

Consolidation of primary care physicians and its impact on healthcare utilization

Jonathan Zhang¹  | Yiwei Chen² | Liran Einav³  | Jonathan Levin⁴ | Jay Bhattacharya⁵

¹Princeton School of Public and International Affairs, Princeton University, Princeton, New Jersey, USA

²Facebook, Menlo Park, California, USA

³Department of Economics, Stanford University, Stanford, California, USA

⁴Graduate School of Business, Stanford University, Stanford, California, USA

⁵School of Medicine, Stanford University Encina Commons, Stanford, California, USA

Correspondence

Jonathan Zhang, Princeton School of Public and International Affairs, Princeton University, 176 Julis Romo Rabinowitz Building, Princeton, NJ 08544, USA.

Email: jxzhang@princeton.edu

Funding information

NIA, Grant/Award Number: #R01 AG032449

Abstract

We use administrative data from Medicare to document the massive consolidation of primary care physicians over the last decade and its impact on patient healthcare utilization. We first document that primary care organizations have consolidated all over the United States between 2008 and 2014. We then show that regions that experienced greater consolidation are associated with greater decline in overall healthcare spending. Finally, in our primary exercise, we exploit transitions of patients across organizations that are driven by changes in the organizational affiliations of their primary care physicians to study the impact of organizational size on overall spending. Our preferred specification suggests that patients switching from small to large physician organizations reduce their overall healthcare spending by 16%, and that this reduction is primarily driven by a 13% reduction in primary care visits and 0.09 (21%) fewer inpatient admissions per year.

KEYWORDS

consolidation, health, healthcare utilization, primary care

1 | INTRODUCTION

Over the last 2 decades, there is a fairly massive consolidation of providers in the US healthcare industry, and this trend has been particularly pronounced in the context of primary care physicians (Cutler & Scott Morton, 2013; Kirchoff, 2013; Muhlestein & Smith, 2016). As primary care serves as the first point of contact of many medical services for patients, such consolidation could have great impacts on healthcare markets and healthcare utilization. On one hand, it may improve efficiency due to economics of scale and scope: as primary care physicians become more consolidated, they may be able to implement more procedures in outpatient settings and reduce wasteful and repetitive visits. On the other hand, it also raises concerns about increased market power: when the market becomes more consolidated, the competitiveness of the market may go down, and providers may be able to raise prices or even induce greater demand from patients.

In recent years, researchers have studied how these two opposing forces impact the price of care in various settings of consolidations in the healthcare industry. However, few have focused specifically on the consolidation of primary care organizations and how they would affect the utilization of medical services. The goal of this paper is to study the effect of consolidation on total patient utilization. Its importance is at least twofold. In the context of public health insurances such as Medicare and Medicaid where prices are regulated, the utilization effects would drive the

healthcare spending for enrollees under those programs instead of the price effects. Second, any analysis of the impact of consolidation should involve both price and utilization effects, so even if prices are affected, it is important to understand the impact on the side of patient utilization. Moreover, primary care is the most frequently utilized health setting, especially for preventive care, and often serves as a gatekeeper and first point-of-contact before specialty care.

In this paper, we study these questions with administrative claims-level data for Traditional Medicare enrollees over the years 2008–2014. This setting is attractive in this context because Medicare prices are regulated, so overall Medicare spending is primarily driven by healthcare utilization. We start by constructing an association between treating primary care physicians and the organizations they use for billing. Using this association, we document the consolidation in primary care across the United States. Then, to study the impact of consolidation on patient healthcare utilization, we exploit the transition of individual primary care physicians from small to large physician organizations and the empirical fact that patients tend to stay with their providers.

We find that consolidations of primary care organizations occurred all over the United States during our observation period. The share of primary care visits to organizations with over 30 providers grew from 32% to 46% between 2008 and 2014. At the regional levels, regions with either higher level or growth of concentration of large primary care organizations display lower regional healthcare utilization. However, at the patient level, we find that the naïve cross-sectional patterns may be misleading. Larger physician organizations tend to attract sicker patients, but once we utilize a within-patient analysis—exploiting the patient's primary care physician changing organizations as an instrumental variable—we find that transition from small to large (where the distinction is 30 or more providers per organization) primary care organizations is associated with a 16% reduction in overall patient spending. This reduction is primarily driven by a reduction in the number of inpatient admissions (0.089 fewer admissions, a 21% reduction, and 5.3 percentage point less likely to have any inpatient encounter, a 25% reduction) as well as a 12.8% reduction in the number of primary care visits. Our results suggest that consolidation of primary care physicians may mitigate growth in healthcare spending in the United States, operating primarily through fewer encounters.

Our paper contributes to the large and growing literature on consolidation in the healthcare industry. Dranove and Lindrooth (2003) found that hospital mergers that share common ownership but do businesses under separate financial licenses do not save money. In contrast, Harrison (2011) found that hospital mergers do save money although the cost savings decline over time. In the physician organization space, McWilliams et al. (2016) examined the early performance of accountable care organization and found a modest cost saving effect.

There is also a related large literature that is focused on vertical integration between hospitals and physician practice groups. Several studies of this topic found that such integration raise prices (Baker et al., 2014; Capps et al., 2018; Conti et al., 2016; Neprash et al., 2015; Robinson & Miller, 2014). Koch et al. (2017) also found that vertical integration also raises utilization. Capps et al. (2018) study vertical integration of providers in private practices, and find that patients' primary care spending increases by 4.9% following vertical integration. Much of this increase is likely due to changes in prices because they find payment rule exploitation in hospital practices. In contrast, we present suggestive evidence that the price channel (at least in the context of Traditional Medicare patients) does not play a major role. This is consistent with the fact that our main empirical strategy is primarily driven by consolidation between physician groups rather than by vertical integration in which hospitals acquire physicians, thus limiting the scope for price changes.

Capps et al. (2018) is also closely related to our work in terms of its empirical strategy. They exploit the fact that enrollee's choice of a primary care physician (PCP) tend to persist and use vertical integration status of each patient's first main PCP as an instrumental variable for the share of visits to a vertically integrated PCP. Our instrumental variable strategy is similar in spirit in the sense that we also leverage persistence in patient-provider relationships, except that we build our instrument based on changes in organization size because we are interested in studying patient spending after consolidation.

Perhaps the most similar work to ours is the concurrent ongoing work by Baker et al. (2019). They analyze Medicare beneficiaries who moved between primary care physicians of single and multi-specialty practices due to geographical market concentration of organizations. They reach similar conclusions to ours, and their preliminary results suggest that consolidation of primary care physicians reduces medical spending and hospitalizations for patients.

The rest of the paper is organized as follows. Section 2 describes the data and the construction of the key variables. Section 3 documents the consolidation of primary care organizations over our observation period. Section 4 introduces our empirical strategies to investigate the impact of this consolidation on healthcare utilization. The results are presented in Section 5, and in Section 6 we report various robustness checks. The last section concludes.

2 | DATA, SAMPLE, AND VARIABLE CONSTRUCTION

We use administrative data from the Centers of Medicare and Medicaid Services (CMS). We primarily rely on the claim-level information associated with the “carrier” of a 20% random sample of fee-for-service, Traditional Medicare enrollees over the years 2008, 2010, 2012, and 2014. These cover all non-institutional Medicare Part B claims, which capture the vast majority of outpatient services provided to Traditional Medicare enrollees. The Appendix 1 provides many more details about the underlying data and the sample construction, and in this section, we highlight the key aspects.

In a given year, we restrict the sample to only those enrollees who are older than 65, who are enrolled in Traditional Medicare for all 12 months of the year, and to enrollees for whom we can calculate risk scores (using prior year data) as of the beginning of the year, which excludes enrollees who were enrolled in Medicare Advantage for any month of the previous year. For those enrollees, the data contain detailed information about the claim, and in particular the total cost associated with the claims and (when relevant) unique identifiers for the billing organization (in the form of Tax Identification Numbers, or TIN), the physician (in the form of her National Provider Identifier, or NPI), and her specialty. Because large parts of our analysis are focused on PCPs, we define PCPs to be those physicians associated with one of the following specialties: general medicine, internal medicine, family medicine, and geriatric medicine.

We organize the analysis separately by calendar year, so all the variables described below are constructed separately for each year of data. We define an organization to be the billing organization, and we define the organization size to be the number of unique providers that use this organization at least once for billing purposes during the year. Using this measure of size, we categorize organizations as large or small based on a threshold size, with a size of 30 or more classified as a “large” organization.¹ This is our key explanatory variable. In some of our analyses, we associate an organization size with the individual physician. If a physician is associated in our data with more than a single organization,² we take a simple, unweighted average over the size of the organizations with which she is associated, and the big organization indicator becomes continuous, indicating the fraction of “large” organizations out of all organizations the physician is associated with.

The patients in our baseline sample include all patients who have at least a single non-institutional Part B claim (outpatient encounter) in a given year that is associated with a primary care physician. We then assign each patient a propensity to be seen by a big organization by computing the fraction of primary care visits with a PCP that is affiliated with a big organization, across all of the patient's PCP visits; for visits to PCPs with multiple organizational affiliations, the average of the continuous analogue is taken. Recall that our unit of observation is at the year level, so patient-level variables and physician-level variables (including big organization) are for that particular year.

We also use the CMS claims files to calculate patients' overall healthcare utilization, including total expenditure and its components by setting (e.g., primary care visit, inpatient hospitalization, etc.). Because patients' health risks vary in predictable ways, for each patient we also calculate a patient's risk score,³ which captures the normalized expected cost—with a risk score of one reflecting an average risk score for all Medicare enrollees in a given year. A patient's risk-adjusted spending then becomes her overall expenditure, divided by her risk score. For the analyses that are done at the Hospital Referral Region (HRR) level, we assign each physician, organization, and patient to their corresponding HRR, and aggregate the above variables to the HRR level. The Appendix 1 provides more details.

Table 1 reports some high-level summary statistics on our sample of patients. The overall sample size is reasonably stable across years, and it includes approximately 3.7 million patients per year. The average patient in the sample is associated with approximately 11,000 dollars of annual healthcare expenditure, split roughly evenly (on average) between inpatient and outpatient services. This even split masks heterogeneity at the patient level, as more than 75% of the patients have no inpatient spending in a given year. Because we require patients to have at least one primary care visit during the year, the average patient in our data is sicker than the average Medicare enrollee, with average risk scores of 1.2–1.25. Similarly, on average, patients have approximately 6 primary care visits and 8 specialist visits in a given year, and in the process, they are associated with six unique organizations.

3 | DESCRIPTIVE PATTERNS

3.1 | Aggregate trends

We begin by documenting the fairly massive consolidation in the way primary care physicians organize themselves over the observation period of 2008–2014. Using our baseline data, the top panel of Figure 1 plots the fraction of primary care

TABLE 1 Summary statistics of patients sample

	2008		2010		2012		2014	
	(N = 3,705,938; 74.8%)		(N = 3,670,235; 73.5%)		(N = 3,707,578; 72.8%)		(N = 3,680,873; 71.9%)	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Total annual spending (2014 \$US)	10,525	3371	11,209	3605	10,882	3578	10,996	3556
Inpatient spending	4955	0	5105	0	4712	0	4711	0
Outpatient spending	5570	3062	6104	3305	6170	3306	6284	3306
Risk-adjusted spending (2014 \$US)	9226	3486	10,097	3837	10,171	3945	10,834	4149
Risk score as of January 1	1.25	0.95	1.23	0.92	1.22	0.89	1.18	0.84
Inpatient visits	0.46	0	0.44	0	0.41	0	0.40	0
Any inpatient	0.23	0	0.22	0	0.21	0	0.20	0
ER visits	0.32	0	0.33	0	0.34	0	0.34	0
Outpatient visits	19.05	15	19.51	15	19.79	15	19.95	15
PCP visits	5.88	4	5.88	4	5.84	4	5.54	4
Specialist visits	7.74	5	8.15	5	8.36	5	8.54	5
Unique organizations	5.68	5	5.91	5	5.99	5	6.07	5
Unique PCPs	1.79	1	1.83	1	1.89	1	1.89	1
Unique specialists	3.27	2	3.49	3	3.65	3	3.73	3

Notes: Table shows summary statistics, year by year, for our baseline patient sample as described in Appendix 1. All dollars are adjusted to 2014 level using CPI-U. Each observation is a patient. The second row displays the number of observation in the patient sample (N) and the share of the patient sample who meet our eligibility criteria in Appendix A5.2 and have at least one non-lab primary care visit. That is, about 25%–28% of patients meet all the eligibility criteria but do not have any non-lab primary care visit in that year.

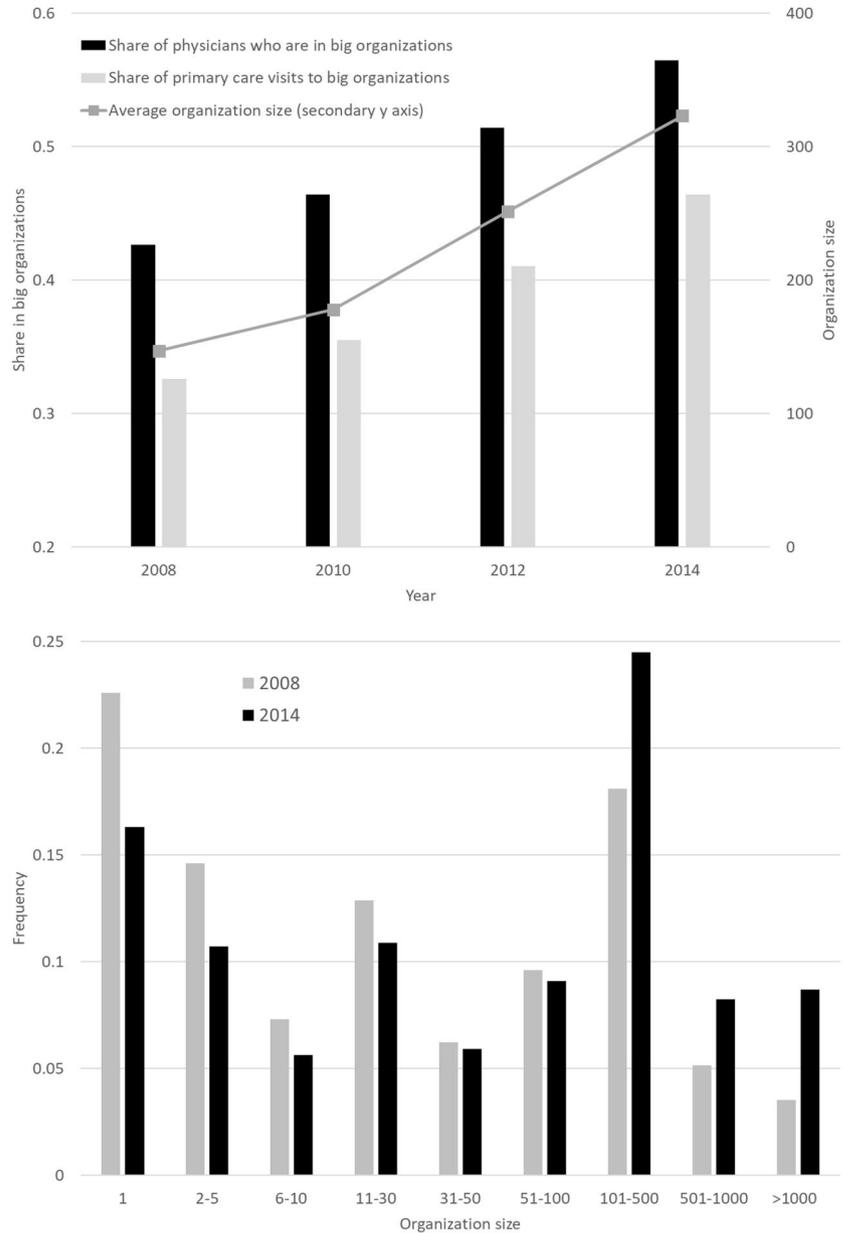
Abbreviation: PCPs, primary care physicians.

physicians who are in big organization⁴ in black bars and the share of primary care visits to big organizations across patients in grey bars, and the grey line shows the distribution of organization sizes across physicians on the secondary y-axis. The trend is quite remarkable, with the average organization size more than doubles (147–324) between 2008 and 2014, and the share of physicians associated with big organizations growing from 43% to 56%. Naturally, a similar trend occurs from the patient perspective, with the share of patient primary care visits to big organizations growing from 32% in 2008 to 46% in 2014.

The bottom panel of Figure 1 shows the distribution of organization sizes across physicians.⁵ It illustrates that this consolidation is not driven by any particular range in organizational sizes, and mostly reflects a fairly widespread shift in the size distribution. Moreover, the pattern cannot be explained by an increased supply of PCPs as the share of medical students entering family medicine residency programs only increased marginally from 8.95% in 2009 to 9.13% in 2014 (Phillips et al., 2019).

Figure 2 presents this trend across HRRs. In each panel, we plot the fraction of primary care physicians belonging to big organizations in each HRR h in a year t ($f_{h,t}^{big}$ as defined in the Appendix 1). The colors are coded according to the four quantiles of $f_{h,2008}^{big}$ in 2008. The top map corresponds to 2008, the first year in our data, and the bottom map to 2014, the last year in our data. Comparing the two maps, the panels illustrate that the consolidation is not concentrated in particular regions of the United States, but rather across virtually the entire country. Furthermore, Figure 2 also illustrates significant heterogeneity in the organizational structure of primary care physicians as of 2008, the first year in our data. In 2008, the majority of physicians in the Midwest states such as Minnesota and Wisconsin were already associated with big organizations, while in other parts of the country most physicians were still working alone or were part of a small physician group. By 2014, almost all HRRs show a larger share of physicians associated with big organization, and leaving almost no HRRs with less than 25% of its physicians working for small organizations.

FIGURE 1 Consolidation of primary care physician, 2008–2014. Top panel displays the average share of primary care physicians who are in big organizations ($\%big_{jt}$ in the Appendix 1) in black bars; the share of Medicare beneficiaries who visit big primary care organizations (w_{jt}^{big} in the Appendix) in gray bars; and the average organization size (Org_{jt} in the Appendix 1) of a primary care physician in the solid line (secondary y-axis). The bottom panel plots the distribution of primary care physician organization size Org_{jt} in 2008 (gray color) and 2014 (black color)



3.2 | Consolidation and trends in healthcare utilization within HRRs

The geographical variation shown in Figure 2 correlates quite strongly with healthcare utilization, but naturally cannot be interpreted as a causal effect. In order to address some of the most obvious endogeneity concerns, we continue by reporting results from a difference-in-differences analysis at the HRR level.

Specifically, we construct observation at the HRR-by-year (h, t) level and run the following regression:

$$y_{ht} = \delta_h + \tau_t + \beta f_{h,t}^{big} + \epsilon_{ht}, \tag{1}$$

in which y_{ht} is an outcome of interest (measured at the HRR-year level), δ_h and τ_t are HRR and year-fixed effects (respectively), and $f_{h,t}^{big}$ is the fraction of primary care physicians in HRR h and year t , who are associated with big organizations. Our main outcomes are spending and utilization counts. The coefficient of interest is β , which captures the relationship between changes in the extent of consolidation of primary care physicians at the HRR level and changes in HRR-specific trends in population health.

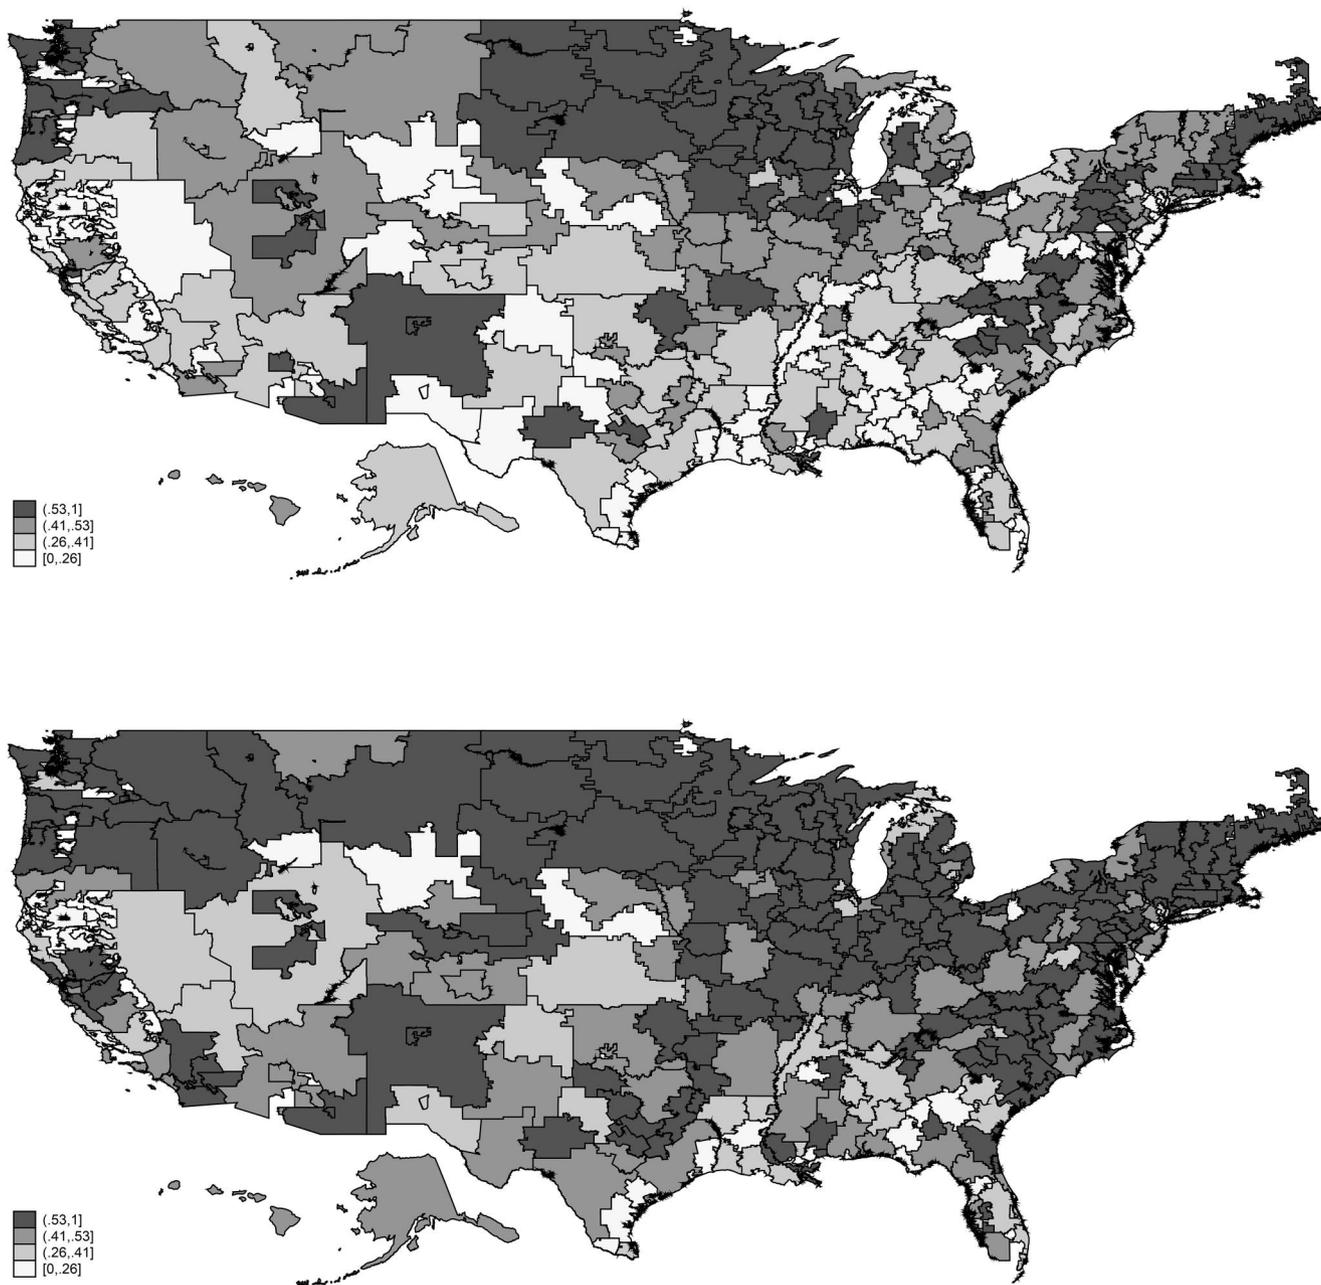


FIGURE 2 Primary care organization size in 2008 and 2014. Figure shows the fraction of primary care physicians belonging to big organizations in each Health Referral Regions ($f_{hrr,t}^{big}$ in the Appendix 1) in 2008 (top panel) and 2014 (bottom panel). The color-coding uses the quantiles of $f_{hrr,t}^{big}$ in 2008.

The regression results are reported in Table 2. Each row reports the estimate of β for a different dependent variable y_{ht} . Column 1 reports results from estimating Equation (1) without HRR-fixed effects (“pooled regressions”) and column 2 reports results from the full difference-in-differences specification. Both the pooled regression and difference-in-differences suggest that greater shares of large PCP organizations are associated with a reduction in healthcare spending and utilization. The difference-in-differences specification suggest that having a PCP that is associated with a big organization is associated with a 6% reduction in healthcare spending (almost entirely due to outpatient spending) and 5% reduction in outpatient encounters. Consolidation is associated (mechanically) with fewer organizations (and with more primary care physicians), which may explain the reduction in healthcare spending.

It is instructive to compare the pooled regression and difference-in-differences coefficients. Including HRR fixed effects significantly attenuates the regression coefficients. This implies that regions where more organizations

TABLE 2 The effect of the HRR-level fraction of physicians associated with big organizations

Dep. Variable	Mean Dep. Var.	Pooled	Diff-In-Diff
Log (adjusted total spending)	9.20	-0.025 (0.028)	-0.057 (0.048)
Log (total spending)	9.26	-0.158*** (0.038)	-0.048 (0.049)
Log (inpatient spending)	8.45	-0.095* (0.053)	-0.007 (0.071)
Log (outpatient spending)	8.66	-0.203*** (0.035)	-0.082* (0.045)
Number of inpatient admissions	0.43	-0.032 (0.020)	0.023 (0.027)
Any inpatient admission	0.22	-0.012 (0.008)	0.012 (0.009)
Log (outpatient visits)	3.00	-0.128*** (0.028)	-0.048** (0.023)
Log (primary care visits)	1.89	-0.182*** (0.035)	0.055 (0.034)
Log (specialist visits)	2.15	-0.231*** (0.043)	-0.078*** (0.028)
Distinct organizations	5.61	-2.901*** (0.238)	-0.460*** (0.164)
Distinct primary care physicians	1.84	0.330*** (0.042)	0.218*** (0.053)
Distinct specialists	3.36	-0.391*** (0.133)	-0.068 (0.092)
Year-fixed effects	-	Yes	Yes
HRR-fixed effects	-	No	Yes

Notes: Table is based on the baseline sample of patients, aggregated to the HRR-year level. The first column represents the dependent variables, which are the HRR-year level healthcare outcomes (all defined precisely in the Appendix 1). The independent variable is the average fraction of medical claims primary care physicians bill from big organizations in the HRR of that year ($f_{h,t}^{big}$). The second column reports the mean of the dependent variables. The third and fourth column report the estimated coefficient from the pooled (no HRR-fixed effects) and DD specifications, respectively. Standard errors are clustered at the HRR level and displayed in parentheses.

Abbreviation: HRR, Hospital Referral Region.

*, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

consolidated tended to have lower spending and utilization. For example, Minnesota and Wisconsin have relatively low spending levels, and are the most consolidated as shown in Figure 2.

We caution against interpreting the difference-in-differences results reported in Table 2 as causal. In the scenario where HRRs experienced greater consolidation based on changes or projections in population health and its associated needs, the identifying assumption of the HRR-level analysis is violated. Moreover, the HRR-level analysis is slightly indirect. It does not tell us whether switching a patient to a new physician organization would impact their utilization; this is ultimately the policy question of interest. In the next section we therefore turn to our primary empirical strategy, which is at the patient level.

4 | EMPIRICAL STRATEGY

In this section, we therefore attempt to address both of the above concerns by relying on a research design that more explicitly attempts to approximate an exogenous move by a patient from being treated by a primary care physician who is affiliated with a small organization to being treated by a physician affiliated with a big organization.

In particular, we use the following regression specification at a person-year (i, t) level:

$$y_{it} = \beta w_{i,t}^{big} + \gamma_{h(i,t)} + \chi_i + \epsilon_{it}, \quad (2)$$

where y_{it} is patient i 's healthcare outcome in year t , $w_{i,t}^{big}$ is the fraction of patient i 's primary care visits that are associated with big organizations, and $\gamma_{h(i,t)}$ is an HRR-year fixed effects, where $h(i,t)$ is the HRR associated with patient i during that year,⁶ and χ_i are patient-fixed effects.

In this specification, identification of the coefficient of interest (β) is driven by within-patient transitions between small and big organizations. While adding patient fixed effects already addresses many possible endogeneity concerns, especially those associated with potential sorting of sicker patients to physicians who are associated with larger organization, one may still be concerned that patient transitions are not random, and could be driven, for instance, by changes to patient health.

In our preferred specification, we address this concern by isolating within-patient changes that are driven by her physician *rather than the patient herself*. That is, the basic idea is that we rely on patient transitions that are primarily driven by her primary care physician transitioning between organizations rather than the patient herself self-selecting physicians (or organizations), taking advantage of the fact that many patients “follow” their primary care physicians even when they change their organizational affiliation. We implement this idea by estimating the above regression using an instrumental variable approach.

Specifically, we assign each patient to her two most common primary care physicians in a given year, and assign each primary care physician to her most common organization size. We then construct an instrument $1_{i,t}^{sb}$, which is an indicator variable that is equal to one if the organizational affiliation of at least one of the two primary care physician of patient i has changed from a small to a big organization between year $t - 2$ (the previous year we observe each patient and physician in our sample) and year t . We then use this variable $1_{i,t}^{sb}$ as an instrument for the independent variable $w_{i,t}^{big}$ in Equation (2). Patient-years that do not have one of their most common two primary care physicians change from a small to big organization receive a $w_{i,t}^{big}$ value of zero. Note that the provider needs to *change* organizations and cases where the organization naturally expands over the years to reach the threshold of 30 *do not* count. The Appendix 1 defines the precise way by which we construct these variables.

The identifying assumptions of this empirical specification (and our instrumental variable) is that organizational changes for PCPs are not directly driven by their individual patients, and the organization change status of any patient's PCP impacts utilization only through organizational size. This is a relatively strong assumption that could be violated in a variety of ways. For example, PCPs may switch organizations because of their patient pool (endogeneity concerns). We check for evidence of this by investigating trends in patient risk among switching PCPs. This assumption may also be violated if PCP switches impact spending and utilization through other channels other than organizational size. For instance, there may be mechanical changes in utilization simply from disruptions in care. We perform various robustness checks to address this concern in Section 6. Moreover, recent work by Sabety et al. (2020) find that Medicare beneficiaries who lose their PCP actually substitute to emergency and specialist visits in the short-run, and thereby increase spending.

One may wonder how often physician change organizational size, but fortunately, given the massive consolidation of primary care over our observation period which we documented earlier, there are many such instances. Appendix Table A1 summarizes the frequency of such transitions. Across years, the patients whose physicians moved from small to big organizations represent approximately 2% of total patients in the sample. It is therefore not surprising that the first stage in the IV regression is quite strong, with a positive coefficient of 0.33 and a standard error of 0.01 (see Appendix Table A2), reflecting our hypothesis that many patients “follow” their physicians, even when those transition organizations. That is, all else equal, of all the patients whose top physicians transitioned from a small to a big organization, about a third “fully transitioned” as well.⁷

5 | RESULTS

Table 3 reports the results from estimating Equation (2). We first report the OLS estimates in column 1, and the IV estimates next in column 2. The comparison is quite remarkable. The OLS estimates, despite including patient fixed effects, are positive for all healthcare outcomes. We interpret it as likely reflecting endogenous selection: as patients' health deteriorates, she is more likely to transition to (or get referred to) big organization, especially if big organizations have a positive effect on patient outcomes. For comparison, reduced form regressions outputs are displayed in Appendix Table A2.

In sharp contrast, the IV estimates are all of a negative sign. Our main estimate implies a large and statistically significant effect of 16% reduction in risk-adjusted (annual) total healthcare spending and 13% reduction in unadjusted (annual) healthcare spending. This is compared to a 19-20% increase in annual spending from the OLS estimates.

It is worth noting that even adjusting for patient risk, changes in patient spending can reflect factors beyond utilization, such as changes to facility reimbursements or payment schedules (whether through exploitation as in Capps

TABLE 3 The effect of big versus small organization on patient outcomes

Dep. Variable	Mean Dep. Var.	OLS estimates	IV estimates
Log (adjusted total spending)	8.35	0.201*** (0.017)	-0.163*** (0.019)
Log (total spending)	8.27	0.193*** (0.018)	-0.128*** (0.017)
Number of inpatient admissions	0.42	0.182*** (0.016)	-0.089*** (0.013)
Any inpatient admission	0.21	0.086*** (0.007)	-0.053*** (0.006)
Log (inpatient spending)	2.03	0.734*** (0.049)	-0.490*** (0.048)
Log (inpatient spending) cond. on any	9.54	0.064*** (0.014)	0.051 (0.046)
Log (outpatient spending)	8.05	0.099*** (0.010)	-0.074*** (0.011)
Log (outpatient visits)	2.75	0.044*** (0.006)	-0.030*** (0.007)
Log (primary care visits)	1.65	0.020** (0.008)	-0.128*** (0.010)
Log (specialist visits)	1.77	0.073*** (0.007)	-0.020** (0.008)
Distinct organizations	5.99	0.102** (0.051)	-0.452*** (0.047)
Distinct primary care physicians	1.87	0.518*** (0.032)	-0.230*** (0.029)
Distinct specialists	3.62	0.466*** (0.036)	-0.138*** (0.031)
OLS/IV?	-	OLS	IV
HRR-Year fixed effects	-	Yes	Yes
Patient-fixed effects	-	Yes	Yes

Notes: Table is based on the baseline sample. Each observation is at the patient-year level. The first column represents the dependent variables, which are the patient-year level healthcare outcomes (all defined precisely in the Appendix 1). The independent variable is the fraction of primary-care visits that are associated with big organizations during the year (w^{big}). The second column reports the mean of the dependent variables. The regression includes the HRR-by-year fixed effects and patient-fixed effects. The third and fourth columns are the OLS and IV estimates, as described in the main text, where the IV (for w^{big}) is an indicator variable that is equal to one if one of the two most frequently visited primary care physician moved from a small to a big organizations between 2 year $t - 2$ and year t (1_{it}^{sb}). Standard errors are clustered at the HRR level and displayed in parentheses.

Abbreviation: HRR, Hospital Referral Region.

*, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

et al., 2018 or other means). For example, if organizational consolidations also come with exploitation of payment rules, then our 16% reduction estimate will be an overestimate. However, in addition to reductions in overall spending, consolidations may also impact the number of visit encounters.

The other rows of Table 3 show that this reduction cannot simply be explained by changes in facility reimbursements or payment schedules, but rather reductions in spending are largely driven by a decline in the number of inpatient admissions and the number of primary care visits. The number of specialist visits is much less affected. The probability of having any inpatient admission is lower by 5 percentage points (relative to an average of 21%), presumably reflecting greater substitution towards outpatient care. The number of primary care visits decreases by 13% and the number of specialist visits decreases by only 2%, suggesting a possible reduction in unnecessary, duplicate care. The number of distinct PCPs decline by 0.23, which means that patients are more likely to receive continued care via a smaller set of providers (even though organization sizes are larger). The medical literature has found that generally primary care and continuity of care (the main mechanism in which our instrument operates) reduce hospitalizations (Kao et al., 2019; Shi, 2012), and our results are consistent with this view.

As mentioned earlier, our paper is closely related to Capps et al. (2018), who find that vertical integration of PCPs increase overall patient spending by 4.9% for commercially insured patients. While it may appear that this finding is at odds with ours, we note that Capps et al. (2018) cannot disentangle price effects from quantity effects and actually find substantial evidence for exploitation of payment rules. They find suggestive (nonsignificant) evidence of an increase in imaging, which can often be wasteful, of poor quality, and associated with redundant information (Oren et al., 2019) and with an increase in overall testing; similarly a positive coefficient on evaluations and management suggest a potential increase in utilization.

Importantly, vertical integration (in which hospitals acquire physician practices) and (horizontal) consolidation (in which physician group merge and/or expand) are not entirely the same. While some of the channels could be similar,

the “exploitation” of Medicare pricing rules is less relevant in the context of consolidation. Indeed, as we show in Appendix Figure A2, where we plot average outpatient and inpatient spending per visit for years before and after a PCP switches from a small to big organizations, the price effect does not seem to play an important role in our analysis. To isolate more directly vertical integration from “natural” organizational growth, we approximate mergers and acquisitions (M&As) in our data in Appendix Table A4.⁸ We show that the vast majority of PCP moves in our data are not driven by M&As, suggesting that vertical integration does not play an important role in our analysis. Indeed, our main results are qualitatively very similar once we remove the 16% of PCPs whose move from small to big organizations is likely driven by M&As. It is important to note that this approximation of M&As exercise would fail for very small acquisitions and scenarios where an original pre-acquisition organization continues to bill with the same TIN. Therefore, one limitation of this study is that we cannot distinguish between the utilization effects associated with growing and merging practices.

6 | ROBUSTNESS

In this section we perform several additional robustness checks, lending credibility to our main results and our identifying assumption. We begin by changing the definition of a “big” organization from a cutoff of 30 physicians per organization to 50 and 100. These cutoffs are selected following Muhlestein and Smith (2016). The results, presented in Appendix Table A3 and columns 2 and 3 of Table 4, are qualitatively similar to our baseline results in Table 3 (also displayed in column 1 of Table 4 for comparison), albeit smaller in magnitude. For example, crossing the cutoff of 30 PCPs leads to a 16% reduction in log adjusted total spending (our baseline estimate), whereas only 8.8% and 10.5% for cutoffs of 50 and 100 PCPs. The majority of the reduction in spending and visits are driven by consolidations of 30 or more providers, and the effects decline at larger organization sizes.

Recall that our main endogenous variable is a continuous variable which represents a patient's share of PCPs who belong to a big organization, which makes the interpretation difficult. We run the same specification as in Table 3, but after we restrict the sample to those patient-year observations who are *only* associated with big or small organizations, that is $w_{i,t}^{big} \in \{0, 1\}$. We thus do not use patient-year observations who (within the year) visit primary care physicians from both small and big organizations. In this case, identification is only driven by “complete” transitions, directly addressing this concern. Appendix Table A5 and column 4 of Table 4 reports the results, which are qualitatively similar to our baseline estimates.

Another potential concern is the our estimates may reflect a mechanical bias that is associated with physician transition (e.g., an adjustment period during which healthcare utilization is lower) rather than the specific effect of transitioning from a small to a big organization. It is worth noting that our instrument is based off physician consolidation and not physician exits. In recent work, Sabety et al (2020) find that provider exits are associated with higher spending due to an immediate shift to increased specialist and emergency care. Instead, we find the exact opposite because we leverage moves associated with consolidations. Nevertheless, to address this potential concern, we estimate a specification that restricts the sample only to patients whose main primary care physicians moved from small to big or from small to small organizations once and only once during the sample period (roughly 8% of the main sample). This exercise now compares patients who all have physicians who change organizations, and all change from a small organization, therefore, every observation in our sample experiences any potential disruption to care. The instrumental variable is the same as in the main specification. The results are reported in Appendix Table A6 and column 5 of Table 4. The OLS results are almost the same, while the IV results are still large and negative, and statistically significant, although the point estimates are about half as in the main specification. This robustness check is still based on moves from small organizations. We estimate a reduced form regression on a sample of patients whose providers all switch to a big organization (small-to-big or big-to-big). The results of that regression are comparable to the reduced form results on our main sample which includes non-movers (Appendix Table A2). This provides additional evidence that our results are driven off changes to organizational size rather than mechanical disruptions associated with practice ownership.

Finally, we note that our identification strategy works off variation in physician moves, and while any endogenous organization change decision on part of the physician may violate our research design, it is unlikely that the physician changes organization for any particular patient out of their entire client base. To illustrate this point, Appendix Figure A1 displays the result of a difference-in-differences regression of risk score of patients whose providers switch from small to big organizations in 2012, compared to those who never switch organizations in our

TABLE 4 Robustness and sensitivity for select main outcomes

	IV Estimates				
	Baseline results (1)	Threshold: 50 (2)	Threshold: 100 (3)	No partial affiliation (4)	Starting Org.: small (5)
Log (adjusted total spending)	-0.163*** (0.019)	-0.088*** (0.010)	-0.105*** (0.010)	-0.146*** (0.019)	-0.053** (0.023)
Log (total spending)	-0.128*** (0.017)	-0.042*** (0.013)	-0.059*** (0.012)	-0.117*** (0.018)	-0.057*** (0.021)
Number of inpatient admissions	-0.089*** (0.013)	-0.031*** (0.011)	-0.041*** (0.011)	-0.086*** (0.012)	-0.032* (0.018)
Any inpatient admission	-0.053*** (0.006)	-0.025*** (0.005)	-0.029*** (0.005)	-0.044*** (0.006)	-0.025*** (0.008)
Log (outpatient visits)	-0.030*** (0.007)	0.004 (0.005)	-0.005 (0.005)	-0.032*** (0.007)	-0.015* (0.009)
Log (primary care visits)	-0.128*** (0.010)	-0.091*** (0.007)	-0.107*** (0.009)	-0.125*** (0.009)	-0.054*** (0.013)
Log (specialist visits)	-0.020** (0.008)	0.013 (0.010)	0.007 (0.006)	-0.017* (0.009)	-0.007 (0.012)

Notes: Table combines various robustness and sensitivity checks. Column 1 presents the baseline IV results from the main paper. Column 2 and 3 show sensitivity to various definitions of big organizations based on the number of physicians in a given practice. Column 4 focuses on patients with no partial organization size affiliation. Column 5 restricts the sample to only patients whose physicians moved from small organizations.

sample period. We find that patient risk scores do not evolve differentially (prior or after the move, relative to 2008) for providers who change organizations. Relatedly, we perform a “placebo” test for patients who are in our sample throughout the entire length of study and run a reduced form regression of patient outcome in 2008–2010 on small to big PCP switches between 2012 and 2014. If these patients experience declines in utilization, then our main findings will be due to differences in pre-trends and not physician consolidation. The results of those exercise are presented in Appendix Table A8. Our main outcomes are either not statistically significant or in the other direction: those who experience their PCP move from a small to a big organization experience a utilization increase prior. This could in part reflect certain situations where physicians are more likely to switch to larger organizations following unusually busy years.

7 | CONCLUSIONS

The US healthcare system has experienced a remarkable consolidation across providers over the last 2 decades. While much of the existing research has focused on the impact of consolidation on provider market power and healthcare prices in various consolidation forms, in this paper we use data from Traditional Medicare—where prices are administratively set—and thus can largely ignore any impact on prices and market power, and focus on the impact of primary care consolidation on patient healthcare utilization.

In our preferred specification we use patient-level panel data and rely on variation in organizational-size associated with the patient that originates in her physicians moving from small to big organizations. We find a fairly large and statistically significant reduction in patient healthcare spending once she is being treated by a primary care physician who is associated with a big, rather than small organization. The effect is driven by a reduction in the number of inpatient admissions and in the frequency of primary care visits, which is consistent with the ability of larger physician groups or healthcare organizations to substitute inpatient services to (presumably cheaper) outpatient services, and by the avoidance of duplicative primary care visits.

Our results provide new evidence about the impact of primary care consolidation on patient utilization, and suggest that if the consolidation continues it may generate a non-negligible reduction in overall healthcare cost through utilization effects. These results are consistent with the motivation behind recent policy reforms for bundled payments and accountable care organizations, which are aimed to reduce the extent of duplicative treatments and wasteful spending. Future research should study the welfare and health outcome impacts of such consolidations.

ACKNOWLEDGMENT

We thank the Editor, two anonymous referees, and seminar participants at Stanford for useful comments, and acknowledge support from the Stanford Institute of Economics and Policy Research.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from Stanford Center for Population Health Sciences (PHS). Restrictions apply to the availability of these data, which were used under license for this study. Data are available Stanford PHS following their data use agreement.

ORCID

Jonathan Zhang  <https://orcid.org/0000-0003-0185-8992>

Liran Einav  <https://orcid.org/0000-0003-3349-5356>

ENDNOTES

- ¹ We chose 30 as it approximately captures the median physician; an organization size of 30 corresponds to the 57th percentile in 2008 and to the 44th percentile in 2014. As shown in Appendix Table A3, results remain qualitatively similar when we use other size thresholds.
- ² This happens for 9.4% of the physicians in our data.
- ³ CMS uses an algorithm to compute a patient's risk score based on claims from the previous 12 months. The algorithm has been updated over time, and we use a fixed formula (from 2015) to keep the risk score comparable across years.
- ⁴ Recall, to qualify as "large," the organization needs to have at least 30 providers in our data in a given year.
- ⁵ Muhlestein and Smith (2016) also study size changes in group practices and find that most of the growth happens in group sizes above 100 PCPs; however, their study period only covers 2013–2015.
- ⁶ As before, the precise details of the variable construction are available in the Appendix 1.
- ⁷ Recall that $w_{i,t}^{big}$ is a continuous variable in cases when patients see multiple primary care physicians who may work for multiple organizations; hence the interpretation of "a third fully transitioned" is an underestimate.
- ⁸ We do not directly observe M&As, so we approximate it by looking for new TINs that are formed by groups of physicians (five or more) from two or more separate old TINs.

REFERENCES

- Baker, L. C., Kate Bundorf, M., & Royalty, A. B. (2019). "The effects of multispecialty group practice on health care spending and use." NBER Working Paper No. 25915.
- Baker, L. C., Bundorf, M. K., & Kessler, D. P. (2014). Vertical integration: Hospital ownership of physician practices is associated with higher prices and spending. *Health Affairs*, 33(5), 756–763.
- Capps, C., Dranove, D., & Ody, C. (2018). The effect of hospital acquisitions of physician practices on prices and spending. *Journal of Health Economics*, 59, 139–152.
- Conti, R. M., Beth Landrum, M., & Jacobson, M. (2016). *The impact of provider consolidation on outpatient prescription drug-based cancer care spending*. Mimeo, Health Care Cost Institute. http://www.siteneutral.org/wp-content/uploads/2016/06/3_HCCI-Issue-Brief-Impact-of-Provider-Consolidation.pdf
- Cutler, D. M., & Scott Morton, F. (2013). Hospitals, market share, and consolidation. *JAMA*, 310(18), 1964–1970.
- Dranove, D., & Lindrooth, R. (2003). Hospital consolidation and costs: Another look at the evidence. *Journal of Health Economics*, 22(6), 983–997.
- Harrison, T. D. (2011). Do mergers really reduce costs? Evidence from hospitals. *Economic Inquiry*, 49(4), 1054–1069.
- Kirchhoff, S. M. (2013). *Physician practices: Background, organization, and market consolidation*. CRS Report for Congress No R42880.
- Koch, T. G., Wendling, B. W., & Wilson, N. E. (2017). How vertical integration affects the quantity and cost of care for Medicare beneficiaries. *Journal of Health Economics*, 52, 19–32.
- Kao, Yu-H., Lin, W.-T., Chen, W.-H., Wu, S.-C., & Tseng, T.-S. (2019). Continuity of outpatient care and avoidable hospitalization: A systematic review. *American Journal of Managed Care*, 25(4), 126–134.
- McWilliams, J. M., Hatfield, L. A., Chernew, M. E., Landon, B. E., & Schwartz, A. L. (2016). Early performance of accountable care organizations in Medicare. *New England Journal of Medicine*, 374(24), 2357–2366.
- Muhlestein, D. B., & Smith, N. J. (2016). Physician consolidation: Rapid movement from small to large group practices, 2013–15. *Health Affairs*, 35(9), 1638–1642.
- Neprash, H. T., Chernew, M. E., Hicks, A. L., Gibson, T., & McWilliams, J. M. (2015). Association of financial integration between physicians and hospitals with commercial health care prices. *JAMA Internal Medicine*, 175(12), 1932–1939.
- Oren, O., Kebebew, E., Ioannidis, J. P. A., & Ioannidis, A. (2019). Curbing unnecessary and wasted diagnostic imaging. *JAMA*, 321(3), 245–246.
- Phillips, J. P., Wendling, A., Bentley, A., Marsee, R., Morley, C. P., & Morley (2019). Trends in US medical School contributions to the family physician workforce: 2018 update from the American academy of family physicians. *Family Medicine*, 51(3), 241–250.
- Robinson, J. C., & Miller, K. (2014). Total expenditures per patient in hospital-owned and physician-owned physician organizations in California. *JAMA*, 312(16), 1663–1669.
- Sabety, A. H., Jena, A. B., & Barnett, M. L. (2021). Changes in health care use and outcomes after turnover in primary care. *JAMA Internal Medicine*, 181, 186. <https://doi.org/10.1001/jamainternmed.2020.6288>
- Shi L. (2012). The Impact of Primary Care: A Focused Review. *Scientifica*, 2012, 1–22. <https://doi.org/10.6064/2012/432892>

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Zhang J, Chen Y, Einav L, Levin J, Bhattacharya J. Consolidation of primary care physicians and its impact on healthcare utilization. *Health Economics*. 2021;30:1361–1373. <https://doi.org/10.1002/hec.4257>