-specific factors, such as the distribution of cases arriving at the ED, hospital diagnostic quality, capacity, and incentives for filling unoccupied beds. Both decisions may also be driven by patient-specific factors, such as preferences for walk-in care over scheduled care and population morbidity that affects the need for urgent and emergent care. Patient-specific factors generally correspond to what we think of as “demand” and place-specific factors to what we think of as “supply.”

In this paper, we analyze the relative importance of supply- and demand-side factors in driving ED use and ED

[☆] We thank the Editor and two anonymous referees for many useful comments, and Joseph Rashba for excellent research assistance. Dan Zeltzer acknowledges funding from the Foerder Institute for Economic Research at Tel Aviv University.

* Corresponding author.

E-mail addresses: dzeltzer@tauex.tau.ac.il (D. Zeltzer), leinav@stanford.edu (L. Einav), avichaiha@clalit.org.il (A. Chasid), rbalicer@clalit.org.il (R.D. Balicer).

admission decisions. Separating demand and supply factors is an important first step for both understanding the nature of the variation in such use and for assessing the need for and potential impact of policies aimed at changing the variation.

To implement this exercise, we study the response to changes in ED access using two complementary empirical approaches. Our primary approach uses the variation in ED utilization across space. To separate the variation in ED utilization that is due to place-specific factors from variation due to patient-specific factors, we use within-patient variation in ED utilization as patients move, arguably for non-healthcare reasons, across regions that vary in their access to EDs.¹ In a secondary approach, we use an event-study design, which takes advantage of the entry of a new hospital (and thus a new ED) into a large market, thus increasing ED access sharply over time.

We implement our empirical strategy using detailed and comprehensive longitudinal medical data from about half of the Israeli population. The data come from Clalit Health Services, the largest of four Health Maintenance Organizations (HMOs) in Israel that provide universal, standardized, tax-funded health insurance to all residents. We study migration across Israel's nine districts, which are geographic units comparable to US counties, with an average population of about one million residents per district and a fairly wide range of annual ED use rates (the share of covered individuals who visit the ED in a given year), ranging from 20% to over 30%. ED visits account directly for 3% of all healthcare spending, and indirectly, as a portal for unplanned admission, for more than a quarter of total spending. Our main sample consists of all 150,000 adult beneficiaries (aged 25 and older) who moved across districts between 2011 and 2017.

The use of the Israeli setting is quite attractive for this purpose. While health insurance coverage broadly resembles that of Medicare in the United States, annual churn rates are lower than 1%, so we can track the entire enrollee population over many years, with minimal attrition. Moreover, Clalit provides uniform coverage and fairly homogeneous clinics across the country, so health insurance coverage does not change when patients move, and unobservable provider characteristics are less likely to drive the nature of variation across districts. Finally, movers in our sample are relatively young, thus mitigating concerns that may arise with an older population (e.g., Medicare) that moves might be driven by changes in health status.

Moving across districts is disruptive and likely affects short-run patterns of healthcare utilization. For this reason, as is now standard in the literature, our primary empirical strategy does not compare movers to non-movers but instead relies on estimating the extent to which the change in movers' ED use around the time of a move correlates with the destination-origin difference in ED use. Intuitively, if all ED use is driven by place-specific factors, we expect it to be perfectly correlated with this difference. Conversely, if it is

entirely driven by individual demand, we expect this correlation to be zero. Because we observe both moves from high to low ED-use areas and vice versa, we can separately and flexibly control for trends in ED utilization that are common to all movers and thus account for any disruption associated with the move itself.

Our main results reveal that the destination district's ED use rate has an important impact on the change in ED utilization of the mover. After moving to a district with an annual ED use rate that is one percentage point higher (lower) than the origin district, movers are about half a percentage point more (less) likely to visit the ED per year. Remarkably, all of the correlation between an individual ED use rate and the origin-destination difference in average ED use rates is concentrated in ED visits that did not result in admission to the hospital. In contrast, for visits that resulted in an admission, we find no correlation between individual rates and the destination-origin difference in ED utilization rate.

Taken together, the results suggest that at least half of the geographic variation in ED use is attributable to supply-side factors, but that such factors influence only cases that are discharged home, which are likely less severe. Conversely, the entire variation in ED visits that result in a subsequent hospital admission is attributable to demand-side factors. Using data on non-ED healthcare utilization, we present additional evidence consistent with the fact that these incremental low-severity ED visits are associated with a reduction in the utilization of other, presumably substitutable healthcare services.

We exploit the fact that in the Israeli context we observe comprehensive data on the affiliation of patients and providers, as well as detailed data on patients, to document that distance from the patient's "home clinic" to the nearest hospital and the propensity of primary care providers to refer patients to the ED are important channels that drive our main results. Both channels change upon a move, and (as we show) impact mainly outpatient ED use: patients living (or moving) closer to a hospital, or those with (or switching to) a high-propensity-to-refer physician are more likely to visit the ED, but not more likely to be admitted from the ED.

In a complementary approach, we exploit a distinctly different, albeit imperfect source of variation in ED access: the impact on ED use and admissions that results from the opening of a new general hospital in a previously underserved city. Like patient migration, hospital entry also leads to a sharp change in supply-side factors because it significantly reduces the distance and travel time to the nearest ED for the "treated" population of nearby residents (while also improving access to other inpatient services). We use a difference-in-differences framework to account for potential time trends by comparing outcomes of the treated population against outcomes of "untreated" patients residing in similar cities that did not experience a change in ED availability. Conforming with our earlier results, we find that an ED opening that occurs when a general hospital enters a local market results in a sharp increase in the monthly probability of an ED visit of nearby residents compared to residents of comparable cities. Furthermore, once again, the increase in ED use following the supply-side

¹ Our approach is similar to that used by Finkelstein et al. (2016a, 2018, 2019), Cutler et al. (2019), Agha et al. (2019) among others.

change is entirely concentrated in visits that do not result in hospital admission.

Taken together, the results consistently suggest that supply-side variation in ED access affects only the less severe ED cases—for which alternative settings might be a good substitute—but not the more severe cases that result in an admission.² In contrast, our finding that the entire effect is driven by ED visits that are less severe suggests that prior evidence of large variations in admissions through the ED may reflect mostly demand-side, not supply-side, factors. Specifically, our evidence militates against explanations of variation in ED admissions in terms of factors that vary geographically. In particular, the findings do not fit with the idea that the mere decision to visit the ED is an important driver of inpatient hospital utilization rates and the findings do not lend support to the idea that admission decisions (from the ED) depend on hospital differences in available capacity, diagnostic quality, or financial incentives.

A large body of existing work has documented substantial variation in ED use and admissions, both across and within countries. For example, [Shoff et al. \(2018\)](#) find that the ED admission rate of Medicare beneficiaries in 2012 varied tremendously across US counties; [Abualenain et al. \(2013\)](#) and [Sabbatini et al. \(2014\)](#) document large variation in hospital-level admission rates; [Dawson et al. \(2008\)](#) document variation across Canadian provinces; and [Van Parys \(2016\)](#) documents variation in admission rates among ED physicians in Florida. By separating supply- and demand-side factors, our analysis extends our understanding of the potential nature of this variation. Indeed, in the concluding section, we discuss the potential policy implications of our results.

The rest of the paper proceeds as follows: Section 2 presents our data source and setting. Section 3 presents our main sample and empirical strategy using patient migration. Section 4 presents our main results decomposing ED utilization rates to place and person effects, and Section 5 explores some channels that drive these results. Section 6 presents additional results analyzing the impact of hospital entry on ED use. Section 7 concludes and discusses some potential policy implications.

2. Setting and data

2.1. Institutional background

Our data come from Clalit Health Services, one of Israel's four non-profit HMOs that provide universal tax-funded healthcare coverage from birth onward to all Israeli residents, in accordance with the National Health Insurance Law (1995). Coverage broadly resembles that of Medicare parts A, B, and D, and includes hospital admissions, outpatient services, physician consults, drugs, and durable

medical equipment. The set of services covered under the universal coverage tier is reviewed and expanded every year by a professional committee that ranks new technologies to match a predetermined budget increase. The universal coverage tier is fully subsidized (HMOs receive risk-adjusted capitated payments from the government for each individual they enroll) but includes copay rates that are updated annually by the Ministry of Health.³ ED visits, which are the focus of our analysis, require a copay, with certain exemptions.⁴ At the time of our observation period, the copay was NIS 900 (about \$240) during the day and NIS 200 (about \$55) between 1:00 a.m. and 6:00 a.m.⁵

All four HMOs offer identical universal coverage but use distinct provider networks to do so, with the exception of hospitals, which admit all Israeli patients regardless of their choice of HMO. All four HMOs also offer optional supplemental coverage for items not covered by the universal tier, such as second-opinion specialist consults and additional paramedical treatments for children. The supplemental coverage requires subscribers to pay an additional, regulated, age-rated premium. In principle, enrollees can switch HMOs every other month and maintain their universal coverage, but the annual switching rate is extremely low (less than 1%), so each HMO covers a very stable population of enrollees.

Clalit is the largest of these four HMOs, covering approximately 4.5 million enrollees of all ages, or just over one half of the entire Israeli population. It is an integrated provider and insurer, providing most of the services it finances, and reimbursing pre-authorized services purchased from external providers. Its enrollees are admitted to all of Israel's thirty general hospitals, eight of which Clalit directly owns and operates.⁶ Clalit employs over 11,000 physicians and 10,000 nurses, operates over 1500 clinics across the country, and provides multiple outpatient services. Most of Clalit's physicians are salaried. Enrollees also have access to urgent care centers (UCCs) that provide after-hours walk-in care for urgent but non-emergent conditions. Most UCCs are operated by external providers, charge a copay of about NIS 90, and close before midnight.

Clalit is organized around nine districts, which follow geographic boundaries and are all subject to general operational guidelines from Clalit management. Patients consume the majority of their healthcare within the boundaries of their home district. Specifically, patients use local

³ Copay rates differ by service, with a service-specific quarterly out-of-pocket maximum. There are no copays for inpatient services.

⁴ ED visits that result in admission are exempt. In addition, copay exemption applies for patients who were referred to the ED by a physician or for one of several specific emergent conditions, such as traffic accidents, work- or school-related injuries, new bone fractures, and injuries that require stitches.

⁵ Prices and payments are denoted in current New Israeli Shekels (which had an average exchange rate of 3.7 NIS per USD during the observation period).

⁶ Hospitals are reimbursed per diem, except for a set of procedures (such as surgeries), for which hospital reimbursement is based on a procedure-related grouping of services. Patients can also utilize services from external providers, which in non-emergent cases require pre-authorization. Our data include detailed claims information for all these services.

² The finding that supply-side variation explains a substantial share of ED utilization is consistent with prior evidence that suggests a scope for substituting ED with other types of urgent care providers in low severity cases, using research designs that are based on the variation in entry and opening hours of retail clinics and urgent-care centers ([Alexander et al., 2019](#); [Allen et al., 2019](#); [Llovera et al., 2019](#)).

services in about 95% of the cases for primary care visits and diagnostic tests, more than 80% for specialist visits, and more than 70% for outpatient and inpatient services.

2.2. Data

Clalit maintains detailed and comprehensive claim-level data associated with all the services it provides to its universe of enrollees. The data contain demographic information about enrollees, information that is typical of billing data in the United States, as well as additional information about tests and test results, sometimes referred to as electronic medical records. The data we use for this study are based on Clalit's claims data from 2010–2018.

For our analysis, we construct health utilization variables at the patient-year level. Our primary variable of interest is the propensity of ED visits and whether the visit resulted in the patient being discharged or admitted to the hospital. We exclude a small number of cases that resulted in death during the ED visit; this is not common (accounting for 0.016% of patient-year observations among movers), so grouping them with ED visits that result in admission does not change any of the results. We refer to ED visits in which the patient was discharged without being admitted to the hospital as *outpatient* ED visits. We construct additional health utilization variables, including healthcare spending for different categories of services and UCC utilization.⁷ The main spending categories are prescription drugs and inpatient services, which together account for two thirds of total spending.

2.3. Cross-sectional variation across districts

The focus of our study is the extent of geographic variation in ED use and how it may be related to healthcare outcomes. In this section, we describe this variation, which will then be used as a key input in our analysis in the next section. The geographic unit for this analysis is a district, which is the administrative unit around which Clalit's healthcare services are organized. There are nine such districts that cover the entire country of Israel, with an average population of one million residents per district, so they can be thought of as similar in size to counties in the United States. Unlike the US, all districts in Israel are roughly similar in size, with the exception of the Eilat district, which is much smaller.⁸

Similar to patterns documented in other countries, Israel's healthcare spending in general and ED use in particular exhibit marked geographic variation. Appendix Table A1 reports summary statistics that illustrate the variation in healthcare utilization across these districts. For each district, Panel A summarizes annual ED use. For the population aged 25 and above, the probability of an ED visit ranges from 20.1% to 30.8% across Israel's districts. The

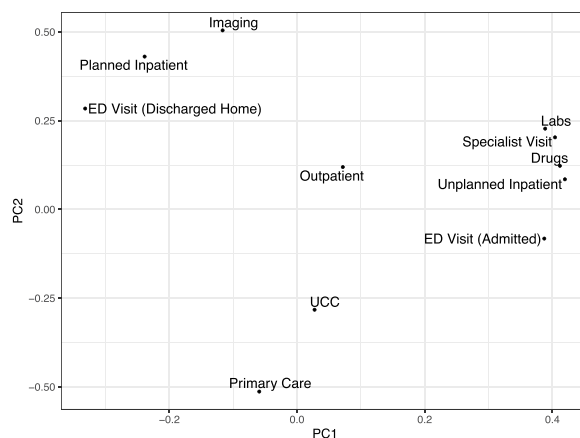


Fig. 1. Main principal components of cross-district ED use and spending. Notes: Figure shows, for the sample of three million covered adults, the loading of different ED use and spending categories on the two main principal components of between-district variation in ED use and spending on different service categories. The figure is based on the data on district means shown in Appendix Table A1 with variables standardized prior to the analysis. The two main principal components explain 75.1 percent of the total variation in shown measures. UCC stands for the fraction of residents with any urgent care center visit. ED Visit stands for the fraction of all residents with at least one emergency department visit, by visit outcome (home or hospital).

fraction of ED visits resulting in an admission also exhibits cross-sectional variation, ranging from 17.6% to 39.1% of ED visits across districts. Panel B summarizes spending by service category and shows a similar variation in spending across districts. Panel C summarizes statistics related to two potential channels: average distance to the nearest hospital and average referral rates to the ED by primary care physicians. We discuss these in detail in Section 5.

Fig. 1 summarizes the regional variation of different measures of ED use and spending. For each such measure, the figure shows its loading on each of two principal components of the regional variation in all measures shown. There are two noteworthy observations. First, there seems to be a clear distinction between districts with more unplanned admissions (through the ED), drugs, specialist visits, and labs on the one hand, and districts with heavier hospital use through planned inpatient visits and outpatient ED visits. Second, primary care and UCC services vary across districts and are unrelated to most services but are negatively correlated with outpatient ED use. These services load almost entirely on the second principal component.

Of course, it is difficult to draw sharp conclusions from descriptive evidence for variation in ED use and its correlation with other services, partly because it is unknown whether this variation reflects regional differences in the patient population's morbidity or preferences, or in provider access or choices. Next, we turn to decomposing this variation along these lines.

3. Empirical strategy and baseline sample

Our primary analysis relies on a “movers” research design (in the spirit of Finkelstein et al., 2016a), which ana-

⁷ Because we do not reliably observe negotiated prices for UCC visits, we only measure their utilization. UCC visits account for less than 1% of total spending.

⁸ Eilat is an Israeli port and resort town located in the very south of the country and is quite isolated from other population centers.

lyzes within-individual changes in healthcare utilization and outcomes before and after a move from one geographic unit (district in our context) to another. In this section we describe the empirical strategy and the baseline sample we assemble to take advantage of this research design.

3.1. Research design and empirical specification

To decompose the variation in average ED utilization across districts into a demand-side component that reflects regional differences in patient health or preferences and a supply-side component that reflects differences across places, we use patient migration. To the extent that migration is not driven by health status, studying patients who move locations provides an opportunity to study the ED utilization of the same patients in different regions.

Our main focus is an event-study analysis of changes in patient ED use associated with the timing and direction of moves. Consider a patient i who moved from origin district $o(i)$ to destination district $d(i)$. Let $\overline{ED}_{o(i)}$, $\overline{ED}_{d(i)}$ denote the average ED utilization in each district, which we estimate using the full population across all periods.⁹ We define the *destination-origin difference* in ED utilization for patient i as

$$\Delta ED_i = \overline{ED}_{d(i)} - \overline{ED}_{o(i)}.$$

We then estimate the following event-study equation:

$$Y_{it} = \alpha_i + \tau_y + \theta_t + \delta_t \Delta ED_i + X_{it} \beta + \varepsilon_{it}, \quad (1)$$

where Y_{it} is one of several measures of utilization of ED and other healthcare services for patient i in year t measured relative to $t = 0$, the year when i moved; α_i and X_{it} are patient fixed effects and controls for time-varying patient characteristics (ACG score and number of chronic conditions); τ_y are calendar year fixed effects; and θ_t represent indicator variables for the number of years before and after the move, which allow us to control flexibly for the (average) disruption in healthcare utilization associated with the move.

The parameter of interest is δ_t , which captures, for each period relative to the move, the relationship between mover i 's utilization and the difference in average ED visit rates between patient i 's destination and origin districts. If all of the regional variation in the outcome (e.g., ED use) is driven by patient-related factors, the estimated coefficients would be zero. If regional variation is fully driven by supply-side factors, the estimated coefficients would be one. An intermediate coefficient measures the extent to which regional variation can be attributed to patient-versus location-related factors.

The key to separate identification of person and place effects is the observed change in utilization when patients move. Our model permits movers to differ from non-movers in both their level of ED utilization (as fixed effects α_i are included) and in the trends of ED utilization around the move (as both δ_t and θ_t are allowed to vary with t).

⁹ For movers, we exclude the year of the move to avoid the need to partially attribute spending or utilization to the origin and destination districts.

The identification assumption is that any shocks to healthcare utilization that exactly coincide with the timing of the move are not correlated with the origin-destination differences in utilization. The assumption might be violated if patients who receive adverse health shocks move to high ED-utilization areas. The fact that movers are younger and healthier than the (relatively young) general population suggests that is not commonly the case. To the extent that such violations occur, it would cause us to overstate the role of place relative to patients.

To estimate the total effect of a move on the mover utilization, we aggregate the estimated period-by-period effects obtained using Eq. (1) to obtain the mean difference between post- and pre-move outcomes, weighted by sample size to account for the fact that our panel is imbalanced. We scale this estimated total difference in mover utilization by the mean absolute destination-origin difference in utilization. The details of this calculation are discussed in the appendix.

3.2. Baseline sample of movers

The Clalit data include the residential address of each enrollee, which is regularly updated by Israel's Population and Immigration Authority. We map each address to one of Clalit's nine districts and construct our baseline sample by including all adult patients aged 25 and above who moved across districts exactly once during the period 2011–2017. We study moves across—not within—districts, to exclude local moves in which access to ED and other healthcare services does not substantially change.

During the study period, out of the entire population of three million adult Clalit enrollees, 150,676 enrollees moved exactly once. These enrollees constitute our baseline “movers” sample. We exclude 78,932 individuals who moved twice or more during the study period as well as the majority of individuals who did not move across districts during the observation period. To guarantee that we observe utilization at least one year before and one year after each move, we include in this sample only enrollees with continuous coverage for at least one calendar year prior to the year of the move and for the year following the year of the move, as long as the individuals are alive. To satisfy this requirement for patients who moved in 2011 or 2017, we extract data from 2010 and 2018 as well. Given the low churn rate of enrollees, nearly all patients are observed throughout our nine-year study period (2010–2018), with the average patient being observed for more than eight years. The resulting (unbalanced) baseline sample has a total of 1.27 million enrollee-year observations. The rate of migration is stable over the study period (Appendix Table A2).

Movers are on average younger and slightly healthier than the overall adult population. Table 1 compares the average characteristics of movers in our sample and the population of non-movers, calculated over the entire study period. Panel A of this table compares demographic and health status. The average mover is 40 years old at the time of the move—ten years younger than the average adult Clalit enrollee who did not move. Movers have a higher socioeconomic status than non-movers, a lower number of

Table 1
Characteristics of movers and non-movers.

	Movers (1)	Non-movers (2)
<i>A. Demographics and health status</i>		
Age (mean)	40.6	50.8
Female (percent)	49.5	52.4
High SES (percent)	52.1	32.7
Number of chronic conditions (mean)	2.28	3.41
ACG high resource (percent)	6.8	12.4
Annual mortality rate (percent)	0.42	1.22
<i>B. Annual emergency department use</i>		
Percent with ED visit, by visit outcome		
Any outcome	22.5	24.3
Discharged home (outpatient ED)	18.8	18.8
Admitted (inpatient ED)	6.9	9.3
Died	0.016	0.057
Fraction of ED visits resulting in admission	25.8	34.3
Percent with UCC visit	5.5	5.4
<i>C. Annual spending, by category (mean, NIS)</i>		
Total (all categories)	4946	6689
Emergency department	200	190
Prescription drugs	1173	1441
Inpatient – unplanned	1222	1796
Inpatient – planned	670	1113
Outpatient services	564	862
Primary care visits	221	293
Specialist visits	243	266
Imaging services	275	302
Laboratory services	196	175
Other	175	243
Number of observations (patient-year)	1,274,445	20,085,457
Number of patients	150,676	2,723,196

Descriptive statistics comparing our main sample of patients who moved and patients who did not move during 2011–2017. Panel A shows demographic and health status information. High SES is the top tercile of zipcode-based socioeconomic status scores. ACG High Resource is a dummy for high predicted healthcare utilization. Panel B shows ED use—having any visits during the year and the number of visits conditional on any visits, by discharge status. Panel C shows healthcare utilization by service category. Non-mover spending was calculated using a random subsample of 100,000 patients. See Section 2.2 for detailed sample and variable definitions.

chronic conditions, and a lower annual mortality rate. The fact that movers are younger and healthier suggests that health is less likely to be the main reason for the move.¹⁰

Consistent with their being slightly younger and healthier, movers' annual probability of ED use is lower than non-movers'. Panel B of Table 1 compares the average ED utilization during the study period between the sample of movers and the non-mover population. Movers had a lower probability of any ED visit during the year (22.5% versus 24.3% for non-movers) and a lower rate of admission from

¹⁰ In a survey conducted in 2008 among a representative sample of the adult Israeli population, the top reported reasons for moving to a new apartment were to improve housing or school quality (27%); family reasons, including family expansion, marriage, and divorce (27%), switching from renting to owning a house or an apartment (19%), changes in personal economic circumstances (6%), and changes in employment (4%). Source: Central Bureau of Statistics, 2008 Social Survey. https://www.cbs.gov.il/he/publications/DocLib/2010/seker_hevrati08/pdf/h_print.pdf. Accessed February 2020.

the ED, conditional on visiting (25.8% versus 34.4% for non-movers). However, average ED utilization rates of movers prior to migration are highly correlated with ED use of non-movers in the origin districts: the correlation coefficient is 0.95 (0.90 when weighting by district size).

Movers also spend less, on average, than non-movers. Panel C of Table 1 shows average annual healthcare spending per patient by type of service for movers and for non-movers. The average annual healthcare spending of a patient in our sample of movers is NIS 4900, which is lower than the average adult enrollee of Clalit who did not move (NIS 6700 per year). The breakdown of spending is fairly similar between movers and non-movers, with movers devoting a slightly higher fraction of their total spending to drugs (24% versus 22% by non-movers) and a slightly lower fraction to inpatient services (38% versus 43%). Movers and non-movers spend a similar amount on ED services, which reflects a slightly higher share of movers' overall lower spending (4% versus 3%).

4. Results

4.1. Decomposing ED utilization using patient migration

Fig. 2 shows the distribution of the destination-origin differences in ED utilization rates for movers—the key right-hand-side variable in the event-study analysis. This distribution is centered at zero and is approximately symmetric, suggesting that moves from low to high ED-utilization districts are as common as moves from high to low. This is also confirmed by the matrix of the origin and destination of migrations (Appendix Table A3).

As a first look at the way ED use changes around moves, Fig. 3 plots the change in any ED use rate (the average share of years with ED use post-move minus the average share




Fig. 2. Destination-origin differences in ED utilization rates. Notes: The figure shows, for the sample of 150,676 movers, the distribution of the difference between each mover's origin and destination districts' average ED visit rate. Averages for each district are calculated for the entire population over the entire study period (excluding the year of the move for movers). An average move involves an absolute origin-destination difference in the ED use rate of 2.6 percentage points. See Appendix Table A1 for additional information on the distribution of ED visit rates across districts. See Appendix Table A3 for additional information on the distribution of origin and destination of moves.




Fig. 3. Change in movers' ED use by size of destination-origin difference in ED visit rates. *Note:* The figure shows the change in movers' ED use by the size of the destination-origin difference in average district ED visit rates. We group movers by deciles of ΔED , namely the destination-origin difference in average ED use. For each decile, the x-axis shows the average of ΔED and the y-axis shows the average post-move ED use (average annual probability of ED visit) minus the average pre-mover ED use. The slope shown is for the linear fit line of the ten plotted points. The sample is all movers ($N = 1,274,445$ patient-years).

of years with ED use pre-move) against the destination-origin difference in average ED use rates. If all geographic variation were due to supply-side factors, we would expect this plot to have a slope of 1. If all variation were due to demand-side factors, we would expect this plot to have a slope of 0. Fig. 3 shows that the slope is 0.45, suggesting roughly equal shares for both supply and demand factors. The relationship is symmetric above and below zero.

Fig. 4 shows our main event-study results, plotting the estimates of the coefficient δ_t from estimating Eq. (1) with the probability of an ED visit during a given year as a dependent variable. Results for all ED visits are shown in Panel A. It shows that there is a sharp discontinuous jump at the time of the move in the probability of any ED visit, from 0 to approximately 0.5.¹¹ Under the identification assumptions discussed in Section 3, this jump implies that half of all ED utilization is attributable to supply-side factors associated with the place of residence, with the remaining half attributable to demand-side factors associated with the patient.¹²

Panel B of Fig. 4 shows results separately by discharge status. Remarkably, supply-side factors that impact patient ED utilization only affect outpatient ED visits. We find no relationship at all between the rate of patients' ED visits

¹¹ While there is a small pre-trend, it is only statistically significant several years before the move, where we have fewer observations per year (because of our unbalanced panel). It is small and insignificant in years where we can more precisely estimate it.

¹² Recall that our specification flexibly controls for the time trend in a mover's ED use around the time of the move, denoted by θ_t . Appendix Figure A1 shows estimates of θ_t . Indeed, separately from the above discussed correlation with the destination-origin difference in average ED use, all moves are also associated with an increase in ED use, which may reflect the disruption in care continuity that is associated with a move.

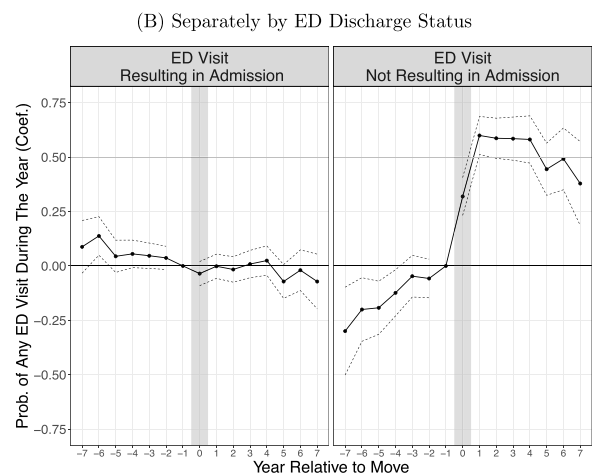
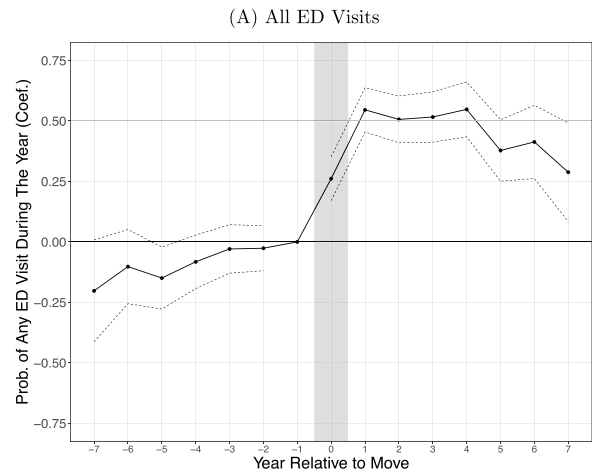


Fig. 4. Change in movers' ED use related to the destination-origin difference in ED visit rates. *Notes:* The figure shows event-study estimates of the fraction of movers' ED use that is associated with the destination-origin difference in ED visit rates, namely, δ_t from Eq. (1). Year 0 is the year of the move. The coefficient of the year before the move is normalized to 0. The sample is all movers ($N = 1,274,445$ patient-years).

that result in an admission and the place of residence; the probability of such visits is entirely attributable to patient factors. This finding further supports the identification assumption that migration is not driven by health status because if it were, as discussed in Section 3, we would expect coefficients to be upward biased. Furthermore, this evidence suggests that hospitals do not indiscriminately admit patients who visit the ED but rather that the decision to admit patients is made consistently across regions. In sum, all of the supply-side driven variation in ED utilization is concentrated within outpatient ED cases.

Table 2 scales our estimates and standard errors from Eq. (1) to show the change in movers' ED use that is related to an "average move"—a move across two hypothetical regions with a destination-origin difference in ED use equal to the sample average of the absolute value of such difference—and compare that change to the pre-move mean of each outcome. Column (1) shows the pre-move

Table 2
The change in ED use and in spending related to destination-origin difference in ED use.

	Pre-move mean	Change related to average absolute origin-destination difference in ED use	S.E.
	(1)	(2)	(3)
<i>A. Annual emergency department use</i>			
Any ED visit	0.207	0.0146	0.00080
Any ED visit not resulting in admission	0.175	0.0170	0.00075
Any ED visit resulting in admission	0.058	−0.0013	0.00048
Any UCC visit	0.065	−0.0323	0.00125
<i>B. Annual spending (by category)</i>			
Emergency department	162.5	21.0	1.3
Primary care visits	209.3	7.4	0.5
Imaging services	180.9	4.8	1.9
Other	128.1	−2.8	3.1
Prescription drugs	882.2	−9.7	28.5
Specialist visits	199.4	−9.9	0.8
Inpatient – planned	447.1	−11.6	17.6
Laboratory services	155.4	−11.7	1.6
Outpatient services	347.4	−12.0	17.6
Inpatient – unplanned	775.8	−94.6	23.5
Total (all categories)	3494.5	−122.2	54.5

Table shows estimates of the change in movers' ED use, following a move. Column 1 shows the average pre-move mean of each outcome. Column 2 shows estimates of the change in ED use that is related to an "average move"—a move across two hypothetical regions with a destination-origin difference in ED use equal to the sample average of the absolute value of such difference. These estimates were obtained using Eq. (1), collapsed to the difference between all post- and pre-periods (weighted by sample size) and scaled to show the effect on spending associated with the average absolute origin-destination difference in ED utilization in our sample, which is 2.6 percentage points. Details of these calculations are discussed in the Appendix. Column 3 shows estimated standard errors, clustered by patient. Panel A shows estimates for the probability of ED and UCC use. Panel B shows estimates for spending on different types of services. Spending is denominated in New Israeli Shekels (NIS). The sample is all movers ($N = 1,274,445$ patient-years).

means among movers.¹³ Columns (2) and (3) show a summary of the weighted post- minus pre-move difference in ED use related to an average move. Panel A shows estimates for ED use. The average move is associated with a change of 1.5 percentage points in the annual probability of any ED visit (which is approximately 7% of the pre-move baseline). The increase is concentrated in outpatient ED visits. We estimate a very small change in the probability of an ED visit resulting in admission (a −0.1 percentage point change with a standard error of 0.048; the pre-move baseline probability is 5.8%). Increased outpatient ED use is mirrored by a 3.23 percentage point decrease in the probability of any UCC visit (approximately a half of the pre-move baseline probability of 6.5%).

4.2. Other outcomes

Appendix Figure A2 and Appendix Figure A3 show additional event-study results relating patient spending—overall and by type of services—to destination-origin differences in ED use. Appendix Figure A2 shows estimates of Eq. (1) with spending in each service category on the left-hand side using the same destination-origin difference in ED visit rates, ΔED_i , on the right-hand side. Consistent with our main result, following a move, a patient's spending on ED becomes more similar to the

average spending in the destination area. In addition, we find that moving to an area with a higher (lower) average ED use rate is positively correlated with spending on primary care services and negatively correlated with spending on urgent care centers, specialist visits, and outpatient laboratory services.

Panel B of Table 2 shows the estimated change in utilization associated with the average place-related change in the ED utilization rate. Column (1) shows the average pre-move spending in each category. Column (2) shows the estimated weighted mean difference between post- and pre-move estimates shown in Appendix Figure A2, scaled by the average absolute destination-origin difference in ED use, and column (3) shows standard errors (details of this calculation are discussed in the Appendix). Appendix Figure A3 summarizes these same results visually. For each category, the left panel of Appendix Figure A3 shows the average annual pre-move spending among movers, and the right panel shows the estimated change in utilization associated with an average move.

Place-related differences in ED use are most strongly—and negatively—associated with spending on unplanned admissions. They are positively associated with spending on primary care and imaging. These findings are consistent with ED visits that complement primary care visits and substitute for office-based consults and diagnostics. Overall, moving to an area with higher (lower) ED use is associated with a decrease (increase) in spending overall, as lower spending on unplanned hospital admission offsets the increased spending on ED and complementary services.

We perform additional heterogeneity analyses whereby we re-estimate the event-study model specified in Eq. (1)

¹³ Note that pre-move mean outcomes for movers are lower than their mean outcomes over the entire study period due to both aging and time trends in utilization. For example, movers' average annual spending is NIS 3500 pre-move (Panel B of Table 2) and NIS 4900 over the entire study period (Panel C of Table 1). Our empirical specification flexibly accommodates such trends.

separately for moves from low to high ED-use districts and from high to low ED-use districts (Appendix Figure A4). We also separately estimate the same model for moves associated with above-median and below-median change in the absolute destination-origin difference in ED use (Appendix Figure A5). In both cases, we find that results hold similarly across subsamples.

5. Channels

In this section, we explore the importance of potential supply-side channels that may contribute to the geographic variation in ED use. We provide evidence in support of two such channels: distance to the nearest hospital and the level of complementarity between local physician services and ED services, as reflected by physicians' ED referral propensity. Our strategy does not provide a definitive quantification of these factors and is certainly not exhaustive, and we conclude by discussing additional potential contributors.

5.1. Distance

One potential channel that could affect the use of health-care services in general and hospital services in particular is distance. The distance between a patients' residence and the hospital has been previously shown to affect hospital choice in the expected direction (Kessler and McClellan, 2000; Gaynor and Vogt, 2003; Romley and Goldman, 2011; Beckert et al., 2012; Baker et al., 2016). In our context, a natural question is whether variation across space in the distribution of patients' distance to hospitals may contribute to geographic variation in ED utilization.

To explore this possibility, we calculate, for each patient in our sample, the distance between their assigned primary care clinic and the nearest hospital. Clalit assigns clinics based on the patient's current place of residence, and since there are over 1500 such clinics and only 30 hospitals, the distance we calculate reflects well the variation in a patient's residence relative to the nearest hospital. The average distance to the nearest hospital also varies considerably between districts, ranging between 3 and 18 km (see Panel C of Appendix Table A1). Panel (A) of Appendix Figure A6 further shows, for the movers sample, the distribution of pre-move distance to the nearest hospital. Two thirds of the patients in our sample have a hospital within 10 km and nine out of ten have a hospital within 20 km, but some have to travel as many as 80 km. Further, because patients who move change their location, our data provide a unique within-patient variation in residential distance to the nearest hospital. Panel (B) of Appendix Figure A6 shows the distribution of the change in this distance associated with a move in our sample, which is well dispersed and fairly symmetric.

Our strategy consists of three parts. First, we describe the relationship between the patient's distance to the nearest hospital, which we denote by D_{it} , and different measures of ED use, denoted as above by Y_{it} . Second, we describe the relationship between the patient's change in the distance to the nearest hospital associated with a move at $t = 0$, $\Delta D_i = D_{i,t>0} - D_{i,t<0}$, and the patient's change in

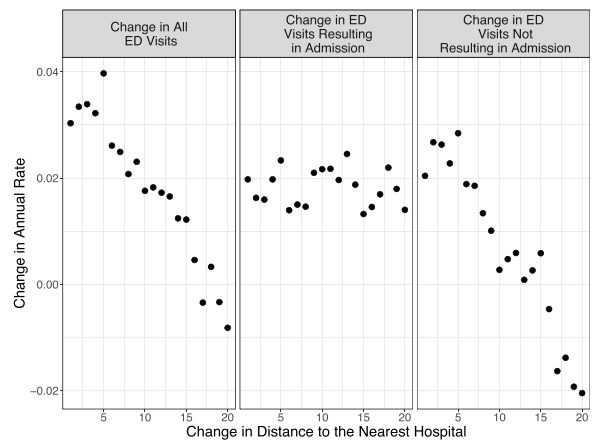


Fig. 5. Change in ED utilization and the change in distance to the nearest hospital around a move. *Notes:* The figure shows the relationship between the change in a mover's distance to the nearest hospital and the change in their ED use rates. The x-axis shows ventiles of the change in a mover distance to the nearest hospital: their distance to the nearest hospital before the move minus their distance to the nearest hospital after the move (see Section 5 for details). The y-axis shows, for each of these bins, the average change in ED use rates, defined as the average ED use rate in the years after the move minus the average ED use rate in the years before the move (the year of the move is excluded). The different facets show results for all ED visits, and separately by discharge status. The sample is all movers ($N = 1,274,445$ patient-years).

ED use around the time of the move, $\Delta Y_i = \text{Avg}_{t>0}(Y_{i,t}) - \text{Avg}_{t<0}(Y_{i,t})$.

Clearly, a patient's place of residence is endogenous and can potentially be correlated with the patient's health. Although this is arguably a smaller concern for the relatively young movers in our sample than it is for older populations, we nonetheless address this concern in a third specification that instruments for the endogenous change of distance over a patient's move using the difference in average distance between the origin and destination districts. $\bar{D}_{o(i)}$ and $\bar{D}_{d(i)}$ denote the average distance to the hospital across all residents at the origin and destination districts of patient i . We define the destination-origin difference in average distance to the hospital for patient i to be $\Delta \bar{D}_i = \bar{D}_{d(i)} - \bar{D}_{o(i)}$. Then, we estimate the following two-step least square model:

$$\begin{aligned} \Delta D_i &= \alpha \Delta \bar{D}_i + X'_{it} \delta + v_{it}, \\ \Delta Y_{it} &= \beta \Delta \hat{D}_i + X'_{it} \delta + \varepsilon_{it}, \end{aligned} \tag{2}$$

where X_{it} is a vector of patient controls. The parameter of interest is β , which captures the change in ED use associated with an increase in distance to the hospital induced by the move to an area where this distance is higher (lower) on average.

Appendix Figure A7 shows the relationship between distance to the hospital and different measures of ED utilization, before and after a move. There is a clear correlation between distance and ED use in the expected direction: patients living closer to a hospital are more likely to use the ED during a given year. Moreover, this correlation is entirely concentrated in ED visits that do not result in an admission. ED visits that result in admissions exhibit no correlation with distance. It is likely that decisions to visit

Table 3

IV estimates of the change in ED use associated with a change in distance to the nearest hospital and with a change in PCP referral propensity.

	Dependent variable			
	Change in ED visits (1)	Change in ED visits not resulting in admission (2)	Change in ED visits resulting in admission (3)	Change in UCC visits (4)
<i>A. Distance IV</i>				
Change in distance to nearest hospital	-0.02223*** (0.00155)	-0.02339*** (0.00146)	-0.00317*** (0.00098)	0.02891*** (0.00092)
<i>B. Referrals IV without controls for distance</i>				
Change in PCP referral propensity	0.01417*** (0.00082)	0.01616*** (0.00078)	-0.00037 (0.00050)	-0.01838*** (0.00052)
<i>C. Referrals IV with controls for distance</i>				
Change in PCP referral propensity	0.01420*** (0.00084)	0.01620*** (0.00080)	-0.00042 (0.00051)	-0.01840*** (0.00053)
Change in distance to nearest hospital	0.00085 (0.00103)	0.00127 (0.00097)	-0.00126** (0.00062)	-0.00044 (0.00065)
Pre-move mean	0.207	0.175	0.058	0.052
ACG resource utilization band	Yes	Yes	Yes	Yes
Number of chronic conditions	Yes	Yes	Yes	Yes
Age group	Yes	Yes	Yes	Yes
Gender	Yes	Yes	Yes	Yes

The table shows IV estimates of the change in a mover ED and UCC utilization associated with a change in their distance to the hospital on a move and, separately, with a change in their PCP propensity to refer to the ED. Panel A shows estimates of Eq. (2) for the change in a mover ED and UCC utilization associated with a change in their distance to the hospital on a move. Panels B and C shows estimates of Eq. (5) of for the change in a mover ED and UCC utilization associated with a change in their PCP propensity to refer to the ED, with and without controls for the patient's change in distance from the hospital. The different columns show estimates of the effect for different utilization measures. Change in Distance to the Nearest Hospital is the difference in the mover's distance to the nearest hospital before the move minus their distance to the nearest hospital after the move, in multiples of 10 km (see Section 5 for details). Primary care physicians' referral propensity is based on estimates of Eq. (3). We averaged these measures, separately for the years before and the years after the move, and calculated the change in this propensity, defined as the average propensity in all years after a move minus the average propensity in all years before a move. Change in PCP Referral Propensity is the ventile of these changes, scaled between 1 and 20. The first-stage estimates for Panel A, from regressing individual change in distance on destination-origin difference in average distance is 1.036 (s.e. 0.005); the F statistic is 1721. The first-stage estimates for Panel C, from regressing individual change in PCP referral propensity (ventile) on destination-origin difference in average such propensity is 0.00010 (s.e. 0.00003); the F statistic is 277. The sample for Panel A includes 149,262 movers, excluding a small number of movers for which location information is missing. The sample for Panels B and C includes 146,928 movers, excluding in addition a small number of movers whose physicians had fewer than 100 visits in 2018, and for which we did not have referral propensity estimates. Standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

the ED with high-severity conditions (that result in admissions) are not sensitive to variation in distance, whereas the decision to visit the ED for the diagnosis and treatment of less severe conditions is.

Fig. 5 shows the relationship between the *change* in distance to the nearest hospital and *change* in ED use following a move. Once again, we find a strong relationship between distance and ED use: moving closer to (farther from) a hospital is associated with an increase (decrease) in the probability of ED use. Yet again, this relationship is concentrated in ED visits not resulting in an admission; the change in the probability of an ED admission is unrelated to the change in the patient's distance to the hospital.¹⁴

Panel A of Table 3 shows the results of estimating the 2SLS model specified in Eq. (2). A 10 km increase in the distance to the nearest hospital (induced by a move across districts) is associated with a 2.2 percentage point decrease

in the annual probability of an ED visit. Nearly all of this change reflects a change in outpatient ED admissions not resulting in an admission, which decrease by 2.3 percentage points. The probability of a patient being admitted through the ED in a given year decreases by only 0.3 percentage points. Overall, these findings suggest that spatial variation in the distribution of patients and hospitals, and consequently in a patient's distance to the nearest hospital, contribute to the geographic variation in ED use. The remaining panels of Table 3 show results of the analysis of the impact of physician referrals on ED use, which we discuss next.

5.2. Physician referrals

Referrals are an important source for ED visits. Office-based physicians are reported to be making growing use of EDs to perform complex workups and expedite non-elective admissions, and EDs are increasingly supporting primary care practices by performing diagnostic workups and handling after-hours demand for care (Morganti et al., 2013). In the United States, referral rates have increased over time (Barnett et al., 2012).

¹⁴ Regardless of the change in distance to the hospital, patients are 2% more likely to be admitted through the ED post-move, relative to pre-move. This reflects a combination of aging (patients are mechanically several years older during the post-move period) and the fact that ED use increases uniformly following a move, in addition to the association with destination-origin difference in average ED use.

To explore whether physician referrals contribute to geographic variation in ED use, we proceed as follows. First, we estimate the propensity of individual physicians to refer patients to the ED. Second, exploiting the fact that more than 95% of movers change their primary care physician after they move, we study the relationship of the change in a mover's primary care physician propensity to refer to the ED and the mover's changed ED use around a move. Finally, we relate the variation in physician propensity to refer and regional variation in ED utilization using an instrumental variable approach similar to the one in Eq. (2).

To estimate the variation in individual physicians' propensity to refer to the ED while controlling for the potential variation in their case mix, we extract an auxiliary sample of 19 million visits to 4200 primary care physicians by (mover and non-mover) Clalit patients in 2018 and use the following fixed effects model to estimate physicians' propensity to refer patients to the ED, conditional on case characteristics¹⁵:

$$Referral_{ij} = X_i \delta + \eta_j + \varepsilon_{ij} \quad (3)$$

where i indexes patients and j primary care physicians, $Referral$ is a dummy for whether the physician referred the patient to the ED, and X_i is a set of controls for case mix, which includes the number of chronic conditions that the patient has, future healthcare utilization predicted by Johns Hopkins ACG classifier (see footnote 19 for details), and the patient gender and age. Appendix Table A4 shows descriptive statistics for the auxiliary sample used in this estimation. Fig. 6 shows the distribution of the estimated fixed effect, $\hat{\eta}_j$, revealing substantial variation in the propensity of different physicians to refer to the ED, even conditional on case mix. This variation exists not just within regions, but also between them. Appendix Table A1 (Panel C) shows that average physician propensity to refer ranges between 1.9 and 3.9 across districts. Using these estimates of PCP propensity to refer to the ED, we then estimate the following model:

$$\Delta Y_i = \kappa_1 \Delta Q_{j(i)} + \gamma \Delta D_i + X'_i \delta + \varepsilon_i \quad (4)$$

where ΔY_i is defined as above, $\Delta Q_{j(i)}$ is the difference in the quantile of the patient's primary care physician' referral propensity (η_j from Eq. (3)) before and after the move, and ΔD_i is a control for the difference in the distance to the nearest hospital from before to after the move.¹⁶ The parameter of interest is κ_1 , which measures how the mover's change in ED utilization around the move is associated with the change between the mover's old and new primary care physician's propensity to refer to the ED.

Finally, to examine the contribution of primary care physicians to the regional variation in ED use and to account for the potential endogeneity of the choice of physician

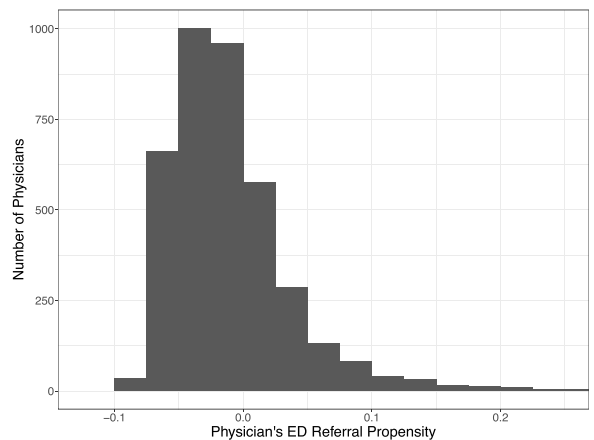


Fig. 6. The distribution of physician propensity to refer to the ED, adjusting for case characteristics. *Notes:* The figure shows the distribution of primary-care physician propensity to refer cases to the ED, captured by the physician fixed effect estimated using Eq. (3), which adjusts for case characteristics. The underlying sample used in the estimation contains 19 million primary care visits of 4200 doctors in 2018. The sample and method are described in more detail in Section 5.

over a move, we estimate the following two-step least squares model, similar to Eq. (2):

$$\begin{aligned} \Delta Q_{j(i)} &= \alpha_1 \Delta \bar{Q}_i + \gamma_1 \Delta D_i + X'_{it} \delta_1 + v_{it}, \\ \Delta Y_{it} &= \kappa_2 \Delta \hat{Q}_{j(i)} + \gamma_2 \Delta D_i + X'_{it} \delta_2 + \varepsilon_{it}, \end{aligned} \quad (5)$$

where $\Delta \bar{Q}_i = \bar{Q}_{d(i)} - \bar{Q}_{o(i)}$ is the difference between the average propensity to refer to the ED of *all* physicians in patient i 's destination and origin districts. The parameter of interest here is κ_2 , which captures the change in ED utilization associated with the change in a mover's PCP referral propensity that is correlated with the destination-origin difference in average local physicians' propensity to refer. A positive κ_2 implies that variation in physicians' propensity to refer to the ED contributes to regional variation in ED use.

Appendix Table A5 shows estimates of Eq. (4). A greater propensity of a mover's new physician, relative to the mover's old physician, to refer to the ED is significantly associated with an increase in the mover's annual probability of an ED visit. This change is concentrated in outpatient ED visits. In contrast, the probability of an admission through the ED is not related to the change in the mover's primary care physician's propensity to refer to the ED. Panels B and C of Table 3 show IV estimates from Eq. (5) for the impact of a change in a mover's PCP referral propensity on different outcomes with and without controls for the change in the patient's distance to the nearest hospital. Even when instrumented by the average change in PCP referral, a switch toward a physician with greater propensity to refer to the ED is associated with a significant increase in the mover's ED utilization. Once again, the effect is fully concentrated in outpatient ED visits. There is no effect of the mover's (instrumented) physician propensity to refer to the ED on the annual probability of an ED admission.

¹⁵ We restrict attention to physicians with at least 100 visits in 2018, resulting in exclusion of a small number of physicians.

¹⁶ That is, $\Delta Q_{j(i)} = Q(\hat{\eta}_{j(i)_{post}}) - Q(\hat{\eta}_{j(i)_{pre}})$, where $j(i)_{pre}$ and $j(i)_{post}$ are patient i 's primary care physician before and after i 's move, and $Q(\hat{\eta}_j)$ denotes the quantile of $\hat{\eta}_j$.

Comparing the different panels of [Table 3](#) reveals that once physicians' propensity to refer to the ED is accounted for, a change in distance has little impact on outcomes. These findings suggest that changes in distance to hospital matter primarily through their impact on the relationships between providers rather than directly through the patient likelihood of visiting a closer ED.

Taken together, variation in physician propensity to refer non-severe cases to the ED contributes to geographic variation in ED use. More broadly, evidence from physician referrals suggests that in some areas, primary care services, which are often viewed as substitutes for ED services, are in fact complements to them. Such complementarity is consistent with prior findings from the Oregon health insurance expansion showing that insurance coverage increased both primary care use and ED use ([Taubman et al., 2014](#)), and that such increases were persistent ([Finkelstein et al., 2016b](#)).

5.3. Other factors

While office-based visits are complementary to ED visits, other urgent care services are potential substitutes.¹⁷ Therefore, spatial variation in the availability of UCCs may contribute to the variation in ED use. We document two findings consistent with this explanation. First, in our main results, shown in [Table 2](#) and discussed in [Section 4](#), we find that an increase in outpatient ED use is associated with a destination-origin difference in average ED use that is mirrored by a decrease in the probability of any UCC visit. A similar pattern also emerges from our analysis of the change in the patient's distance to the nearest hospital associated with a move. Column (4) of Panel A of [Table 3](#) shows that an increase in the distance to the hospital is associated with an increase in the probability of using an urgent care center, mirroring the decrease in ED use as a result of the same move documented in Column 3 of that table. These findings are consistent with the supply of urgent care centers being a contributing factor for variation in (outpatient) ED use.

We emphasize that the above factors do not exhaust the potential supply-side factors that drive the observed variation in ED use. Other factors that may also contribute to this variation include: the distance to other services, the number of primary care physicians and specialists per capita, and ED wait times (to the extent they are not entirely demand driven), among others.

6. Decomposing ED utilization via hospital entry

In this section, we complement the previous analysis with an event-study design, in which we estimate the contribution of supply- versus demand-side factors to ED utilization using an alternative supply-side shock, or

¹⁷ For example, using the closing time of urgent care centers, [Allen et al. \(2019\)](#) show that when UCCs close, ED use increases. [Myong et al. \(2020\)](#) find that increased funding for federally qualified health centers is associated with reduced ED use. [Alexander et al. \(2019\)](#) find that opening of retail clinics is associated with a decrease in nearby ED use.

“treatment”: the opening of a new general hospital, which sharply improved ED access for nearby residents.

We acknowledge that the research design of the event study we report on in this section is imperfect. The opening of the new hospital is surely driven by increasing latent demand for hospital services. It seems plausible that one can attribute sharp changes in ED utilization around the entry to supply-side changes, but any estimated longer-run effect may be confounded by changes in underlying demand.

Yet, despite these concerns, we feel that the analysis provided in this section is useful. It relies on a very different setting, using non-movers rather than movers, and the event of hospital entry is correlated with sharp changes in a host of other hospital services (beyond ED). Despite all these differences (and the imperfect research design), the qualitative results in this section are quite similar to the results in our primary analysis; thus, in our view, making it more likely that the results generalize well beyond our specific setting.

6.1. Data sample and empirical specification

In 2017, a new general hospital was opened in Ashdod—the sixth largest city in Israel, with a population of 220,000 in 2017—introducing to the city its first emergency department. Prior to its opening, the closest emergency departments serving the Ashdod population (which we refer to as “the treated population”) were located in general hospitals in nearby cities, each located 32 km (20 miles) away. Therefore, the opening of a new general hospital resulted in a sharp change in local availability of emergency services.

As a comparison group, we pick the five Israeli cities most similar in their total population size to the treated city.¹⁸ Throughout the study period, each of these cities had been served by at least one local hospital and had not experienced hospital entry or exit. We sample the entire Clalit patient population aged 25 and older in these cities.

[Table 4](#) shows summary statistics for this sample. Residents of both the treated and comparison cities have similar mortality rates, ages, and gender composition. Compared to Ashdod, other cities have higher average socioeconomic statuses and lower morbidity burdens (measured by the Johns Hopkins Adjusted Clinical Group commercial risk classifier).¹⁹ Nonetheless, both the treated and comparison populations have similar average total annual healthcare spending (NIS 9200 versus NIS 8950; the breakdown by

¹⁸ We chose the Israeli cities ranked third to eighth in Israel by 2017 total population size (ranging between 200,000 and 280,000 residents). These cities are: Haifa, Rishon LeZion, Petah Tikva, Netanya, and Be'er Sheva. We did not include in the comparison group Jerusalem and Tel Aviv (ranked first and second), which are substantially bigger, and the cities ranked ninth through eleventh (with populations of between 150,000 and 200,000 residents), which are all part of the Tel Aviv metropolitan area.

¹⁹ The ACG system is used by both commercial insurers and non-commercial healthcare organizations worldwide (as well as by Clalit) to describe or predict a population's past or future healthcare utilization and costs. For more information, see the Johns Hopkins ACG System Version 11.0 Technical Reference Guide (2014).

Table 4
Summary statistics: treated and comparison cities.

	Treated city (1)	Comparison cities (2)
<i>A. Demographics and health status</i>		
Age (mean)	54.8	54.8
Female (percent)	53.7	53.9
High SES (percent)	25.8	43.4
Number of chronic conditions (mean)	4.48	4.05
ACG high resource (percent)	19.4	16.0
Annual mortality rate (percent)	1.62	1.60
<i>B. Annual emergency department use</i>		
Percent with ED visit, by visit outcome		
Any outcome	21.6	25.4
Discharged home (outpatient ED)	15.6	19.6
Admitted (inpatient ED)	9.7	9.9
Died	0.058	0.070
Fraction of ED visits resulting in admission	40.5	34.0
Percent with UCC visit	13.6	5.9
<i>C. Annual spending, by category (mean, NIS)</i>		
Total (all categories)	9193	8956
Emergency department	171	212
Prescription drugs	2165	2185
Inpatient – Unplanned	2249	2219
Inpatient – planned	1729	1421
Outpatient services	1149	1284
Primary care visits	319	278
Specialist visits	366	357
Imaging services	420	438
Laboratory services	215	226
Other	387	328
Number of observations (patient-year)	160,116	1,340,494
Number of patients	41,013	343,332

Descriptive statistics for the sample of patients who resided in treated (experiencing hospital entry) and comparison cities (with no hospital entry) during 2015–2019. See Section 6 for sample definitions. Panel A shows demographic and health status information. High SES is the top tercile of zipcode-based socioeconomic status scores. ACG High Resource is a dummy for high predicted healthcare utilization. Panel B shows ED use—having any visits during the year and the number of visits conditional on any visits, by discharge status. Panel C shows healthcare utilization by service category. For comparability with Table 1, the numbers shown are annual, aggregated at the patient-year level. To better reveal time trends, our analysis below aggregates the same data at the patient-month level.


category is also similar). In line with Ashdod's relatively limited access to ED services, throughout the beginning of the sample period, its residents visited the ED less and visited urgent care centers more than residents of comparison cities.

Using individual-level claims data on the sample of treated and comparison cities from 2015–2018 aggregated at a quarterly or monthly frequency, we estimate a difference-in-differences specification:

$$Y_{it} = \alpha_i + \beta X_{it} + \gamma T_i + \delta_t + \eta_t T_i + \varepsilon_{it}, \tag{6}$$

where i and t index individuals and the number of periods since hospital opening; Y_{it} is one of the same outcomes as in Eq. (1), including ED use and cost by category; α_i and X_{it} are individual fixed effects and controls for time-varying individual characteristics (including ACG resource utilization band); and T_i is a “treatment” dummy for whether the individual lives in an area where a hospital opened.

(A) All ED Visits



(B) ED Visits by Discharge Status

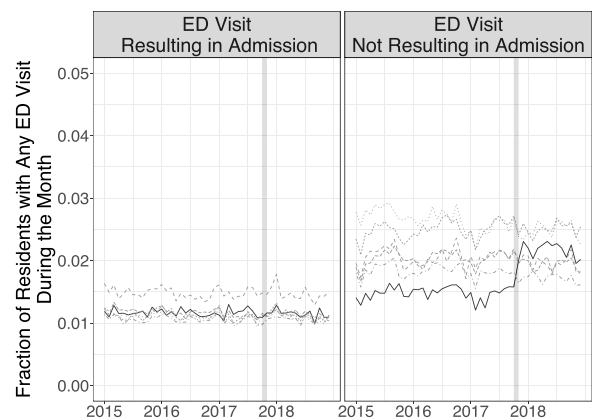


Fig. 7. Monthly ED utilization before and after hospital opening. *Notes:* For the treated city (solid black) and five comparison cities (gray), the figure shows the fraction of residents with one or more ED visits during each month. The gray vertical shade highlights November 2017, the month in which the index city opened its first ED. Before that time, it had been served only by the emergency departments in adjacent cities. Comparison cities were served by at least one local emergency department throughout the study period. The sample consists of 41,031 treated and 343,332 untreated patients ($N = 17,870,139$ patient-months).

The parameter of interest is η_t , the impact of a positive shock to ED access on ED use by nearby residents, relative to residents of other cities that did not experience such a shock. To the extent that the increased availability of an ED induces greater ED use, we expect an increase in ED use in the treated city post-treatment. In addition to this specification, which we use to produce event-study plots, we also estimate another variant substituting the monthly dummies with one dummy for all periods t after the hospital opening for reporting the overall effects in a more succinct tabulated form. The identification assumption is that trends in the treated and other cities would have been common if not for the shock we observe. Given that planning and constructing a hospital is a multi-year process, it seems plausible that the exact opening day is determined by factors such as construction and licensing and not in exact anticipation of local healthcare utilization (even if it

Table 5
Change in ED use and in spending related to ED opening.

	Estimated impact of hospital entry	S.E.	Pre-treatment mean dep. var (for treated)	Estimated impact of entry as percent of pre-treatment mean dep. var
	(1)	(2)	(3)	(4)
<i>A. Monthly emergency department use</i>				
Any ED visit	0.0070	0.0003	0.0255	27.5
ED visit not resulting in admission	0.0075	0.0003	0.0150	49.8
ED visit resulting in admission	0.0002	0.0002	0.0117	1.3
Any UCC visit	-0.0015	0.0002	0.0141	-10.4
<i>B. Monthly spending (by category)</i>				
Emergency Department	8.6	0.3	11.9	71.9
Prescription drugs	3.7	5.4	175.5	2.1
Imaging services	2.5	0.5	33.1	7.7
Laboratory services	-0.5	0.3	17.7	-2.6
Other	-0.8	1.0	31.5	-2.5
Outpatient services	-4.3	3.8	94.6	-4.6
Inpatient – planned	35.7	6.1	134.6	26.5
Primary care visits	-0.5	0.1	27.1	-1.7
Specialist visits	0.9	0.2	29.9	3.0
Inpatient – unplanned	21.4	7.1	179.3	11.9
Total (all categories)	66.7	13.5	737.2	9.0

Table shows estimates of the change in ED use due to ED opening. Column 1 shows estimates of Eq. (6) for the impact of the hospital opening on ED use and healthcare spending of (“treated”) nearby residents relative to residents of comparison cities that did not have any new hospital open during the study period. Column 2 shows standard errors for these estimates, clustered by patient. Column 3 shows the average pre-treatment outcome for residents of the treated city. Column 4 shows the magnitude of the impact (from column 1) as a percent of the pre-move baseline (from column 3). The sample consists of 41,031 treated and 343,332 untreated patients ($N = 17,870,139$ patient-months).

is correlated with long-run healthcare needs of the local population).²⁰

6.2. Results

As may be expected, the hospital opening led to a sharp increase in the treated city’s ED use, moving it from the bottom to the middle of the range of ED use across all comparison cities. This change is evident in Fig. 7, which shows the average monthly probability of ED use in each of the comparison cities in our sample (the treated city in black; comparison cities in gray).

Panel A of Table 5 summarizes the estimated impact of hospital entry on ED use, from estimates of Eq. (6) at a quarterly frequency; disaggregated estimates are shown in Appendix Figure A8. The improved access brought by the opening of a local ED was associated with a sharp increase in the probability of ED visits by nearby residents. Post-treatment, treated residents were 27.5% more likely to visit the ED in a given month (a 0.7 percentage-point increase in the probability of a treated patient’s visiting the ED each month, over the pre-period baseline of 2.5%).

Remarkably, consistent with our earlier findings from patient moves described in Section 4, the increased ED use associated with hospital entry is entirely concentrated in outpatient ED cases, the probability of which increased by almost 50% following the event. The use of urgent care centers—an ED substitute for less severe cases—decreased by more than 10%. In contrast, there was virtually no change in the share of patients with ED visits resulting in an

admission (a 0.01 percentage point increase over the pre-treatment baseline of 1.17%, with a standard error of 0.02 percentage points).

Panel B of Table 5 shows estimates of Eq. (6) with spending on each category and in total as the outcome (Appendix Figure A9 shows the same data as the table in a visual form; Appendix Figure A10 shows disaggregated estimates.) Average spending on ED services sharply increased among the treated population immediately after the opening (by 72%) and remained high throughout our sample period, which includes four quarters after the opening. Other spending categories that saw large increases were planned and unplanned inpatient services, which increased by 27% and 12%, respectively. Such increases are consistent with increased availability of inpatient services as a result of the hospital’s opening. Overall, the hospital’s opening was associated with a 9% increase in total average spending.

It is important to highlight that, naturally, a hospital opening reflects a different type of shock to access than a patient move, which we analyzed in the previous section. While a move is associated with a change in the composition (and distance to) multiple health services available to movers, a hospital opening is primarily associated with a sharp change in access to emergency and inpatient services to nearby residents. Therefore, there is no reason the two changes should have the same association with spending. Indeed, our findings suggest that they differ in this regard.

All results are robust to the definition of the comparison group, as shown in Appendix Figure A11 and Appendix Figure A12, which show estimates of Eq. (6) using five different samples obtained by separately pairing the treated city with each one of the comparison cities.

²⁰ The planned hospital construction was announced in 2002, and the bidding for its construction closed in 2011, six years before its inauguration.

7. Conclusion

We document large geographic variation in ED use in Israel in 2011–2017, similar to that observed in other countries. Using longitudinal data covering a large population, we study the role of supply- and demand-side factors in explaining this variation using two approaches that exploit two distinct supply-side shocks to ED access: patients who moved to an area with a different ED use rate and residents who had a nearby hospital opening. Overall, we find that about half of this variation can be attributed to supply-side factors and the rest to demand-side factors. Namely, we estimate that the probability of visiting the ED changed sharply after a move and following a hospital (and ED) opening.

In contrast, the variation in unscheduled admissions through the ED is entirely explained by person-specific demand factors. In particular, the (unconditional) probability of a patient being urgently admitted through the ED in a given year did not change upon a move or following hospital entry. The fact that the marginal cases—those most affected by supply-side factors—are of low severity is unsurprising because there are many reasons for less severe cases to be the most responsive to variation in ED access. What is, perhaps, more surprising is that even though a patient's probability of utilizing the ED is affected by supply-side factors, the admission threshold is not. That is, if unplanned admissions are driven by demand-side factors alone, as we find, it does not appear that ED physicians in low versus high ED-access locations have different criteria associated with admitting patients. This finding militates against a host of explanations that attribute the observed variation in ED admissions to place- or hospital-specific factors. For example, unless the following forces exactly offset each other in our sample, our main finding is inconsistent with different hospitals having different admission thresholds due to varying capacity, quality of on-site diagnostic, or financial incentives.

Analyzing potential channels, we show that three factors: the average distance of patients to the nearest hospital, the propensity of local physicians to refer (otherwise similar) cases to the ED, and the availability of urgent-care centers all contribute to the spatial, supply-side variation in ED use. While these channels are not exhaustive, they do highlight the fact that substitutes and complements for ED services, such as community-based urgent care and office-based primary care, determine in part the variation in ED use.

Taken together, our results highlight the dual purpose of Emergency Departments (in Israel, but presumably more broadly): a gateway to hospital admissions for more severe cases and an outpatient facility focused on treating less severe cases. The evidence we presented shows, repeatedly, that ED access has essentially no impact on healthcare utilization for patients who are sick enough; their demand is inelastic and they will arrive at the ED and be admitted regardless of how far it is or how likely their primary care physician is to refer patients to the ED. In contrast, less severe cases respond to supply conditions in predictable ways, and may substitute to alternate services

(UCC, community care, or no care at all), when ED access is more limited.

While not entirely surprising, these results present clear insights for policy. First, consider the policy concern that easier access to ED would lead to an inefficient increase in hospital admissions, and subsequent medical spending. Our results suggest that this is not the case. They suggest that variation in ED use, while important in itself, may be concentrated on the lower end of the severity spectrum, and as such does not necessarily affect variation in hospital admission. Therefore, efforts to influence admissions from the ED should be considered separately from efforts to reduce ED use for low-severity cases.

Second, consider public concerns about crowded EDs and long waiting times. Our results accentuate that low-severity ED visits have high elasticity of substitution with other services, and with the spatial distribution of patients, making it likely that a relatively straightforward way to reduce crowding at EDs, and allow them to focus on the more severe cases, is to increase the availability of substitute services, such as UCCs, which would capture demand for less severe cases.

Finally, our results may inform discussion about the (infrequent) decision to open or close new hospitals. They suggest that severe enough cases are inelastic, and thus would likely find their way to the ED regardless of how dense the hospital network is, suggesting that a new hospital or an expansion of an existing hospital would have a similar effect. In contrast, denser hospital networks may induce demand for less severe cases, thereby increasing healthcare utilization and spending, without noticeable effects on measurable health outcomes (such as admissions or mortality). Naturally, this conclusion holds within the density levels observed in Israel and in other developed countries. How sparse a hospital network could be before these results change—a question that may be critical for more developing countries or more rural areas—is an important question for further work.

Conflict of interest

The authors declare no conflict of interest.

Acknowledgement

We acknowledge financial support from the Koret Foundation.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:[10.1016/j.jhealeco.2021.102453](https://doi.org/10.1016/j.jhealeco.2021.102453).

References

- Abualenain, J., Frohna, W.J., Shesser, R., Ding, R., Smith, M., Pines, J.M., 2013. Emergency department physician-level and hospital-level variation in admission rates. *Ann. Emerg. Med.* 61 (6), 638–643.
- Agha, L., Frandsen, B., Rebitzer, J.B., 2019. Fragmented division of labor and healthcare costs: evidence from moves across regions. *J. Public Econ.* 169, 144–159.

- Alexander, D., Currie, J., Schnell, M., 2019. Check up before you check out: retail clinics and emergency room use. *J. Public Econ.* 178, 1–18.
- Allen, L., Cummings, J.R., Hockenberry, J., 2019. Urgent care centers and the demand for non-emergent emergency department visits. NBER Working Paper No. 25428.
- Baker, L.C., Kate Bundorf, M., Kessler, D.P., 2016. The effect of hospital/physician integration on hospital choice. *J. Health Econ.* 50, 1–8.
- Barnett, M.L., Song, Z., Landon, B.E., 2012. Trends in physician referrals in the United States, 1999–2009. *Arch. Intern. Med.* 172 (2), 163–170.
- Beckert, W., Christensen, M., Collyer, K., 2012. Choice of NHS-funded hospital services in England. *Econ. J.* 122 (560), 400–417.
- Cutler, D., Skinner, J., Stern, A.D., Wennberg, D., 2019. Physician beliefs and patient preferences: a new look at regional variation in health care spending. *Am. Econ. J.: Econ. Policy* 11 (1), 192–221.
- Dawson, H., Weerasooriya, J., Webster, G., 2008. Hospital admissions via the emergency department: implications for planning and patient flow. *Health Care Q.* 11, 2–20.
- Finkelstein, A., Gentzkow, M., Williams, H.L., 2019. Place-based drivers of mortality: evidence from migration. NBER Working Paper No. 25975.
- Finkelstein, A., Gentzkow, M., Williams, H., 2016a. Sources of geographic variation in health care: evidence from patient migration. *Q. J. Econ.* 131 (4), 1681–1726.
- Finkelstein, A.N., Taubman, S.L., Allen, H.L., Wright, B.J., Baicker, K., 2016b. Effect of Medicaid coverage on ED use—further evidence from Oregon's experiment. *N. Engl. J. Med.* 375 (16), 1505–1507.
- Finkelstein, A., Gentzkow, M., Williams, H., 2018. What drives prescription opioid abuse? Evidence from migration. SIEPR Working Paper 18-028.
- Gaynor, M., Vogt, W.B., 2003. Competition Among Hospitals. National Bureau of Economic Research.
- Kessler, D.P., McClellan, M.B., 2000. Is hospital competition socially wasteful? *Q. J. Econ.* 115 (2), 577–615.
- Llovera, I., Loscalzo, K., Gao, J., Li, T., Brave, M., Becker, L., Barata, I., 2019. Increased access to urgent care centers decreases low acuity diagnoses in a nearby hospital emergency department. *Am. J. Emerg. Med.* 37 (3), 486–488.
- Morganti, K.G., Bauhoff, S., Blanchard, J.C., Abir, M., Iyer, N., Smith, A., Vesely, J.V., Okeke, E.N., Kellermann, A.L., 2013. The evolving role of emergency departments in the United States. *RAND Health Q.* 3 (2).
- Myong, C., Hull, P., Price, M., Hsu, J., Newhouse, J.P., Fung, V., 2020. The impact of funding for federally qualified health centers on utilization and emergency department visits in Massachusetts. *PLOS ONE* 15 (12), e0243279.
- Romley, J.A., Goldman, D.P., 2011. How costly is hospital quality? A revealed-preference approach. *J. Ind. Econ.* 59 (4), 578–608.
- Sabbatini, A.K., Nallamothu, B.K., Kocher, K.E., 2014. Reducing variation in hospital admissions from the emergency department for low-mortality conditions may produce savings. *Health Aff.* 33 (9), 1655–1663.
- Shoff, C., Caines, K., Pines, J.M., 2018. Geographic variation in predictors of emergency department admission rates in US Medicare fee-for-service beneficiaries. *Am. J. Emerg. Med.*
- Taubman, S.L., Allen, H.L., Wright, B.J., Baicker, K., Finkelstein, A.N., 2014. Medicaid increases emergency-department use: evidence from Oregon's Health Insurance Experiment. *Science* 343 (6168), 263–268.
- Van Parys, J., 2016. Variation in physician practice styles within and across emergency departments. *PLOS ONE* 11 (8).